



# Manual Set

# Getting Started

Documentation Roadmap >

Glossary >

General Description >

Release Notes >

Program Updates >

Software Release Notes >

# Installation/Setup

Server >

Client >

Server Assistant >

# DevSecOps

Application Development Guide >

Operation Guide >

Security Operation Guide >

Cluster Operation Guide >

Connection Manager User Guide >

# Reference

Glossary >

Reference Guide >

Message Guide >

FJQSS User Guide >





# Getting Started

Documentation Roadmap >

Glossary >

General Description >

Release Notes >

Program Updates >

Software Release Notes >





# Documentation Roadmap

# Preface

#### Purpose of this document

This document is intended for users of "FUJITSU Software Enterprise Postgres" (hereafter referred to as "FUJITSU Enterprise Postgres"), and explains how to read the manuals.

#### Structure of this document

The structure and content of this manual is shown below.

Chapter 1 How to Read the Manuals

This section explains the notational conventions in FUJITSU Enterprise Postgres manuals.

Chapter 2 Trademarks

This section explains the trademarks.

#### **Export restrictions**

Exportation/release of this document may require necessary procedures in accordance with the regulations of your resident country and/or US export control laws.

#### Issue date and version

Edition 1.0: July 2020

#### Copyright

Copyright 2015-2020 FUJITSU LIMITED

# Contents

Chapter 1 How to Read the Manuals	1
Chapter 1 How to Read the Manuals	1
1.2 Abbreviations of Manual Titles	1
1.3 System of Manuals and How to Use the Manuals	1
1.3.1 System of Manuals	2
1.3.2 Documentation Road Man	3
1.4 Notational Conventions in the Manuals	5
1.4.1 Platform-specific Information.	5
1.4.1 Platform-specific Information	5
1.4.3 FUJITSU Enterprise Postgres Conventions	7
1.4.3.1 Server	7
1.4.3.2 Client	7
1.4.4 Symbol Convention.	7
1.5 Notes about Manuals	7
Chapter 2 Trademarks	0

# Chapter 1 How to Read the Manuals

The FUJITSU Enterprise Postgres manuals use certain notational conventions and rules. Pay attention to these conventions and rules when reading the FUJITSU Enterprise Postgres manuals.

### 1.1 Intended Products

The manuals apply to the following products:

- FUJITSU Enterprise Postgres Advanced Edition
- FUJITSU Enterprise Postgres Standard Edition

### 1.2 Abbreviations of Manual Titles

The following tables list abbreviations of the titles of manuals for FUJITSU Enterprise Postgres as they appear in the manuals.

Formal manual title	Abbreviation in FUJITSU Enterprise Postgres manuals
FUJITSU Enterprise Postgres Release Notes	Release Notes
FUJITSU Enterprise Postgres General Description	General Description
FUJITSU Enterprise Postgres Installation and Setup Guide for Server	Installation and Setup Guide for Server
FUJITSU Enterprise Postgres Installation and Setup Guide for Client	Installation and Setup Guide for Client
FUJITSU Enterprise Postgres Installation and Setup Guide for Server Assistant	Installation and Setup Guide for Server Assistant
FUJITSU Enterprise Postgres Operation Guide	Operation Guide
FUJITSU Enterprise Postgres Cluster Operation Guide (Database Multiplexing)	Cluster Operation Guide (Database Multiplexing)
FUJITSU Enterprise Postgres Cluster Operation Guide (PRIMECLUSTER)	Cluster Operation Guide (PRIMECLUSTER)
FUJITSU Enterprise Postgres Security Operation Guide	Security Operation Guide
FUJITSU Enterprise Postgres Application Development Guide	Application Development Guide
FUJITSU Enterprise Postgres Connection Manager User's Guide	Connection Manager User's Guide
FUJITSU Enterprise Postgres Reference	Reference
FUJITSU Enterprise Postgres Java API Reference	Java API Reference
FUJITSU Enterprise Postgres Npgsql API Reference	Npgsql API Reference
FUJITSU Enterprise Postgres Glossary	Glossary
FUJITSU Enterprise Postgres Messages	Messages
PostgreSQL 12.1 Documentation	PostgreSQL Documentation

# 1.3 System of Manuals and How to Use the Manuals

This section describes the system of manuals for FUJITSU Enterprise Postgres.

# 1.3.1 System of Manuals

### **FUJITSU Enterprise Postgres manuals**

The table below shows the manuals on FUJITSU Enterprise Postgres.

Use/Purpose	Manual title	Content	When to read
Deciding whether to upgrade the product.	Release Notes	Overview of upgraded features and incompatibility information.	When learning about features upgraded from earlier versions and incompatibility information.
Acquiring an overview of the product and the basic information required for work and operation.	General Description	Description of all available functions associated with each intended purpose or use, and screenshots of operations.	When learning basic information and restrictions that system engineers and operators must know to actually operate the product.
Installing and setting up FUJITSU Enterprise Postgres correctly to enable its use.	Installation and Setup Guide for Server	Procedure for installing and setting up FUJITSU Enterprise Postgres.	When installing and setting up FUJITSU Enterprise Postgres.
Installing the FUJITSU Enterprise Postgres client function correctly to enable its use.	Installation and Setup Guide for Client	Installing the FUJITSU Enterprise Postgres client function.	When installing the FUJITSU Enterprise Postgres client function.
Installing and setting up the FUJITSU Enterprise Postgres Server Assistant.	Installation and Setup Guide for Server Assistant	Procedure for installing and setting up the FUJITSU Enterprise Postgres Server Assistant.	When installing and setting up the FUJITSU Enterprise Postgres Server Assistant.
Operating and managing FUJITSU Enterprise Postgres.	Operation Guide	Description of the tasks required in FUJITSU Enterprise Postgres management and operation.	When learning how to operate and manage the databases.
Performing switchover using database multiplexing mode.	Cluster Operation Guide (Database Multiplexing)	Description of the tasks required for database multiplexing operation.	When using database multiplexing mode to create operating environment for switchover and perform it.
Performing failover using PRIMECLUSTER.	Cluster Operation Guide (PRIMECLUSTE R)	Description of how to set up failover and perform operations using PRIMECLUSTER.	When performing failover using PRIMECLUSTER.
Performing security operation.	Security Operation Guide	Description of the tasks required for security operations.	When using security features and performing security operation in FUJITSU Enterprise Postgres.
Applications using the interface provided by FUJITSU Enterprise Postgres.	Application Development Guide	Procedure for creating an application using embedded SQL, JDBC driver, ODBC	When developing an application using the interface provided by FUJITSU Enterprise Postgres.

Use/Purpose	Manual title	Content	When to read
		driver, and .NET Data Provider.	
Performing high availability system using the Connection Manager feature.	Connection Manager User's Guide	Description of the features, setup, and usage of Connection Manager.	When using Connection Manager features by FUJITSU Enterprise Postgres.
Usage of FUJITSU Enterprise Postgres commands.	Reference	Description of the FUJITSU Enterprise Postgres commands expanded on from PostgreSQL.	When learning FUJITSU Enterprise Postgres command functions, options, and examples of use.
Learning the syntax of the JDBC API.	Java API Reference	Description of the syntax of the JDBC API.	When learning the syntax of the JDBC API.
Learning the syntax of the Npgsql API.	Npgsql API Reference	Description of the syntax of the Npgsql API.	When learning the syntax of the Npgsql API.
Learning the meaning of the terms of FUJITSU Enterprise Postgres.	Glossary	Description of the terms used in the FUJITSU Enterprise Postgres manuals.	When checking the meaning of terms used in the FUJITSU Enterprise Postgres manuals.
Referring to messages from FUJITSU Enterprise Postgres and taking measures for them.	Messages	Description of each message and description of any measures to be taken for it.	When finding out the specific measures for dealing with messages from FUJITSU Enterprise Postgres.

#### PostgreSQL manual

The table below shows the manual on PostgreSQL-compatible features.

Use/Purpose	Manual title	Content	When to read
Learning about PostgreSQL features.	PostgreSQL Documentation	Official PostgreSQL documentation.  Explains all features officially supported by the relevant version of PostgreSQL.	When learning how to use PostgreSQL.

# 1.3.2 Documentation Road Map

This section provides a documentation roadmap, broken down by user role.

#### **Database administrator**

 $The \ database \ administrator \ is \ a \ user \ who \ performs \ FUJITSU \ Enterprise \ Postgres \ installation \ and \ setup, \ and \ who \ operates \ and \ monitors \ the \ database.$ 

Refer to the manuals in the table below, according to purpose:

	Purpose	Manual name
Required reading	To learn about upgraded features and incompatibility information	Release Notes
	To read an overview of the software	General Description
	To perform installation and setup	Installation and Setup Guide for Server
		Cluster Operation Guide (Database Multiplexing)
		Cluster Operation Guide (PRIMECLUSTER)
		Connection Manager User's Guide
	To install the Server Assistant	Installation and Setup Guide for Server Assistant
	To operate and monitor	Operation Guide
		Security Operation Guide
		Cluster Operation Guide (Database Multiplexing)
		Cluster Operation Guide (PRIMECLUSTER)
		Reference
	Using Connection Manager features	Connection Manager User's Guide
	Reference	Messages
		Glossary
Refer to as required	To learn about PostgreSQL features	PostgreSQL Documentation

#### **Application developer**

The application developer is a user who defines the database and develops applications.

Refer to the manuals in the table below, according to purpose:

	Purpose	Manual name
Required reading	To learn about upgraded features and incompatibility information	Release Notes
	To read an overview of the software	General Description
	To perform installation and setup	Installation and Setup Guide for Client
	To define a database	Operation Guide
	To develop applications	Application Development Guide
		Java API Reference

	Purpose	Manual name
		Npgsql API Reference
	Using Connection Manager features	Connection Manager User's Guide
	Reference	Messages
		Glossary
Refer to as required	To learn about PostgreSQL features	PostgreSQL Documentation

### 1.4 Notational Conventions in the Manuals

Manual titles and product names in the manual are abbreviated.

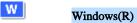
This section explains the notational conventions for abbreviations and platform-specific information in the manuals.

### 1.4.1 Platform-specific Information

Even manuals whose title has a platform name contains content common to all the platforms supported by FUJITSU Enterprise Postgres. In such cases, the platform-specific information is marked as shown below. Refer to only the necessary information.



Indicates content concerning Linux.



Indicates content concerning Windows(R).

#### 1.4.2 Abbreviation of Product Names

The following table lists abbreviations of the names of products related to FUJITSU Enterprise Postgres as they appear in the manuals.

Formal name	Abbreviation
Oracle Solaris 10 and Oracle Solaris 11	Solaris
Red Hat(R) Enterprise Linux(R) 7 Red Hat(R) Enterprise Linux(R) 8.1 and SUSE Linux Enterprise Server 12	Linux
Red Hat(R) Enterprise Linux(R) 7	RHEL7
Red Hat(R) Enterprise Linux(R) 8.1	RHEL8.1
SUSE Linux Enterprise Server 12	SLES 12
Windows(R) 8.1, Windows(R) 8.1 Pro and Windows(R) 8.1 Enterprise	Windows(R) 8.1
Windows(R) 10 Home, Windows(R) 10 Education, Windows(R) 10 Pro and Windows(R) 10 Enterprise	Windows(R) 10
Microsoft(R) Windows Server(R) 2012 Datacenter, Microsoft(R) Windows Server(R) 2012 Standard,	Windows Server(R) 2012

Formal name	Abbreviation
Microsoft(R) Windows Server(R) 2012 Essentials and Microsoft(R) Windows Server(R) 2012 Foundation	
Microsoft(R) Windows Server(R) 2012 R2 Datacenter, Microsoft(R) Windows Server(R) 2012 R2 Standard, Microsoft(R) Windows Server(R) 2012 R2 Essentials and Microsoft(R) Windows Server(R) 2012 R2 Foundation	Windows Server(R) 2012 R2
Microsoft(R) Windows Server(R) 2016 Datacenter, Microsoft(R) Windows Server(R) 2016 Standard and Microsoft(R) Windows Server(R) 2016 Essentials	Windows Server(R) 2016
Microsoft(R) Windows Server(R) 2019 Datacenter, Microsoft(R) Windows Server(R) 2019 Standard and Microsoft(R) Windows Server(R) 2019 Essentials	Windows Server(R) 2019
Windows(R) 8.1, Windows(R) 8.1 Pro, Windows(R) 8.1 Enterprise, Windows(R) 10 Home, Windows(R) 10 Education, Windows(R) 10 Enterprise, Microsoft(R) Windows Server(R) 2012 Datacenter, Microsoft(R) Windows Server(R) 2012 Standard, Microsoft(R) Windows Server(R) 2012 Essentials, Microsoft(R) Windows Server(R) 2012 Foundation, Microsoft(R) Windows Server(R) 2012 R2 Datacenter, Microsoft(R) Windows Server(R) 2012 R2 Datacenter, Microsoft(R) Windows Server(R) 2012 R2 Standard, Microsoft(R) Windows Server(R) 2012 R2 Essentials, Microsoft(R) Windows Server(R) 2012 R2 Foundation, Microsoft(R) Windows Server(R) 2016 Datacenter, Microsoft(R) Windows Server(R) 2016 Standard and Microsoft(R) Windows Server(R) 2016 Essentials Microsoft(R) Windows Server(R) 2019 Datacenter, Microsoft(R) Windows Server(R) 2019 Standard and Microsoft(R) Windows Server(R) 2019 Essentials	Windows(R)
Internet Explorer(R) 11	Internet Explorer
Microsoft(R) Edge	Edge
Java Naming and Directory Interface  Java(TM) 2 SDK, Standard Edition,  Java(TM) 2 Platform, Enterprise Edition,  Java(TM) Platform, Standard Edition and  Java(TM) Development Kit	JNDI JDK
Java(TM) 2 Runtime Environment, Standard Edition and Java(TM) Runtime Environment	JRE
Microsoft(R) Visual Basic(R) for Applications	VBA
Microsoft(R) Visual Basic	Visual Basic
Microsoft(R) Visual Basic.NET	Visual Basic.NET
Microsoft(R) Visual Studio	Visual Studio
Microsoft(R) Visual Studio.NET	Visual Studio.NET
Microsoft .NET Framework	.NET Framework or.NET

Formal name	Abbreviation
Interstage Application Server Enterprise Edition, Interstage Application Server Standard-J Edition and Interstage Web Server	Interstage Application Server
FUJITSU Enterprise Postgres Advanced Edition (64bit)	AE or FUJITSU Enterprise Postgres Advanced Edition
FUJITSU Enterprise Postgres Standard Edition (64bit)	SE or FUJITSU Enterprise Postgres Standard Edition

Remarks: The symbols (R) and (TM) may be omitted in this manual.

### 1.4.3 FUJITSU Enterprise Postgres Conventions

The naming conventions for the FUJITSU Enterprise Postgres product names and functions used in the FUJITSU Enterprise Postgres manuals are shown below.

#### 1.4.3.1 Server

The names used in the manuals in explanations regarding FUJITSU Enterprise Postgres functions are shown below.

Product name	Name used in manuals
FUJITSU Enterprise Postgres Advanced Edition (64bit) and	64-bit product
FUJITSU Enterprise Postgres Standard Edition (64bit)	

#### 1.4.3.2 Client

The names used in the manuals in explanations regarding FUJITSU Enterprise Postgres client functions are shown below.

Product name	Name used in manuals
FUJITSU Enterprise Postgres Client (64bit)	64-bit product
FUJITSU Enterprise Postgres Client (32bit)	W 32-bit product

## 1.4.4 Symbol Convention

The symbols shown below are used in the manuals.

Symbol	Meaning
[]	These symbols indicate characters displayed in a window or dialog box or keyboard keys.
	Examples: [Setting] dialogue box, [File] menu, [Item name], [OK] button, [Enter] key.

## 1.5 Notes about Manuals

This section contains notes about the FUJITSU Enterprise Postgres operating environments and manuals.

- Images in figures

The FUJITSU Enterprise Postgres manuals contain figures showing printouts for FUJITSU Enterprise Postgres to provide the reader an idea of what the printouts look like, but since the figures are only examples, they are incomplete.

#### - Explanatory examples

Most of the examples of databases in the FUJITSU Enterprise Postgres manuals are modeled after inventory control databases of retail stores. The design and contents of the databases in the examples are fictitious and do not represent any real database.

#### - UNIX release version number

This system conforms to UNIX System V Rel4.2MP.

# Chapter 2 Trademarks

- Internet Information Services, Microsoft, MS, MS-DOS, Windows, Windows Server, Visual Basic, Visual Studio and Microsoft .NET Framework are registered trademarks or trademarks of Microsoft Corporation in the U.S. and/or other countries.
- Oracle and Java are registered trademarks of Oracle Corporation and its subsidiaries and affiliated companies in the U.S. and other countries. Product and company names mentioned in this manual are the trademarks or registered trademarks of their respective owners.
- Linux(R) is a registered trademark of Linus Torvalds in the U.S. and other countries.
- Red Hat, RPM, and all Red Hat-based trademarks and logos are registered trademarks or trademarks of Red Hat, Inc. in the U.S. and other countries.
- SUSE and the SUSE logo are registered trademarks of SUSE LLC in the United States and other countries.
- UNIX is a registered trademark of Open Group in the U.S. and other countries.
- Interstage, Systemwalker, PRIMECLUSTER, and FUJITSU Software Enterprise Postgres are trademarks or registered trademarks of Fujitsu Limited.

Other product and company names mentioned in this manual are the trademarks or registered trademarks of their respective owners.





Glossary

# Preface

#### Purpose of this document

This document explains FUJITSU Enterprise Postgres terminology.

#### Intended readers

This document is aimed at all users of FUJITSU Enterprise Postgres.

#### **Export restrictions**

Exportation/release of this document may require necessary procedures in accordance with the regulations of your resident country and/or US export control laws.

#### Issue date and version

Edition 1.0: July 2020

#### Copyright

Copyright 2015-2020 FUJITSU LIMITED

# Term definitions

#### Arbitration command

A user exit (user command) called when an abnormality is detected using operating system/server heartbeat monitoring in database multiplexing mode.

#### Arbitration server

A dedicated server on which the Server Assistant program is installed.

#### Archive log

Contains the history of updates made to the database, and is used during recovery.

#### Backup data storage destination

The directory that stores the backup data.

#### Client command

A command that is executed from the client machine and used. Also known as a client application.

#### Connection Manager

The replication operation to continue without knowing where the application is connected. The Connection Manager feature improves the availability.

#### Data storage destination

The directory that stores the database clusters.

#### Database cluster

The database storage area on the database storage disk. Database clusters are a collection of databases managed by an instance.

#### Data masking

A feature that can change the returned data for queries generated by applications, to prevent exposing actual data.

#### Database multiplexing

Mechanism in which a database is made redundant on multiple servers, by transferring transaction logs (WAL) via the network to enable application jobs to be continued.

#### Database superuser

A user defined in the database with access privileges for all database objects.

#### Encoding

Indicates the character set.

#### Fencing

A process that isolates a database server with an unstable status from the cluster system in database multiplexing mode. This process is implemented as a fencing command.

#### Fencing command

A user exit (user command) that implements fencing in database multiplexing mode.

#### Global Meta Cache

The Global Meta Cache feature cache the informations about system catalogs and table definitions information (meta cache) in shared memory. The meta cache on shared memory is called the Global Meta Cache (GMC).

#### Instance

A series of server processes for managing database clusters.

#### Instance administrator

The OS user account that owns the database cluster files and operates the database server processes.

#### Instance name

Indicates the instance name.

#### Masking policy

A method of changing data under specific conditions when it is returned for a query from an application. You can configure masking target, masking type, masking condition and masking format.

#### Mirrored transaction log

The log that mirrors the transaction log at the backup data storage destination.

#### Mirroring Controller arbitration process

A process that performs arbitration and fencing on the arbitration server.

#### Mirroring Controller monitoring process

A process that performs heartbeat monitoring of the Mirroring Controller process. If the Mirroring Controller process returns no response or is down, the Mirroring Controller monitoring process is restarted automatically.

#### Mirroring Controller process

A process that performs operating system/server and process heartbeat monitoring and disk abnormality monitoring between database servers. Additionally, the process issues arbitration requests to the arbitration server and executes arbitration commands.

#### Papool-II connection pooling

The connection pooling feature of Pgpool-II supported by FUJITSU Enterprise Postgres.

This feature maintains the connection established with the database server and reuses that connection each time a new connection with the same properties (user name, database, and protocol version) arrives. By reducing the connection overhead for the database server, throughput of the whole system is improved.

#### Pgpool-II failover

The automatic failover feature of Pgpool-II supported by FUJITSU Enterprise Postgres.

If any of the database servers crashes or can no longer be reached, this feature disconnects the server and continues operation on the remaining servers. The streaming replication feature of PostgreSQL is combined with Pgpool-II to achieve a high-availability system.

#### Pgpool-II load balancing

The load balancing feature of Pgpool-II supported by FUJITSU Enterprise Postgres.

This feature distributes reference queries to multiple database servers, improving throughput of the whole system. The database multiplexing feature or PostgreSQL streaming replication feature is combined with Pgpool-II to reduce the load on the database server.

#### Pgpool-II server

A server for using the failover, connection pooling, and load balancing features of Pgpool-II. It is a dedicated server that has a server program installed for using these features.

#### Primary server

The server that processes the main database jobs during multiplexed database operation.

#### Server Assistant

A feature that objectively determines the status of database servers as a third party, and if necessary, isolates affected databases if the database servers are unable to accurately ascertain their mutual statuses in database multiplexing mode, such as due to a network error between database servers, or server instability.

#### Server Assistant program

A program to be installed on the arbitration server.

#### Server command

A command used on the database server. Also known as a server application.

#### Standby server

A server that generates a replicated database synchronized with the primary server, and that can run as an alternative server in case the primary server fails during multiplexed database operation.

#### State transition command

A user exit (user command) called when Mirroring Controller performs a state transition of a database server in database multiplexing mode. State transition commands include the post-switch command, pre-detach command, and post-attach command.

#### Transaction log

Contains the history of updates made to the database by transactions. Also known as the WAL (Write-Ahead Log).

#### Transaction log storage destination

The directory that stores the transaction log.

#### VCI (Vertical Clustered Index)

An index with columnar data structure suitable for aggregation.

#### WAL (Write-Ahead Log)

Has the same meaning as 'transaction log'.

#### WebAdmin program

A GUI-based program installed on a database server or a dedicated WebAdmin server, used to manage database instances.

#### WebAdmin server

By using the WebAdmin program on a different server to the database server, instances on multiple database servers can be managed from a dedicated WebAdmin server on which the WebAdmin program is installed.





# General Description

# **Preface**

### Purpose of this document

This document explains the FUJITSU Enterprise Postgres concepts to those who are to operate databases using it. This document explains the features of FUJITSU Enterprise Postgres.

#### Intended readers

This document is intended for people who are:

- Considering installing FUJITSU Enterprise Postgres
- Using FUJITSU Enterprise Postgres for the first time
- Wanting to learn about the concept of FUJITSU Enterprise Postgres
- Wanting to see a functional overview of FUJITSU Enterprise Postgres

Readers of this document are also assumed to have general knowledge of:

- Computers
- Jobs
- Linux
- Windows(R)

#### Structure of this document

This document is structured as follows:

Chapter 1 FUJITSU Enterprise Postgres Basics

Explains the features of FUJITSU Enterprise Postgres.

Appendix A List of Features

Lists the main features provided by FUJITSU Enterprise Postgres.

Appendix B OSS Supported by FUJITSU Enterprise Postgres

Explains the OSS supported by FUJITSU Enterprise Postgres.

Appendix C Features that can be Used on Servers Other than the Database Server

Explains features that can be used on servers other than the database server.

#### **Export restrictions**

Exportation/release of this document may require necessary procedures in accordance with the regulations of your resident country and/or US export control laws.

#### Issue date and version

Edition 1.0: July 2020

#### Copyright

Copyright 2015-2020 FUJITSU LIMITED

# Contents

Chapter 1 FUJITSU Enterprise Postgres Basics	1
1.1 Flexible Database Recovery	
1.2 Simple GUI-Based Installation and Operation Management	
1.3 High Reliability with Database Multiplexing	
1.4 High Reliability Using Failover Integrated with the Cluster Software	
1.5 Seamless Migration from Oracle Databases	
1.6 Linkage with Integrated Development Environment	
1.7 Storage Data Protection Using Transparent Data Encryption	
1.8 Data Masking for Improved Security	
1.9 Security Enhancement Using Audit Logs	
1.10 Enhanced Query Plan Stability	
1.11 Increased Aggregation Performance Using the In-memory Feature	
1.12 High-Speed Data Load	
1.13 High availability by using Connection Manager	
1.14 Memory usage reduction by Global Meta Cache	
Appendix A List of Features	15
Appendix B OSS Supported by FUJITSU Enterprise Postgres	16
Appendix C Features that can be Used on Servers Other than the Database Server	18
C.1 WebAdmin.	
C.2 Server Assistant.	18
C.3 Failover, Connection Pooling, and Load Balancing Features of Pgpool-II.	18
to describe the second	0.4

# Chapter 1 FUJITSU Enterprise Postgres Basics

FUJITSU Enterprise Postgres maintains the operating methods, interfaces for application development and SQL compatibility of PostgreSQL, while providing expanded features for enhanced reliability and operability.

This chapter explains the functionality extended by FUJITSU Enterprise Postgres.

Refer to "Appendix A List of Features" for feature differences between editions.

Additionally, FUJITSU Enterprise Postgres supports various open source software (OSS). Refer to "Appendix B OSS Supported by FUJITSU Enterprise Postgres" for information on OSS supported by FUJITSU Enterprise Postgres.

FUJITSU Enterprise Postgres has the following features:

- Flexible database recovery

Not only does FUJITSU Enterprise Postgres recover data to its most recent form when a failure occurs, which is essential for databases, but it can also recover to any point in time. Additionally, backup/recovery can be performed using any copy technology.

- Simple GUI-based installation and operation management
   FUJITSU Enterprise Postgres uses GUI to simplify cumbersome database operations, and allows databases to be used intuitively.
- High reliability by using database multiplexing

  Database multiplexing protects important data and enables highly reliable database operation.
- High reliability by using the failover feature integrated with the cluster software FUJITSU Enterprise Postgres links with PRIMECLUSTER, thereby allowing highly reliable systems to be achieved by using failover.
- Seamless migration from Oracle databases
   FUJITSU Enterprise Postgres provides a compatibility feature with Oracle databases that localizes the correction of existing applications and allows easy migration to FUJITSU Enterprise Postgres.
- Linkage with integrated development environment
   FUJITSU Enterprise Postgres links with Visual Studio, thereby allowing a standard framework to be used to create applications.
- Storage Data Protection using Transparent Data Encryption
   Information can be protected from data theft by encrypting data to be stored in the database.
- Data masking for improved security
   The data masking feature changes the returned data for queries from applications, to prevent exposing actual data. This improves security for handling confidential data such as personal information.
- Audit logs for improved security

  Audit logs can be used to counter security threats such as unauthorized access and misuse of privileges for the database.

- Enhanced query plan stability

The following features can control SQL statement query plans:

- Optimizer hints
- Locked statistics

These features are used for curbing performance deterioration caused by changes in SQL statement query plans, such as with mission-critical jobs that emphasize performance stability over improved SQL statement processing performance.

- Increased aggregation performance using the in-memory feature

  The following features help speed up scans even when aggregating many rows.
  - Vertical Clustered Index (VCI)
  - In-memory data
- High-speed data load

Data from files can be loaded at high speed into FUJITSU Enterprise Postgres tables using the high-speed data load feature.

- High availability by using Connection Manager
   With the Connection Manager features, replication operation can be continued without being aware of the connection destination of the applications.
- Memory usage reduction by Global Meta Cache
   The Global Meta Cache feature loads some of meta cahe information in shared memory. This reduces overall system memory usage.

## 1.1 Flexible Database Recovery

Threats such as data corruption due to disk failure and incorrect operations are unavoidable in systems that use databases. The ability to reliably recover corrupted databases without extensive damage to users when such problems occur is an essential requirement in database systems.

FUJITSU Enterprise Postgres provides the following recovery features that flexibly respond to this requirement:

- Media recovery, which recovers up to the most recent point in time
- Point-in-time recovery, which can recover up to a specific point in time
- Backup/recovery that can integrate with various copy technologies

#### Media recovery, which recovers up to the most recent point in time

When a disk failure occurs, media recovery can recover data to how it was immediately before the failure.

In order to recover the database, FUJITSU Enterprise Postgres accumulates a history of database update operations, such as data additions and deletions, as an update log.

FUJITSU Enterprise Postgres retains a duplicate (mirror image) of the update log after backup execution on the data storage destination and on the backup data storage destination. Therefore, the data on one disk can be used to recover to the most recent state of the database even if a disk failure has occurred on the other.

Media recovery is executed using either a GUI tool provided with FUJITSU Enterprise Postgres (WebAdmin) or server commands.



Recovery using WebAdmin requires less time and effort, since WebAdmin automatically determines the scope of the operation.

Data storage destination

Update log

Same data

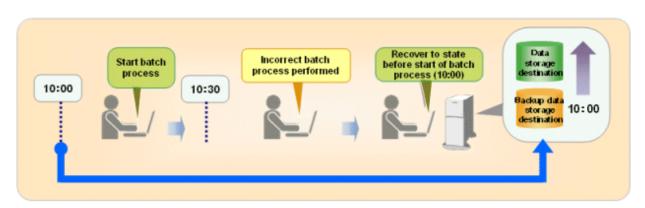
Database data

Can recover to most recent state from either disk

#### Point-in-time recovery, which can recover up to a specific point in time

Point-in-time recovery can be used to recover a database that has been updated by an incorrect operation, for example, by specifying any date and time before the incorrect operation.

Point-in-time recovery is executed using FUJITSU Enterprise Postgres server commands.



#### Backup/recovery that can integrate with various copy technologies

It is possible to back up to the backup data storage destination, or to any backup destination using any copy technology implemented by user exits.

For example, by using the high-speed copy feature of the storage device, the processing time for backup of large databases can be greatly reduced.



See

Refer to "Backup/Recovery Using the Copy Command" in the Operation Guide for information on backup/recovery using user exits.

# 1.2 Simple GUI-Based Installation and Operation Management

FUJITSU Enterprise Postgres provides WebAdmin, which is a GUI tool for a range of tasks, from database installation to operation management. This allows the databases to be used simply and intuitively.

WebAdmin can be used for FUJITSU Enterprise Postgres setup, creating and monitoring a streaming replication cluster, database backups, and for recovery. Depending on the configuration, WebAdmin can be used to manage FUJITSU Enterprise Postgres instances in a single server, or instances spread across multiple servers.

- Setup

To perform setup using WebAdmin, you must create an instance. An instance is a set of server processes that manage a database cluster (database storage area on the data storage destination disk). Instances can be created easily and with only minimal required input, because the tool automatically determines the optimal settings for operation.

- Database backup/recovery

Database backup and recovery can be performed using simple GUI operations.

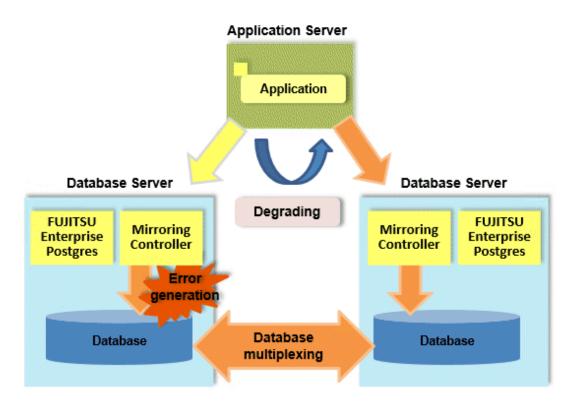
In particular, FUJITSU Enterprise Postgres can automatically identify and isolate the location of errors. This simplifies the recovery process and enables faster recovery.

In addition, FUJITSU Enterprise Postgres provides the following expanded features in pgAdmin:

- NCHAR type
- Expanded trigger definition
  - REPLACE feature
  - Function call feature

# 1.3 High Reliability with Database Multiplexing

It is vital for systems that use databases to protect data from damage or loss caused by a range of factors such as hardware and software errors. Database multiplexing protects important data and enables highly reliable database operation.



FUJITSU Enterprise Postgres not only mirrors a database using the PostgreSQL streaming replication feature, but also provides simplified switchover and standby disconnection features as well as a feature to detect faults in elements that are essential for the continuity of database process, disk, network, and other database operations.

Even if a switchover is performed, the client automatically distinguishes between the primary and standby servers, so applications can be connected transparently regardless of the physical server.

The Mirroring Controller option enables the primary server (the database server used for the main jobs) to be switched automatically to the standby server if an error occurs in the former.

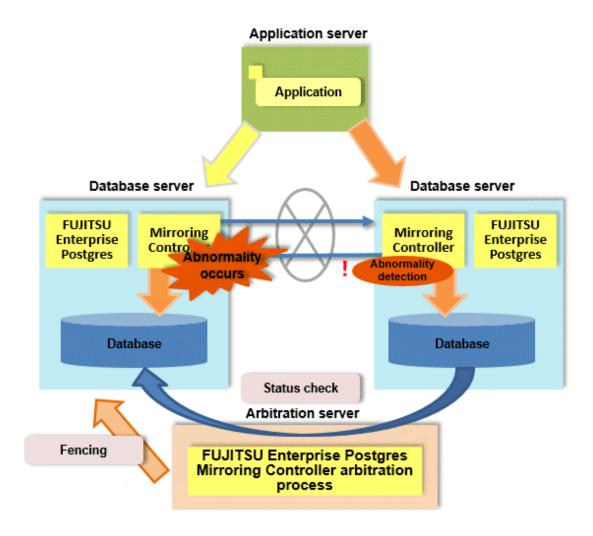
In addition, by using the data on the standby server, reference jobs such as data analysis and form output can be performed in parallel to the jobs on the primary server.

#### Operation using the arbitration server

Mirroring Controller may not be able to correctly determine the status of the other server if there is a network issue between database servers or a server is in an unstable state. As a result, both servers will temporarily operate as primary servers, so it may be possible to perform updates from either server.

The Server Assistant is a feature that objectively checks the status of database servers as a third party and isolates (fences) unstable servers in such cases.

In database multiplexing mode, the Server Assistant is made available by adding a new server (arbitration server) on which the Server Assistant is installed. Using an arbitration server can prevent the issue mentioned above (both servers temporarily operating as primary servers) and enables highly reliable operation.





L

- Refer to "Database Multiplexing Mode" in the Cluster Operation Guide (Database Multiplexing) for information on the database multiplexing.

# 1.4 High Reliability Using Failover Integrated with the Cluster Software

If the system stops, services are interrupted until recovery is complete. In large-scale systems, the interruption takes longer, and may cause significant disruption for many people receiving the services.

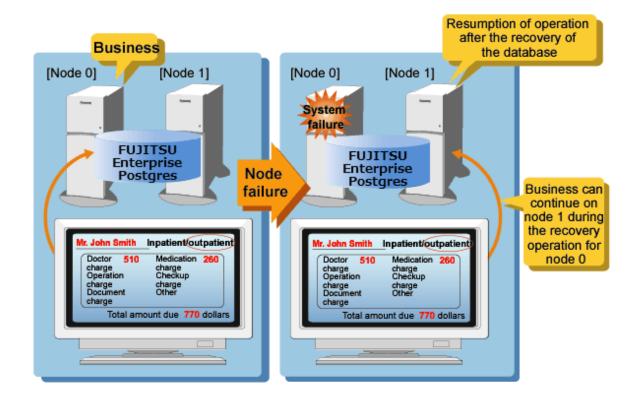
In FUJITSU Enterprise Postgres, the failover feature integrated with the cluster software can minimize the system stoppage time when an issue occurs.

#### Medical accounting system

Some hospitals with a large number of patients manage and operate the various data required for accounting in a database. If this accounting system stops and takes several minutes to recover, it is expected that this will have a significant impact.

But if failover is applied to this kind of system and an issue occurs on the operating server, it is quickly switched and the standby server takes over operation, so that services are provided without interruption.

The example below illustrates a medical accounting system using failover.





Refer to the Cluster Operation Guide (PRIMECLUSTER) for information on the failover feature integrated with the cluster software.

# 1.5 Seamless Migration from Oracle Databases

FUJITSU Enterprise Postgres supports Orafce, to provide compatibility with Oracle databases.

Using the compatibility feature reduces the cost of correcting existing applications and results in easy database migration.



See

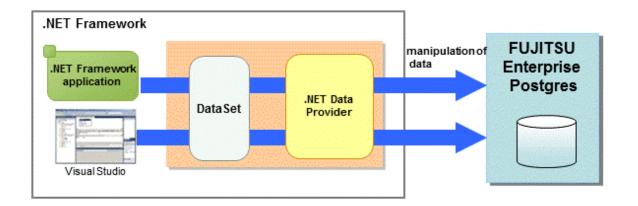
Refer to "Compatibility with Oracle Databases" in the Application Development Guide for information on compatible features.

## 1.6 Linkage with Integrated Development Environment

You can link with Microsoft Visual Studio to create application. And, you can automatically generate applications to access database resources by linking to Visual Studio.

#### Relationship between .NET Framework and FUJITSU Enterprise Postgres

FUJITSU Enterprise Postgres provides .NET Data Provider, which is an interface for ADO.NET of .NET Framework. This enables you to select FUJITSU Enterprise Postgres as the connection destination database of ADO.NET and use the intuitive and efficient application development features of Visual Studio.

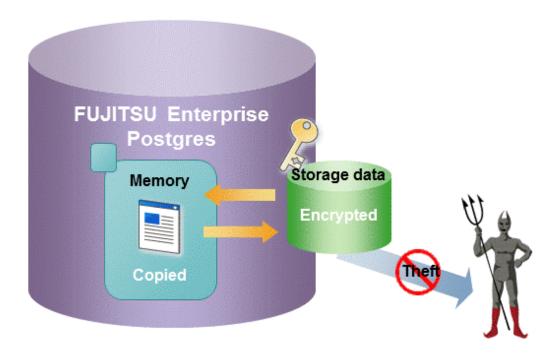


# 1.7 Storage Data Protection Using Transparent Data Encryption

The encryption of data to be stored in a database is essential under the following encryption requirements of PCI DSS (Payment Card Industry Data Security Standard), the data security standard of the credit industry:

- Confidential information (such as credit card numbers) can be encrypted.
- The encryption key and data are managed as separate entities.
- The encryption key is replaced at regular intervals.

To satisfy these requirements, FUJITSU Enterprise Postgres provides a transparent data encryption feature. Note that PostgreSQL uses an encryption feature called pgcrypto, which can also be used in FUJITSU Enterprise Postgres, but requires applications to be modified. Therefore, we recommend using FUJITSU Enterprise Postgres's transparent data encryption feature.





Refer to "Protecting Storage Data Using Transparent Data Encryption" in the Operation Guide for information on stored data encryption.

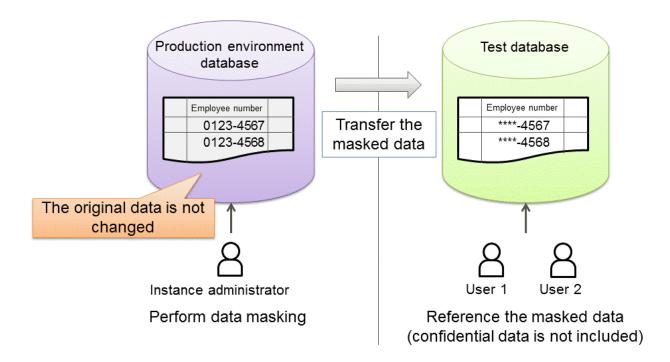
## 1.8 Data Masking for Improved Security

FUJITSU Enterprise Postgres provides a data masking feature that protects data to maintain security of data handled in systems.

The data masking feature changes the returned data for queries from applications and makes it available for reference without exposing the actual data.

For example, for a query of employee data, digits except the last four digits of an eight-digit employee number can be changed to "\*" so that it can be used for reference.

Also, the data changed by the data masking feature can be transferred to a test database so that users who perform testing or development can reference the data. As production data should not be used in a test or development environment because of the risk of data leakage, this feature enables data that is similar to actual production data to be safely used in those environments.





Refer to "Data Masking" in the Operation Guide for information on data masking.

## 1.9 Security Enhancement Using Audit Logs

Details relating to database access can be retrieved in audit logs. The audit log feature can be used to counter security threats such as unauthorized access to the database and misuse of privileges.

In PostgreSQL, logs output as server logs can be used as audit logs by using the log output feature. There are, however, logs that cannot be analyzed properly, such as SQL runtime logs, which do not output the schema name. Additionally, because the output conditions cannot be specified in detail, log volumes can be large, which may lead to deterioration in performance.

The audit log feature of FUJITSU Enterprise Postgres enables retrieval of details relating to database access as an audit log by extending the feature to pgaudit. Additionally, audit logs can be output to a dedicated log file or server log. This enables efficient and accurate log monitoring.



Refer to "Audit Log Feature" in the Security Operation Guide for details.

## 1.10 Enhanced Query Plan Stability

FUJITSU Enterprise Postgres estimates the cost of query plans based on SQL statements and database statistical information, and selects the least expensive query plan. However, like other databases, FUJITSU Enterprise Postgres does not necessarily select the most suitable query plan. For example, it may suddenly select unsuitable query plan due to changes in the data conditions.

In mission-critical systems, stable performance is more important than improved performance, and changes in query plans case to be avoided. In this situation, the following features can stabilize query plans:

- Optimizer hints

You can use pg\_hint\_plan to specify a query plan in each individual SQL statement.

- Locked statistics

You can use pg\_dmbs\_stats to lock statistical information per object, such as a database, schema, or table.



#### See

Refer to "Optimizer Hints" in the Application Development Guide for information on optimizer hints.

Refer to "Locked Statistics" in the Application Development Guide or information on locked statistics.



Use the features provided when FUJITSU Enterprise Postgres is installed for optimizer hints and locked statistical information. FUJITSU Enterprise Postgres does not support other similar open-source features.

# 1.11 Increased Aggregation Performance Using the Inmemory Feature

FUJITSU Enterprise Postgres provides the in-memory feature, which uses columnar index and memory-resident data. This reduces disk I/Os and enhances aggregation performance.



This feature can only be used in Advanced Edition.

#### Columnar index

Many aggregation processes may require a large portion of data in a particular column. However, traditional row data structure reads unnecessary columns, resulting in inefficient use of memory and CPU cache, and slower processing. FUJITSU Enterprise Postgres provides a type of columnar index, VCI (Vertical Clustered Index). This addresses the above issues, and enhances aggregation performance.

VCI provides the following benefits:

- Minimizes impact on existing jobs, and can perform aggregation using job data in real time.
- Provided as an index, so no application modification is required.

- Stores data also on the disk, so aggregation jobs can be quickly resumed using a VCI even if a failure occurs (when an instance is restarted).
- If the amount of memory used by VCI exceeds the set value, aggregation can still continue by using VCI data on the disk.

It also provides the features below:

- Disk compression

  Compresses VCI data on the disk, minimizing required disk space. Even if disk access is required, read overhead is low.
- Parallel scan
   Enhances aggregation performance by distributing aggregation processes to multiple CPU cores and then processing them in parallel.

#### In-memory data

The following features keep VCI data in memory and minimize disk I/Os on each aggregation process.

- Preload feature

  Ensures stable response times by loading VCI data to memory before an application scans it after the instance is restarted.
- Stable buffer feature
   Reduces disk I/Os by suppressing VCI data eviction from memory by other job data.

#### Purposes of this feature

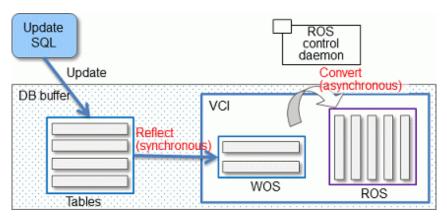
This feature has a data structure that can efficiently use the newly added resources, and aims to enhance the existing aggregation processing in normal operations to be faster than parallel scan. It shares the same purpose of enhancing aggregation performance with the parallel scan feature that is provided separately, but differs in that it speeds up nightly batch processes by utilizing available resources.

#### VCI architecture

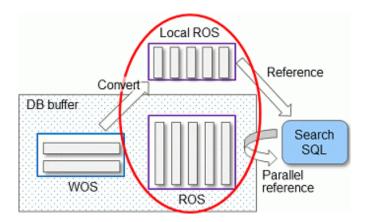
This section briefly explains VCI architecture as it contains basic terminology required, for example, when setting parameters.

Update and aggregation operations to enable real time use of job data are described.

VCI has write buffer row-based WOS (Write Optimized Store) in addition to the columnar data structure ROS (Read Optimized Store). Converting each update into a columnar index has a significant impact on the update process response times. Therefore, data is synchronously reflected to the row-based WOS when updating. After a certain amount of data is stored in WOS, the ROS control daemon asynchronously converts it to ROS. As above, the entire VCI is synchronized with the target table column, minimizing update overhead.



The same scan results can be obtained without a VCI by using WOS in conjunction with ROS. More specifically, WOS is converted to Local ROS in local memory for each aggregation process, and aggregated with ROS.





Refer to "Installing and Operating the In-memory Feature" in the Operation Guide for information on installation and operation of VCI.

Refer to "Scan Using a Vertical Clustered Index (VCI)" in the Application Development Guide for information on scan using a VCI.

# 1.12 High-Speed Data Load

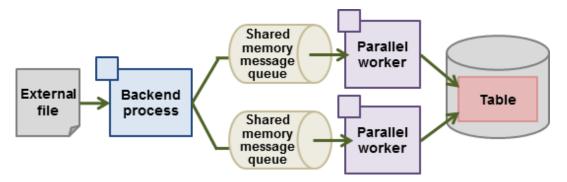
High-speed data load executes COPY FROM commands using multiple parallel workers. Because conversion of data from the external file to the appropriate internal format, table creation, and index creation are performed in parallel, it is possible to load large volumes of data at high speed.



This feature is available only in the Advanced Edition.

#### Architecture of high-speed data load

High-speed data load is required for parameter setting and resource estimation, so a brief description of its architecture is provided below.



High-speed data load uses a single backend process collaborating with multiple parallel workers to perform data load in parallel. Data is exchanged between the backend process and parallel workers via shared memory message queues. The backend process distributes the loaded data of external files to multiple parallel workers. Each parallel worker then converts the data loaded from the shared memory message queue into the appropriate internal format, and inserts it into the table. If the table has indexes, their keys are extracted and inserted into the index page.



Refer to "High-Speed Data Load" in the Operation Guide for details.



## 1.13 High availability by using Connection Manager

The Connection Manager provides the following features. You can use these features to increase system availability.

#### Heartbeat monitoring feature

Detects kernel panics between the server running the client and the server running the instance, physical server failures, and inter-server network link downs, and notifies the client or instance. The client is notified as an error event through the SQL connection, and the instance will be notified in the form of a force collection of SQL connections with clients that are out of service.

#### Transparent connection support feature

When an application wants to connect to an instance of an attribute (Primary/Standby) configured with replication, it can do so without knowing which server the instance is running on.



The available client drivers for Connection Manager are libpq (C language library) and ECPG (embedded SQL in C).



See

Refer to the Connection Manager User's Guide for details.

## 1.14 Memory usage reduction by Global Meta Cache

The Global Meta Cache feature cache the informations about system catalogs and table definitions information (It is called a meta cache.) in shared memory. The meta cache on shared memory is called the Global Meta Cache (GMC).

Without this feature, the meta cache was cached in per-process memory. Therefore, there was a problem increase in memory usage in environments with large databases and large numbers of connections. The Global Meta Cache feature enables sharing of meta caches on shared memory, thereby reducing overall system memory usage.



This feature is available only in the Advanced Edition.

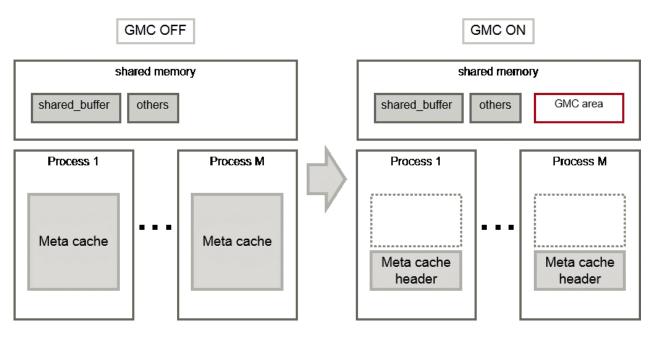
#### Meta cache

Processing a query involves parsing the query, creating the plan, executing the plan, and so on. PostgreSQL process accesses the system catalog to perform these steps. Once accessed, the system catalog tuples are cached in per-process memory. PostgreSQL also aggregates and caches the working information required to access tables and the system catalog information about those tables. The information cached in these memories is called a meta cache. Each process performs faster query processing by searching the meta cache instead of searching for the required tuples in the system catalog each time.

The meta cache usage increases in proportion to the number of tables and columns accessed. It is cached on a per-process basis, so the system's overall meta cache usage increases in proportion to the number of connections.

#### **Architecture of Global Meta Cache Feature**

Describes the architecture of the Global Meta Cache feature.



When the GMC feature is on, the per-process meta cache is cached in the GMC area on shared memory. Reference to the GMC area and process-specific work information is cached in the memory of each process. PostgreSQL process searches the meta cache for each process and accesses the GMC based on the reference information. If there is no reference information in the process's memory, it searches the GMC area. If the GMC area also does not have a corresponding meta cache, it accesses the system catalog to create meta cache.

Also, sharing the meta cache does not cause any loss of data consistency. If the system catalog or table definition changes while a transaction is running, the cache deletion or creation does not affect outside of the process running the transaction. After the transaction commits, the GMC area cache is deleted or created. If other transactions are referencing the cache when GMC is tried to be deleted, the deletion is deferred until there are no more references. After a commit, a new transaction sees the new cache instead of the old one.



Global Meta Cache feature is disabled by default. Refer to "Global Meta Cache" in the Operation Guide for information how to decide whether introduce it or not and usage.

## Appendix A List of Features

The following table lists the main features provided by FUJITSU Enterprise Postgres.

		Lir	Linux		Windows	
Category	Feature	AE	SE	AE	SE	
Fujitsu-developed software technology	WebAdmin (Rapid setup, One-click recovery)	Y	Y	Y	Y	
Improved	Database multiplexing	Y	N	Y	N	
reliability and availability	Failover (integration with PRIMECLUSTER)	Y (*1)	Y (*1)	N	N	
	Backup/recovery using user exits	Y	N	Y	N	
	Connection Manager	Y	N	N	N	
Application development	Embedded SQL integration Java integration ODBC integration .NET Framework integration	Y	Y	Y	Y	
	Features compatible with Oracle databases	Y	Y	Y	Y	
Security	Storage data encryption	Y	Y	Y	Y	
	Data masking	Y	Y	Y	Y	
	Audit log	Y	N	Y	N	
Performance	In-memory feature	Y	N	Y	N	
	High-speed data load	Y	N	Y	N	
	Global Meta Cache		N	Y	N	
Performance	Optimizer hints	Y	Y	Y	Y	
tuning	Fixed statistical information	Y	Y	Y	Y	

Y: Provided

N: Not provided

\*1: Supported on RHEL only.

# Appendix B OSS Supported by FUJITSU Enterprise Postgres

The OSS supported by FUJITSU Enterprise Postgres is listed below.

OSS name	Version and	Pla	atform	Description	Deference
OSS name	level	Linux	Windows	Description	Reference
PostgreSQL	12.1	Y	Y	Database management system	PostgreSQL Documentation
orafce	3.8.0	Y	Y	Oracle-compatible SQL features	"Compatibility with Oracle Databases" in the Application Development Guide
Pgpool-II	4.1	Y	N	Failover, connection pooling, load balancing, etc.	"Pgpool-II" in the Installation and Setup Guide for Server
oracle_fdw	2.2.0	Y	Y	Connection to the Oracle database server	"oracle_fdw" in the Installation and Setup Guide for Server
pg_statsinfo	12.0	Y	N	Collection and accumulation of statistics	"pg_statsinfo" in the Installation and Setup Guide for Server
pg_hint_plan	12.1.3.5	Y	Y	Tuning (statistics management, query tuning)	- "pg_hint_plan" in the Installation and Setup Guide for Server
					"Optimizer Hints" in the     Application     Development Guide
pg_dbms_stats	1.3.11	Y	Y		- "pg_dbms_stats" in the Installation and Setup Guide for Server
					"Locked Statistics" in the     Application     Development Guide
pg_repack	1.4.5	Y	N	Table reorganization	"pg_repack" in the Installation and Setup Guide for Server
pg_rman	1.3.9	Y	N	Backup and restore management	"pg_rman" in the Installation and Setup Guide for Server
pgAdmin4	4.15	N	Y	Operation and development GUI	Operation Guide
pgBadger	11.1	Y	N	Log analysis	"pgBadger" in the Installation and Setup Guide for Server
pg_bigm	1.2	Y	N	Full-text search (multibyte)	"pg_bigm" in the Installation and Setup Guide for Server
PostgreSQL JDBC driver	42.2.8	Y	Y	JDBC driver	"JDBC Driver" in the Application Development Guide

000 nama	OSS name Version and Platform Description		Description	Reference	
OSS name	level	Linux	Windows	Description	Reference
psqlODBC	12.01.0000	Y	Y	ODBC driver	"ODBC Driver" in the Application Development Guide
Npgsql	4.0.10	N	Y	.NET data provider	".NET Data Provider" in the Application Development Guide

Y: Supported

N: Not supported

# Appendix C Features that can be Used on Servers Other than the Database Server

This chapter explains the configuration and operating environment of features to be installed and used on servers other than the database server when used in conjunction with the FUJITSU Enterprise Postgres database server.

In this chapter, FUJITSU Enterprise Postgres programs are referred to as server programs.

Below are features to be installed and used on servers other than the database server:

- WebAdmin
- Server Assistant
- Pgpool-II (failover, connection pooling, and load balancing)

#### C.1 WebAdmin

If there is only one database server, WebAdmin is normally installed on the same server as the database (the WebAdmin program can be installed at the same time as the server program).

If there are multiple database servers, database server instances can be managed collectively if a dedicated WebAdmin server is used. In this case, the WebAdmin program is installed on the WebAdmin server, and the server program and WebAdmin program are installed on the database server.



- Refer to "1.2 Simple GUI-Based Installation and Operation Management" for information on WebAdmin.
- Refer to "Determining the Preferred WebAdmin Configuration" in the Installation and Setup Guide for Server for information on the server configuration when using WebAdmin.
- Refer to "Required Operating System" in the Installation and Setup Guide for Server for information on the operating environment of WebAdmin.

#### C.2 Server Assistant

To use the Server Assistant, the Server Assistant program is installed on a dedicated server (arbitration server).



#### See

- Refer to "Overview of Database Multiplexing Mode" in the Cluster Operation Guide (Database Multiplexing) for information on the Server Assistant and the server configuration.
- Refer to "Required Operating System" in the Installation and Setup Guide for Server Assistant for information on the operating environment of the Server Assistant.

## C.3 Failover, Connection Pooling, and Load Balancing Features of Pgpool-II

Pgpool-II is software that is placed between the database server and database client to relay the connection.

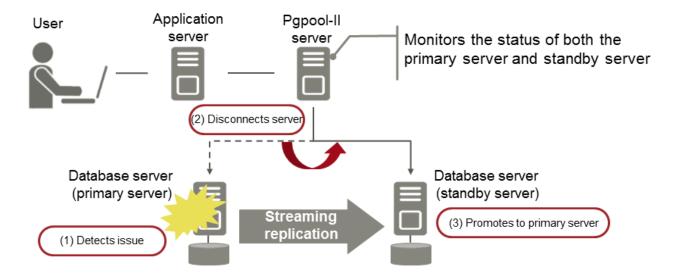
Pgpool-II provides the failover, connection pooling, and load balancing features for use during streaming replication.

#### **Failover**

In PostgreSQL, a database can be made redundant (building a high availability system) using synchronous streaming replication.

If the database server of either the primary server or standby server fails or is no longer accessible when using synchronous streaming replication, jobs will stop.

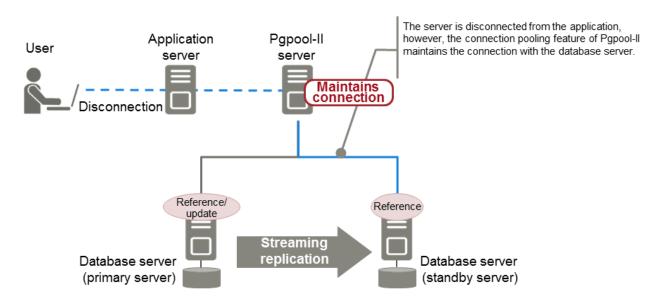
Failover monitors the status of each database and automatically disconnects the server when an error occurs. As a result, jobs can continue uninterrupted on the remaining server.



#### **Connection pooling**

This feature maintains (pools) the connection established with the database server, and reuses that connection each time a new connection with the same properties (user name, database, and protocol version) arrives.

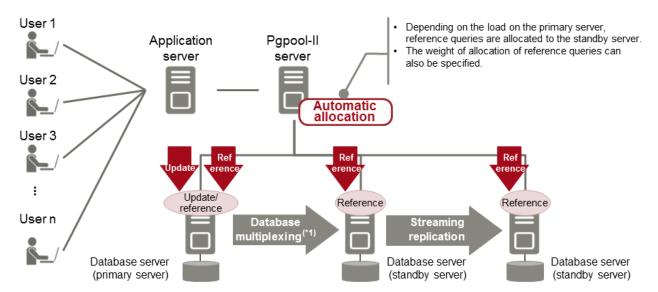
Connection pooling reduces the connection overhead for the database server, improving throughput of the whole system.



#### Load balancing

This feature distributes reference queries to multiple database servers, improving throughput of the whole system.

By combining load balancing with the FUJITSU Enterprise Postgres database multiplexing feature or the PostgreSQL streaming replication feature, load on the database server is reduced.



\*1: The arbitration server used during database multiplexing has been omitted from this document.



- Refer to "System configuration when using Pgpool-II" in the Installation and Setup Guide for Server for information on the server configuration when using Pgpool-II.
- Refer to "Required Operating System" in the Installation and Setup Guide for Server for information on the operating environment of Pgpool-II.

## Index

[Special characters]
.NET Framework7
[C]  Columnar index
[D] Database Multiplexing
[F] Flexible Database Recovery2
[G] Global Meta Cache13
[H] High-Speed Data Load12
[I] In-memory data11
[L] Linkage with Integrated Development Environment7
[M] Media recovery2
[O] Oracle Database7
[P] Point-in-time recovery3
[S] Security Enhancement Using Audit Logs9
[T] Transparent Data Encryption 8



# FUJITSU Enterprise Postgres 12



# Release Notes

Linux > Windows >



# FUJITSU Enterprise Postgres 12



Release Notes (Linux)

## Preface

#### Purpose of this document

This document provides release information for FUJITSU Enterprise Postgres.

#### Structure of this document

This document is structured as follows:

Chapter 1 New Features and Improvements

Explains the new features and improvements in this version.

Chapter 2 Compatibility Information

Provides information regarding compatibility.

Chapter 3 Program Updates

Explains updates incorporated in this version.

#### **Export restrictions**

Exportation/release of this document may require necessary procedures in accordance with the regulations of your resident country and/or US export control laws.

#### Issue date and version

Edition 1.0: July 2020

#### Copyright

Copyright 2015-2020 FUJITSU LIMITED

## **Contents**

Chapter 1 New Features and Improvements	1
1.1 Features Added in 12	1
1.1.1 OSS	1
1.1.1.1 PostgreSQL Rebase	1
1.1.1.2 Update of OSS Provided	1
1.1.2 Platform Enhancement	
1.1.2.1 Additional Operating System Support for Server Feature	
1.1.2.2 Additional Operating System Support for Client Feature	
1.1.2.3 Additional Operating System Support for Server Assistant Feature	
1.1.3 Operation	
1.1.3.1 Database Multiplexing Improvements	
1.1.3.2 Connection Manager	
1.1.4 Performance	
1.1.4.1 Global Meta Cache	
1.1.4.2 VACUUM/TRUNCATE Improvements	
1.1.5 Application Development	
1.1.5.1 Support for Visual Studio 2019	
1.1.5.2 Support for .NET Framework 4.8	4
Chapter 2 Compatibility Information	5
2.1 Installation/Setup Incompatibility	
2.1.1 Removing Operating System Support for Client Feature	
2.1.2 Removing Operating System Support for Server Feature	
2.1.3 Changing the Way OSS is Set Up	
2.1.4 Modifying Pgpool-II Installation Handling.	
2.1.5 Changing Core and Log File Paths when Instance is Created with WebAdmin	
2.1.6 Renaming WebAdmin Services.	
2.1.7 Changing the Features Targeted for Installation in a 64-bit Environment	
2.1.8 Changing the Access Permissions of the Windows Client Installation Folder	
2.1.9 Changing the WebAdmin Installation Method	
2.2 Application Migration Incompatibility	
2.2.1 Changed to Error when Running an Operator or Function that Returns non Data Types for Masking Type	8
2.2.2 Changing the Method of Specifying the Application Connection Switch Feature	9
2.2.3 PostgreSQL Compatibility of Embedded SQL Applications in C and COBOL	9
2.2.4 Changing Vertical Clustered Index (VCI)	9
2.2.5 Changing how to Use the Features Compatible with Oracle Databases	10
2.3 Operation Migration Incompatibility	10
2.3.1 Mirroring Controller no Longer Retries to Monitor Database Processes when they are Detected as Down	11
2.3.2 Changing the Name and Parameter Name of the Mirroring Controller Post-Promote Command	
2.3.3 Changing Mirroring Controller User Command Input Values	
2.3.4 Changing the Maximum Number of Connections per Server	
2.3.5 Changing the Encryption Settings Using the ALTER TABLESPACE Statement	
2.3.6 Changing the Default Configuration of the Cluster System Using Database Multiplexing	
2.3.7 Changing the Default Operation when mc_ctl Command Options are Omitted	
2.3.8 Changing the Connection Settings when Mirroring Controller Connects to an Instance	
2.3.9 Changing the Status Display of the Mirroring Controller Server	
$2.3.10ChangingtheOperationwhenthesynchronous\_standby\_namesParameterisChangedduringDatabaseMultiplexingOperationMultiplexi$	
2.3.11 Changing Masking Policy Definition for Unsupported Data Types	
2.4 .NET Data Provider Incompatibility	
2.4.1 Changing the Server Explorer View in Visual Studio Integration with .NET Data Provider	
2.5 oracle_fdw Incompatibility	
2.5.1 Changing the Oracle Client Version	16
Chapter 3 Program Updates	17

## Chapter 1 New Features and Improvements

This chapter explains FUJITSU Enterprise Postgres new features and improvements added in this version.

Table 1.1 New features and improvements

Version and level	Classification	Feature	Provided in AE	Provided in SE
12	OSS	PostgreSQL Rebase	Y	Y
		Update of OSS Provided	Y	Y
	Platform enhancement	Additional Operating System Support for Server Feature	Y	Y
		Additional Operating System Support for Client Feature	Y	Y
		Additional Operating System Support for Server Assistant Feature	Y	N
	Operation	Database Multiplexing Improvements	Y	N
		Connection Manager	Y	N
	Performance	Global Meta Cache	Y	N
		VACUUM/TRUNCATE Improvements	Y	Y
	Application	Support for Visual Studio 2019	Y	Y
	development	Support for .NET Framework 4.8	Y	Y

## 1.1 Features Added in 12

This section explains new features and improvements in FUJITSU Enterprise Postgres 12.

#### 1.1.1 OSS

This section explains the new feature related to OSS:

- PostgreSQL rebase
- Update of OSS provided

#### 1.1.1.1 PostgreSQL Rebase

The PostgreSQL version that FUJITSU Enterprise Postgres is based on is 12.1.



Refer to "PostgreSQL Version Used for FUJITSU Enterprise Postgres" in the Installation and Setup Guide for Server for details.

#### 1.1.1.2 Update of OSS Provided

The OSS provided by FUJITSU Enterprise Postgres have been updated.



Refer to "OSS Supported by FUJITSU Enterprise Postgres" in the General Description for details.

#### 1.1.2 Platform Enhancement

This section explains the new features related to platform enhancement:

- Additional operating system support for server
- Additional operating system support for client
- Additional operating system support for server assistant

#### 1.1.2.1 Additional Operating System Support for Server Feature

The following additional operating system is supported:

- RHEL8.1



See

Refer to "Required Operating System" in the Installation and Setup Guide for Server for details.

#### 1.1.2.2 Additional Operating System Support for Client Feature

The following additional operating system is supported:

- RHEL8.1



See

Refer to "Required Operating System" in the Installation and Setup Guide for Client for details.

#### 1.1.2.3 Additional Operating System Support for Server Assistant Feature

The following additional operating system is supported:

- RHEL8.1



See

Refer to "Required Operating System" in the Installation and Setup Guide for Server Assistant for details.

### 1.1.3 Operation

This section explains the new feature related to OSS:

- Database Multiplexing Improvements
- Connection Manager

#### 1.1.3.1 Database Multiplexing Improvements

The database multiplexing introduces the following new features.

- Heartbeat monitoring of database processes

Monitoring retries no longer occur when a database process is detected as down.



See

Refer to "Server Configuration File for the Database Servers" in the Cluster Operation Guide (Database Multiplexing) for details.

- Faster detection time for heartbeat monitoring of the operating system or server

For database degradation using the arbitration server, if the database servers are unable to check their mutual statuses (due to a network error between database servers or server instability), then the database server queries the arbitration server for the status of the other database server.

Previously, when an arbitration server received a query, it inquired about the state of the other database server, but now the arbitration server can respond immediately to the query from the database server by periodically heartbeating the database server. Accordingly, reduced error detection time in heartbeat monitoring of the operating system or server.



#### See

Refer to "Tuning Abnormality Monitoring for Operations that Use an Arbitration Server for Automatic Degeneration" in the Cluster Operation Guide (Database Multiplexing) for details.

- WebAdmin Mirroring Controller integration

You can now set the heartbeat\_error\_action parameter to specify what happens when a heartbeat error is detected in the heartbeat monitoring of the operating system or server during Mirroring Controller setup in WebAdmin. You can also edit the settings after you set up Mirroring Controller. Accordingly, the database operation using WebAdmin becomes more flexible.



#### See

Refer to "Mirroring Controller Configuration" in the Cluster Operation Guide (Database Multiplexing) for details.

#### 1.1.3.2 Connection Manager

With the Connection Manager features, replication operation can be continued without being aware of the connection destination of the applications.

The Connection Manager provides the following capabilities:

- When a client error occurs, force a reclaim of the SQL connection with the client.
- When a server error occurs, the application is notified as a SQL error.
- Applications connect to the appropriate server without being aware of the server state (Primary, Standby).



#### See

Refer to Connection Manager User's Guide for details.

#### 1.1.4 Performance

This section explains the new feature added to improve performance:

- Global Meta Cache
- VACUUM/TRUNCATE improvements

#### 1.1.4.1 Global Meta Cache

Before FUJITSU Enterprise Postgres 11, information about system catalogs and table definitions was cached in per-process memory. You can cache some of this information in shared memory by using the Global Meta Cache feature. This reduces overall system memory usage.



#### See

Refer to "Memory usage reduction by Global Meta Cache" in the General Description for details.

#### 1.1.4.2 VACUUM/TRUNCATE Improvements

If VACUUM/TRUNCATE truncates empty pages and relation indexes at the end of a transaction, or if relations are truncated, improve processing time by doing the following:

- Reduce the number of scan loops when deleting a relation buffer
- Reduce the number of shared buffer scans

This reduces WAL apply time independent of the size of the shared buffer.

### 1.1.5 Application Development

This section explains the new features related to application development:

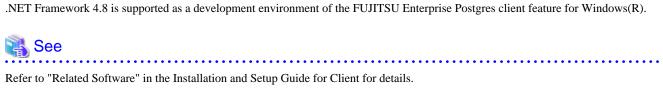
- Support for Visual Studio 2019
- Support for .NET Framework 4.8

#### 1.1.5.1 Support for Visual Studio 2019

Visual Studio 2019 is supported as a development environment of the FUJITSU Enterprise Postgres client feature for Windows(R).

See
Refer to "Related Software" in the Installation and Setup Guide for Client for details

#### 1.1.5.2 Support for .NET Framework 4.8



## Chapter 2 Compatibility Information

This chapter explains incompatible items and actions required when migrating from an earlier version to FUJITSU Enterprise Postgres 12. Check compatibility before migrating and take the appropriate action.

## 2.1 Installation/Setup Incompatibility

Item		Pre-migration version			
nem	9.5	9.6	10	11	
Removing Operating System Support for Client Feature	Y	Y	Y	Y	
Removing Operating System Support for Server Feature	Y	Y	Y	Y	
Changing the way OSS is set up	Y	Y	Y	Y	
Modifying Pgpool-II Installation Handling	Y	Y	Y	Y	
Changing Core and Log File Paths when Instance is Created with WebAdmin	Y	Y	Y	Y	
Renaming WebAdmin Services	Y	Y	Y	Y	
Changing the features targeted for installation in a 64-bit environment	Y	Y	N	N	
Changing the access permissions of the Windows client installation folder	Y	Y	N	N	
Changing the WebAdmin installation method	Y	N	N	N	

Y: Incompatibility exists

## 2.1.1 Removing Operating System Support for Client Feature

#### Incompatibility

In FUJITSU Enterprise Postgres 12 or later, the following operating systems have been removed.

- Windows(R) 7
- Windows Server(R) 2008
- RHEL6

Also, the 32 bit Linux client can no longer be installed because RHEL 6 has been removed from the operating system.

#### **Action method**

None.

## 2.1.2 Removing Operating System Support for Server Feature

#### Incompatibility

In FUJITSU Enterprise Postgres 12 or later, the following operating systems have been removed.

- RHEL6

#### **Action method**

None.

N: Incompatibility does not exist

## 2.1.3 Changing the Way OSS is Set Up

#### Incompatibility

FUJITSU Enterprise Postgres 12 and later do not place OSS extension modules in the executable directory. The OSS extension modules must be placed in the executable directory when you set up OSS.

Refer to "Setting Up and Removing OSS" in the Installation and Setup Guide for Server for details.

#### **Action method**

None.

## 2.1.4 Modifying Pgpool-II Installation Handling

#### Incompatibility

For FUJITSU Enterprise Postgres 12 and later, Pgpool-II is not automatically installed when you install the server. Therefore, if you want to take advantage of Pgpool-II, install it separately from the server installation.

The extension modules required for the database server are shipped with the server program. You should set up Pgpool-II on the database server side, even if Pgpool-II is to be used on a different server than the database server.

Refer to "Setting Up and Removing OSS" in the Installation and Setup Guide for Server for details.

#### **Action method**

None.

# 2.1.5 Changing Core and Log File Paths when Instance is Created with WebAdmin

#### Incompatibility

In FUJITSU Enterprise Postgres 12 and later, change the core and log file paths when creating an instance in WebAdmin.

#### FUJITSU Enterprise Postgres 11 or earlier

Log File Path: /var/tmp/fsep\_version/instanceAdminUser\_instanceNamePortNumber/log

Core File Path: /var/tmp/fsep\_version/instanceAdminUser\_instanceNamePortNumber/core

version: product version\_edition\_architecture

[Example]

Log File Path: /var/tmp/fsep\_110\_AE\_64/naomi\_myinst27599/log

Core File Path: /var/tmp/fsep\_110\_AE\_64/naomi\_myinst27599/core

#### FUJITSU Enterprise Postgres 12

Log File Path: /var/tmp/fsep\_version/instanceAdminUser\_instanceNamePortNumber/log

Core File Path: /var/tmp/fsep\_version/instanceAdminUser\_instanceNamePortNumber/core

version: product version\_WA\_architecture

[Example]

Log File Path: /var/tmp/fsep\_120\_WA\_64/naomi\_myinst27599/core

Core File Path: /var/tmp/fsep\_120\_WA\_64/naomi\_myinst27599/log

#### **Action method**

None.

## 2.1.6 Renaming WebAdmin Services

#### Incompatibility

In FUJITSU Enterprise Postgres 12 and later, change the service name registered when you set up WebAdmin.

FUJITSU Enterprise Postgres 11 or earlier

fsep\_xSPz\_edition\_64\_WebAdmin\_Port1

fsep\_xSPz\_edition\_64\_WebAdmin\_Port2

**FUJITSU Enterprise Postgres 12** 

fsep\_xSPz\_WA\_64\_WebAdmin\_Port1

fsep\_xSPz\_WA\_64\_WebAdmin\_Port2

#### **Action method**

None.

### 2.1.7 Changing the Features Targeted for Installation in a 64-bit Environment

#### Incompatibility

32 bit Linux client can no longer be installed on RHEL7 in FUJITSU Enterprise Postgres 10 and later.

Note: 32 bit Linux client can install on RHEL6 in FUJITSU Enterprise Postgres 11 or earlier.

#### **Action method**

None.

# 2.1.8 Changing the Access Permissions of the Windows Client Installation Folder

#### Incompatibility

When changed for install folder the Windows client from default, the access permissions of the installed files and folder changes to read and execute.

The above changes apply to non-administrators.

This incompatibility occurs with all supported Windows client operating systems. Refer to "Required Operating System" in the Installation and Setup Guide for Client for information on the operating environment.

#### **Action method**

This change was made simply to enhance security, therefore, no action is normally required.

However, if existing access permissions are required, this change can be reverted by running *installDir*\setup\revert\_cacls.bat as an administrator after installation.

## 2.1.9 Changing the WebAdmin Installation Method

#### Incompatibility

In FUJITSU Enterprise Postgres 9.6 or later, WebAdmin is not installed automatically during server installation. Therefore, install WebAdmin separately.

#### **Action method**

None.

## 2.2 Application Migration Incompatibility

ltem		Pre-migration version			
item	9.5	9.6	10	11	
Changed to Error when Running an Operator or Function that Returns non Data Types for Masking Type	Y	Y	Y	Y	
Changing the method of specifying the application connection switch feature	Y	Y	N	N	
PostgreSQL compatibility of embedded SQL applications in C and COBOL	Y	Y	N	N	
Changing Vertical Clustered Index (VCI)	Y	Y	N	N	
Changing how to use the features compatible with Oracle databases	Y	N	N	N	

Y: Incompatibility exists

# 2.2.1 Changed to Error when Running an Operator or Function that Returns non Data Types for Masking Type

#### Incompatibility

In FUJITSU Enterprise Postgres 12, changed to error when running an operator or function that returns non Data Types for Masking type.

#### FUJITSU Enterprise Postgres 11 or earlier

In the following cases, the operator or function in 3) may be executed without masking.

This is an incorrect result because the data containing the masking column is performed without masking.

- 1) Create a Masking policy. and
- 2) Execute a SELECT statement. and
- 3) Execute an operator or function in the SELECT statement of 2). and
- 4) The operator or function argument of 3) includes a subquery. and
- 5) The operator or function argument of 3) contains the column to be protected created in 1). and
- 6) The operator or function of 3) returns a type not listed in "Data Types for Masking" in "Operation Guide".

#### FUJITSU Enterprise Postgres 12 or later

Operator or function results of 3) in the following error.

```
{\tt ERROR:}\  The output data type is incompatible with the confidential policy.
```

HINT: Data type of the result value(s) produced by expression/function using confidential columns is not supported by Data masking module. Consider removing confidential columns from such expressions/functions.

This is correct because the operator or function returns a data type not listed in "Data Types for Masking" in "Operation Guide".

#### **Action method**

 $Do \ one \ of \ the \ following \ to \ ensure \ that \ the \ results \ are \ identical \ to \ those \ of \ FUJITSU \ Enterprise \ Postgres \ 11 \ and \ earlier:$ 

- 1) Modify the Masking policy to prevent masking from being performed for the user executing the SQL.
- 2) Modify SQL to not use operators or functions that return types not listed in "Data Types for Masking" in "Operation Guide".

N: Incompatibility does not exist

# 2.2.2 Changing the Method of Specifying the Application Connection Switch Feature

#### Incompatibility

The target server specified using the application connection switch feature is changed to the same name as that of PostgreSQL.

#### **Action method**

Use the FUJITSU Enterprise Postgres 10 or later client to execute applications that have the target server specified as the application connection switch feature after changing the target server and specified value as listed below and recompiling.

#### Specified name of the target server

Usage target	Before correction	After correction
JDBC driver	No change	
ODBC driver	TargetServer	target_session_attrs
.NET Data Provider	target_server	TargetServerType
Connection service file	target_server	target_session_attrs
Library (libpq) for C	target_server  Environment variable (PGXTARGETSERVER)	target_session_attrs  Environment variable (PGTARGETSESSIONATTRS)
Embedded SQL	target_server	target_session_attrs
psql	Environment variable (PGXTARGETSERVER)	Environment variable (PGTARGETSESSIONATTRS)

#### Specified value for the target server

Server	JDBC driver	.NET Data	Provider	Other driver		
selection order		Before correction	After correction	Before correction	After correction	
Primary server	No change	primary	master	primary	read-write	
Standby server	No change	-	slave	-	-	
Priority given to a standby server	No change	prefer_standby	preferSlave	prefer_standby	prefer-read	
Any	No change	-	any	-	any	

# 2.2.3 PostgreSQL Compatibility of Embedded SQL Applications in C and COBOL

#### Incompatibility

Embedded SQL applications in C and COBOL that were compiled with FUJITSU Enterprise Postgres 9.6 or earlier cannot be used in the FUJITSU Enterprise Postgres 10 client.

#### **Action method**

Use the FUJITSU Enterprise Postgres 10 or later client to execute the applications only after recompiling them with the FUJITSU Enterprise Postgres 10 or later client.

## 2.2.4 Changing Vertical Clustered Index (VCI)

#### Incompatibility

The functions below have been added to the functions for which VCI is not used.

Туре	Function
Date/time functions	age(timestamp)
	current_date
	current_timestamp
	current_time
	localtime
	localtimestamp
Session information	current_user
functions	current_role

#### **Action method**

To use VCI, specify the second argument when using the age function, and specify other functions as subqueries.

[Example]

Before: select age(column A), current\_date from table

After: select age(column A, now()), (select current\_date) from table

## 2.2.5 Changing how to Use the Features Compatible with Oracle Databases

#### Incompatibility

In FUJITSU Enterprise Postgres 9.6 or later, to use the features compatible with Oracle databases, create a new instance and execute the following command for the "postgres" and "template1" databases:

CREATE EXTENSION oracle\_compatible

#### **Action method**

None.

## 2.3 Operation Migration Incompatibility

Item	Pre-migration version			
пен	9.5	9.6	10	11
Mirroring Controller no longer retries to monitor database processes when they are detected as down	Y	N	Y	Y
Changing the Name and Parameter Name of the Mirroring Controller Post-Promote Command		N	Y	Y
Changing Mirroring Controller User Command Input Values	N	N	Y	Y
Changing the maximum number of connections per server	Y	Y	N	N
Changing the encryption settings using the ALTER TABLESPACE statement	Y	Y	N	N
Changing the default configuration of the cluster system using database multiplexing	Y	N	N	N
Changing the default operation when mc_ctl command options are omitted	Y	N	N	N

Item	Pre-migration version			
nem	9.5	9.6	10	11
Changing the connection settings when Mirroring Controller connects to an instance	Y	N	N	N
Changing the status display of the Mirroring Controller server	Y	N	N	N
Changing the operation when the synchronous_standby_names parameter is changed during database multiplexing operation	Y	N	N	N
Changing masking policy definition for unsupported data types	Y	N	N	N

Y: Incompatibility exists

# 2.3.1 Mirroring Controller no Longer Retries to Monitor Database Processes when they are Detected as Down

#### Incompatibility

For the FUJITSU Enterprise Postgres Mirroring Controller, change the heartbeat monitoring of the database process so that it does not retry monitoring when it detects down.

#### **Action method**

None.

# 2.3.2 Changing the Name and Parameter Name of the Mirroring Controller Post-Promote Command

#### Incompatibility

In the FUJITSU Enterprise Postgres 12 Mirroring Controller, change the name of the post-promote command, which is the state transition command, and the parameter name in the server configuration file that specifies the post-promote command.

#### FUJITSU Enterprise Postgres 11 or earlier

- Command Name
  - post-promote command
- The parameter name in the server configuration file that specifies the post-promote command post\_promote\_command

#### FUJITSU Enterprise Postgres 12 or later

- Command Name
  - post-switch command
- The parameter name in the server configuration file that specifies the post-promote command post\_switch\_command

#### **Action method**

The post\_promote\_command parameter in the server configuration file continues to be available in FUJITSU Enterprise Postgres 12 and later.If specified in the server configuration file, it acts as a post-switch command.You cannot specify the post\_promote\_command and post\_switch\_command parameters at the same time.

N: Incompatibility does not exist

## 2.3.3 Changing Mirroring Controller User Command Input Values

#### Incompatibility

Arguments (Fixed value: primarycenter) have been added to the following user commands:

- Fencing command of the database server
- Arbitration command
- Post-switch command
- Pre-detach command
- Post-attach command

#### **Action method**

If you are checking the number of arguments in a user command, increase the number of arguments by one.

### 2.3.4 Changing the Maximum Number of Connections per Server

#### Incompatibility

The maximum number of connections per server is changed from 262,143 to 65,535.

#### **Action method**

None.

## 2.3.5 Changing the Encryption Settings Using the ALTER TABLESPACE Statement

#### Incompatibility

The ALTER TABLESPACE statement can be used to change the tablespace encryption settings if data is not stored in the tablespace.

#### **Action method**

None.

# 2.3.6 Changing the Default Configuration of the Cluster System Using Database Multiplexing

#### Incompatibility

In FUJITSU Enterprise Postgres 10, a split brain will not occur if heartbeat monitoring using an admin network times out, so the default is changed to a system configuration requiring an arbitration server acting as a third party.

FUJITSU Enterprise Postgres 9.6 or earlier

A cluster system comprises two database servers.

FUJITSU Enterprise Postgres 10 or later

By default, a cluster system comprises two database servers and an arbitration server.

#### **Action method**

If selecting the same system configuration as that in FUJITSU Enterprise Postgres 9.6 or earlier, set the parameters below in the server configuration file to perform automatic degradation unconditionally when a heartbeat abnormality is detected during OS/server heartbeat monitoring.

- Parameter: heartbeat\_error\_action

- Value: fallback

## 2.3.7 Changing the Default Operation when mc\_ctl Command Options are Omitted

#### Incompatibility

In FUJITSU Enterprise Postgres 10, the default operation when the mc\_ctl command options below are omitted has been changed:

- The -f option during start mode
- The -w option during start mode

#### FUJITSU Enterprise Postgres 9.6 or earlier

- If the -f option is omitted, automatic switching and disconnection immediately after the startup of Mirroring Controller will not be enabled.
- If the -w option is omitted, the system will not wait for operations to finish.

#### FUJITSU Enterprise Postgres 10 or later

- Even if the -f option is omitted, automatic switching and disconnection immediately after the startup of Mirroring Controller will be enabled.
- Even if the -w option is omitted, the system will wait for operations to finish.

#### Action method

If selecting the same operation as that in FUJITSU Enterprise Postgres 9.6 or earlier, take the action below when executing start mode of the mc\_ctl command.

- If the -f option is omitted, specify the -F option.
- If the -w option is omitted, specify the -W option.

# 2.3.8 Changing the Connection Settings when Mirroring Controller Connects to an Instance

#### Incompatibility

In FUJITSU Enterprise Postgres 10, Mirroring Controller changes the connection setting below when connecting to an instance to detect failure of each database element.

- Application name

#### FUJITSU Enterprise Postgres 9.6 or earlier

- The application name is an empty string.

#### FUJITSU Enterprise Postgres 10 or later

- The application name is 'mc\_agent'.

#### **Action method**

There is no method for changing the application name.

Therefore, if there is an application that identifies a database connection session of Mirroring Controller with the application name being an empty string, modify the process so that identification is performed using 'mc\_agent'.

## 2.3.9 Changing the Status Display of the Mirroring Controller Server

#### Incompatibility

In FUJITSU Enterprise Postgres 10, the condition for displaying an abnormality for the process (WAL receive process) that receives transaction logs has been changed in regard to the server status display performed in status mode of the mc\_ctl command of Mirroring Controller.

- Role of the applicable server (host\_role): standby (standby)
- Display item: DBMS process status (db\_proc\_status)
- Display content: abnormal (abnormal process name)

  The applicability for incompatibility is determined based on whether 'wal\_receiver' is included in the abnormality process name.

#### FUJITSU Enterprise Postgres 9.6 or earlier

Displays an abnormality when Mirroring Controller detects that the WAL send process has stopped.

#### FUJITSU Enterprise Postgres 10 or later

Displays an abnormality when Mirroring Controller detects that a streaming replication connection has not been established with the primary server using the WAL receive process.

#### **Action method**

If there is an application that determines the WAL receive process status from the display content retrieved using status mode of the mc\_ctl command, modify as below to determine at the same level of accuracy as previously. However, this is not recommended due to incompatibilities accompanying the improvement in accuracy.

- If 'wal\_sender' is included in the abnormal process name displayed in the 'abnormal' db\_proc\_status item of the 'primary' host\_role, it is determined that an abnormality has occurred in the WAL receive process.
- If 'wal\_sender' is not included in the abnormal process name displayed in the 'abnormal' db\_proc\_status item of the 'primary' host\_role, it is determined that an abnormality has not occurred in the WAL receive process.

# 2.3.10 Changing the Operation when the synchronous\_standby\_names Parameter is Changed during Database Multiplexing Operation

#### Incompatibility

In FUJITSU Enterprise Postgres 9.6 or earlier, Mirroring Controller periodically accesses the database to retrieve the synchronous\_standby\_names parameter value using the SHOW command, and automatically recovers if changes are detected, in case the user accidentally changes the synchronous\_standby\_names parameter in postgresql.conf during database multiplexing mode. However, because this process uses the CPU of the database server for redundancy and SQL statements are executed with high frequency, these processes are not executed by default in FUJITSU Enterprise Postgres 10.

#### FUJITSU Enterprise Postgres 9.6 or earlier

Mirroring Controller checks if the synchronous standby names parameter in postgresql.conf has been mistakenly changed by the user.

#### FUJITSU Enterprise Postgres 10 or later

By default, Mirroring Controller does not check if the synchronous\_standby\_names parameter in postgresql.conf has been mistakenly changed by the user.

#### **Action method**

If selecting the same operation as that in FUJITSU Enterprise Postgres 9.6 or earlier, set the parameter below in the server definition file.

- Parameter: check\_synchronous\_standby\_names\_validation
- Value: 'on'

## 2.3.11 Changing Masking Policy Definition for Unsupported Data Types

#### Incompatibility

The data masking feature of FUJITSU Enterprise Postgres is updated so that an error occurs when a masking policy is defined for a column of an unsupported data type (array type or timestamp with timezone type).

#### FUJITSU Enterprise Postgres 9.5

The definition of a masking policy for a column of array type or timestamp with timezone type ends normally, but an error occurs when accessing a column of a table for which the policy is defined.

#### FUJITSU Enterprise Postgres 9.6 or later

An error occurs when defining a masking policy for a column of array type or timestamp with timezone type.

#### **Action method**

If a policy that was defined in FUJITSU Enterprise Postgres 9.5 includes masking of a column of an unsupported data type (array type or timestamp with timezone type), perform one of the following in FUJITSU Enterprise Postgres 9.6 or later:

- Ignore the error if the masking policy contains columns of only these data types, or delete the policy with the pgx\_drop\_confidential\_policy system administration function.
- If the masking policy target contains columns of these data types and a supported data type, delete the problematic columns with the pgx\_alter\_confidential\_policy system function.

## 2.4 .NET Data Provider Incompatibility

Item		Pre-migration version			
		9.6	10	11	
Changing the Server Explorer View in Visual Studio Integration with .NET Data Provider	Y	Y	N	N	

Y: Incompatibility exists

N: Incompatibility does not exist

# 2.4.1 Changing the Server Explorer View in Visual Studio Integration with .NET Data Provider

#### Incompatibility

FUJITSU Enterprise Postgres 10 and later, the following changes occur when you view database resources in Server Explorer.

- "Schema Name. Table Name" display under the Tables folder.
- The Indexes folder is not displayed.
- The Triggers folder does not displayed.
- The Sequences folder is not displayed.
- The Procedures folder does not displayed.

#### **Action method**

None.

## 2.5 oracle\_fdw Incompatibility

ltoro	Pre-migration version			
ltem		9.6	10	11
Changing the Oracle Client Version	N	N	Y	Y

Y: Incompatibility exists

N: Incompatibility does not exist

## 2.5.1 Changing the Oracle Client Version

#### Incompatibility

FUJITSU Enterprise Postgres 12 change the version of the Oracle client used to build oracle\_fdw to 11.2.

#### **Action method**

Use Oracle client version 11.2 or later.

Also, if a file named libclntsh.so.11.1 does not exist in OCI library, create a symbolic link named libclntsh.so.11.1.

## Chapter 3 Program Updates

For program fix information, see "Program Updates".

## <u>Index</u>

Compatibility Information	[C]
	[F]
	[P]



# FUJITSU Enterprise Postgres 12



Release Notes (Windows)

## Preface

#### Purpose of this document

This document provides release information for FUJITSU Enterprise Postgres.

#### Structure of this document

This document is structured as follows:

Chapter 1 New Features and Improvements

Explains the new features and improvements in this version.

Chapter 2 Compatibility Information

Provides information regarding compatibility.

Chapter 3 Program Updates

Explains updates incorporated in this version.

#### **Export restrictions**

Exportation/release of this document may require necessary procedures in accordance with the regulations of your resident country and/or US export control laws.

#### Issue date and version

Edition 1.0: July 2020

#### Copyright

Copyright 2015-2020 FUJITSU LIMITED

## Contents

Chapter 1 New Features and Improvements	1
1.1 Features Added in 12	
1.1.1 OSS	
1.1.1.1 PostgreSQL Rebase	
1.1.1.2 Update of OSS Provided	1
1.1.2 Platform Enhancement.	
1.1.2.1 Additional Operating System Support for Client Feature	2
1.1.2.2 Additional Operating System Support for Server Assistant Feature	
1.1.3 Operation.	
1.1.3.1 Database Multiplexing Improvements	
1.1.4 Performance	
1.1.4.1 Global Meta Cache	3
1.1.4.2 VACUUM/TRUNCATE Improvements	
1.1.5 Application Development	
1.1.5.1 Support for Visual Studio 2019	
1.1.5.2 Support for .NET Framework 4.8	
Chapter 2 Compatibility Information	5
2.1 Installation/Setup Incompatibility	4
2.1.1 Removing Operating System Support for Client Feature	4
2.1.2 Removing Operating System Support for Server Feature	4
2.1.3 Changing the Way OSS is Set Up	<i>(</i>
2.1.4 Changing Core and Log File Paths when Instance is Created with WebAdmin	6
2.1.5 Renaming WebAdmin Services	6
2.1.6 Changing the Features Targeted for Installation in a 64-bit Environment	
2.1.7 Changing the Access Permissions of the Windows Client Installation Folder	
2.1.8 Changing the WebAdmin Installation Method	
2.2 Application Migration Incompatibility	
2.2.1 Changed to Error when Running an Operator or Function that Returns non Data Types for Masking Type	8
2.2.2 Removing UTL_FILE for Oracle Database Compatibility Features	8
2.2.3 Changing the Method of Specifying the Application Connection Switch Feature	9
2.2.4 PostgreSQL Compatibility of Embedded SQL Applications in C and COBOL	9
2.2.5 Changing Vertical Clustered Index (VCI)	9
2.2.6 Changing how to Use the Features Compatible with Oracle Databases	10
2.3 Operation Migration Incompatibility	10
2.3.1 Mirroring Controller no Longer Retries to Monitor Database Processes when they are Detected as Down	1
2.3.2 Changing the Name and Parameter Name of the Mirroring Controller Post-Promote Command	1
2.3.3 Changing Mirroring Controller User Command Input Values	12
2.3.4 Changing the Maximum Number of Connections per Server	12
2.3.5 Changing the Encryption Settings Using the ALTER TABLESPACE Statement	12
2.3.6 Changing the Default Configuration of the Cluster System Using Database Multiplexing	12
2.3.7 Changing the Default Operation when mc_ctl Command Options are Omitted	13
2.3.8 Changing the Connection Settings when Mirroring Controller Connects to an Instance	13
2.3.9 Changing the Status Display of the Mirroring Controller Server	13
2.3.10 Changing the Operation when the synchronous_standby_names Parameter is Changed during Database Multiplexing Operation	peration
	14
2.3.11 Changing Masking Policy Definition for Unsupported Data Types	
2.4 .NET Data Provider Incompatibility	
2.4.1 Changing the Server Explorer View in Visual Studio Integration with .NET Data Provider	
Chapter 3 Program Updates	16
To allow	4-

## Chapter 1 New Features and Improvements

This chapter explains FUJITSU Enterprise Postgres new features and improvements added in this version.

Table 1.1 New features and improvements

Version and level	Classification	Feature	Provided in AE	Provided in SE
12	OSS	PostgreSQL Rebase	Y	Y
		Update of OSS Provided	Y	Y
	Platform enhancement	Additional Operating System Support for Client Feature	Y	Y
		Additional Operating System Support for Server Assistant Feature	Y	N
	Operation	Database Multiplexing Improvements	Y	N
	Performance	Global Meta Cache	Y	N
		VACUUM/TRUNCATE Improvements	Y	Y
	Application	Support for Visual Studio 2019	Y	Y
	development	Support for .NET Framework 4.8	Y	Y

## 1.1 Features Added in 12

This section explains new features and improvements in FUJITSU Enterprise Postgres 12.

#### 1.1.1 OSS

This section explains the new feature related to OSS:

- PostgreSQL rebase
- Update of OSS provided

#### 1.1.1.1 PostgreSQL Rebase

The PostgreSQL version that FUJITSU Enterprise Postgres is based on is 12.1.



Refer to "PostgreSQL Version Used for FUJITSU Enterprise Postgres" in the Installation and Setup Guide for Server for details.

#### 1.1.1.2 Update of OSS Provided

The OSS provided by FUJITSU Enterprise Postgres have been updated.



See

Refer to "OSS Supported by FUJITSU Enterprise Postgres" in the General Description for details.

#### 1.1.2 Platform Enhancement

This section explains the new features related to platform enhancement:

- Additional operating system support for client

- Additional operating system support for server assistant

#### 1.1.2.1 Additional Operating System Support for Client Feature

The following additional operating system is supported:

- RHEL8.1



#### See

Refer to "Required Operating System" in the Installation and Setup Guide for Client for details.

#### 1.1.2.2 Additional Operating System Support for Server Assistant Feature

The following additional operating system is supported:

- RHEL8.1



#### See

Refer to "Required Operating System" in the Installation and Setup Guide for Server Assistant for details.

#### 1.1.3 Operation

This section explains the new feature related to OSS:

- Database Multiplexing Improvements

#### 1.1.3.1 Database Multiplexing Improvements

The database multiplexing introduces the following new features.

- Heartbeat monitoring of database processes

Monitoring retries no longer occur when a database process is detected as down.



#### See

Refer to "Server Configuration File for the Database Servers" in the Cluster Operation Guide (Database Multiplexing) for details.

- Faster detection time for heartbeat monitoring of the operating system or server

For database degradation using the arbitration server, if the database servers are unable to check their mutual statuses (due to a network error between database servers or server instability), then the database server queries the arbitration server for the status of the other database server.

Previously, when an arbitration server received a query, it inquired about the state of the other database server, but now the arbitration server can respond immediately to the query from the database server by periodically heartbeating the database server. Accordingly, reduced error detection time in heartbeat monitoring of the operating system or server.



#### Sec

Refer to "Tuning Abnormality Monitoring for Operations that Use an Arbitration Server for Automatic Degeneration" in the Cluster Operation Guide (Database Multiplexing) for details.

- WebAdmin Mirroring Controller integration

You can now set the heartbeat\_error\_action parameter to specify what happens when a heartbeat error is detected in the heartbeat monitoring of the operating system or server during Mirroring Controller setup in WebAdmin. You can also edit the settings after you set up Mirroring Controller. Accordingly, the database operation using WebAdmin becomes more flexible.



Refer to "Mirroring Controller Configuration" in the Cluster Operation Guide (Database Multiplexing) for details.

#### 1.1.4 Performance

This section explains the new feature added to improve performance:

- Global Meta Cache
- VACUUM/TRUNCATE improvements

#### 1.1.4.1 Global Meta Cache

Before FUJITSU Enterprise Postgres 11, information about system catalogs and table definitions was cached in per-process memory. You can cache some of this information in shared memory by using the Global Meta Cache feature. This reduces overall system memory usage.



Refer to "Memory usage reduction by Global Meta Cache" in the General Description for details.

#### 1.1.4.2 VACUUM/TRUNCATE Improvements

If VACUUM/TRUNCATE truncates empty pages and relation indexes at the end of a transaction, or if relations are truncated, improve processing time by doing the following:

- Reduce the number of scan loops when deleting a relation buffer
- Reduce the number of shared buffer scans

This reduces WAL apply time independent of the size of the shared buffer.

## 1.1.5 Application Development

This section explains the new features related to application development:

- Support for Visual Studio 2019
- Support for .NET Framework 4.8

#### 1.1.5.1 Support for Visual Studio 2019

Visual Studio 2019 is supported as a development environment of the FUJITSU Enterprise Postgres client feature for Windows(R).



Refer to "Related Software" in the Installation and Setup Guide for Client for details.

#### 1.1.5.2 Support for .NET Framework 4.8

.NET Framework 4.8 is supported as a development environment of the FUJITSU Enterprise Postgres client feature for Windows(R).



Refer to "Related Software" in the Installation and Setup Guide for Client for details.

# **Chapter 2 Compatibility Information**

This chapter explains incompatible items and actions required when migrating from an earlier version to FUJITSU Enterprise Postgres 12. Check compatibility before migrating and take the appropriate action.

## 2.1 Installation/Setup Incompatibility

Item		Pre-migrat	ion version	
iteiii	9.5	9.6	10	11
Removing Operating System Support for Client Feature	Y	Y	Y	Y
Removing Operating System Support for Server Feature	Y	Y	Y	Y
Changing the way OSS is set up	Y	Y	Y	Y
Changing Core and Log File Paths when Instance is Created with WebAdmin	Y	Y	Y	Y
Renaming WebAdmin Services	Y	Y	Y	Y
Changing the features targeted for installation in a 64-bit environment	Y	Y	N	N
Changing the access permissions of the Windows client installation folder	Y	Y	N	N
Changing the WebAdmin installation method	Y	N	N	N

Y: Incompatibility exists

## 2.1.1 Removing Operating System Support for Client Feature

#### Incompatibility

In FUJITSU Enterprise Postgres 12 or later, the following operating systems have been removed.

- Windows(R) 7
- Windows Server(R) 2008
- RHEL6

Also, the 32 bit Linux client can no longer be installed because RHEL 6 has been removed from the operating system.

#### **Action method**

None.

## 2.1.2 Removing Operating System Support for Server Feature

#### Incompatibility

In FUJITSU Enterprise Postgres 12 or later, the following operating systems have been removed.

- Windows Server(R) 2008

#### **Action method**

None.

N: Incompatibility does not exist

### 2.1.3 Changing the Way OSS is Set Up

#### Incompatibility

FUJITSU Enterprise Postgres 12 and later do not place OSS extension modules in the executable directory. The OSS extension modules must be placed in the executable directory when you set up OSS.

Refer to "Setting Up and Removing OSS" in the Installation and Setup Guide for Server for details.

#### **Action method**

None.

# 2.1.4 Changing Core and Log File Paths when Instance is Created with WebAdmin

#### Incompatibility

In FUJITSU Enterprise Postgres 12 and later, change the core and log file paths when creating an instance in WebAdmin.

FUJITSU Enterprise Postgres 11 or earlier

Log File Path: userProfileFolder\localSettingsFolder\Fujitsu\fsep\_version\instanceNamePortNumber\log

 $Core\ File\ Path:\ \textit{userProfileFolder} \ local SettingsFolder \ \ Fujitsu \ \ fsep\_\textit{version} \ \ instanceNamePortNumber \ \ core$ 

version: product version\_edition\_architecture

[Example]

 $Log\ File\ Path:\ C:\ Users\ naomi\ App Data\ Local\ Fujitsu\ fsep\_110\_AE\_64\ myinst 27599\ log\ naomi\ App Data\ naomi\ naomi$ 

Core File Path: C:\Users\naomi\AppData\Local\Fujitsu\fsep\_110\_AE\_64\myinst27599\core

#### FUJITSU Enterprise Postgres 12

userProfileFolder\localSettingsFolder\Fujitsu\fsep\_version\instanceNamePortNumber\log

Core File Path: userProfileFolder\localSettingsFolder\Fujitsu\fsep\_version\instanceNamePortNumber\core

version: product version\_WA\_architecture

[Example]

Log File Path: C:\Users\naomi\AppData\Local\Fujitsu\fsep\_120\_WA\_64\myinst27599\log

Core File Path: C:\Users\naomi\AppData\Local\Fujitsu\fsep\_120\_WA\_64\myinst27599\core

#### **Action method**

None.

## 2.1.5 Renaming WebAdmin Services

#### Incompatibility

In FUJITSU Enterprise Postgres 12 and later, change the service name registered when you set up WebAdmin.

FUJITSU Enterprise Postgres 11 or earlier

fsep\_xSPz\_edition\_64\_WebAdmin\_Port1

fsep\_xSPz\_edition\_64\_WebAdmin\_Port2

FUJITSU Enterprise Postgres 12

fsep\_xSPz\_WA\_64\_WebAdmin\_Port1

fsep\_xSPz\_WA\_64\_WebAdmin\_Port2

#### **Action method**

None.

### 2.1.6 Changing the Features Targeted for Installation in a 64-bit Environment

#### Incompatibility

32 bit Linux client can no longer be installed on RHEL7 in FUJITSU Enterprise Postgres 10 and later.

Note: 32 bit Linux client can install on RHEL6 in FUJITSU Enterprise Postgres 11 or earlier.

#### **Action method**

None.

# 2.1.7 Changing the Access Permissions of the Windows Client Installation Folder

#### Incompatibility

When changed for install folder the Windows client from default, the access permissions of the installed files and folder changes to read and execute.

The above changes apply to non-administrators.

This incompatibility occurs with all supported Windows client operating systems. Refer to "Required Operating System" in the Installation and Setup Guide for Client for information on the operating environment.

#### **Action method**

This change was made simply to enhance security, therefore, no action is normally required.

However, if existing access permissions are required, this change can be reverted by running *installDit*\setup\revert\_cacls.bat as an administrator after installation.

## 2.1.8 Changing the WebAdmin Installation Method

#### Incompatibility

In FUJITSU Enterprise Postgres 9.6 or later, WebAdmin is not installed automatically during server installation. Therefore, install WebAdmin separately.

#### **Action method**

None.

## 2.2 Application Migration Incompatibility

Itam		Pre-migration version				
Item	9.5	9.6	10	11		
Changed to Error when Running an Operator or Function that Returns non Data Types for Masking Type	Y	Y	Y	Y		
Removing UTL_FILE for Oracle Database Compatibility Features	Y	Y	Y	Y		
Changing the method of specifying the application connection switch feature	Y	Y	N	N		

ltem		Pre-migration version			
item	9.5	9.6	10	11	
PostgreSQL compatibility of embedded SQL applications in C and COBOL	Y	Y	N	N	
Changing Vertical Clustered Index (VCI)	Y	Y	N	N	
Changing how to use the features compatible with Oracle databases	Y	N	N	N	

Y: Incompatibility exists

# 2.2.1 Changed to Error when Running an Operator or Function that Returns non Data Types for Masking Type

#### Incompatibility

In FUJITSU Enterprise Postgres 12, changed to error when running an operator or function that returns non Data Types for Masking type.

#### FUJITSU Enterprise Postgres 11 or earlier

In the following cases, the operator or function in 3) may be executed without masking.

This is an incorrect result because the data containing the masking column is performed without masking.

- 1) Create a Masking policy. and
- 2) Execute a SELECT statement. and
- 3) Execute an operator or function in the SELECT statement of 2). and
- 4) The operator or function argument of 3) includes a subquery. and
- 5) The operator or function argument of 3) contains the column to be protected created in 1). and
- 6) The operator or function of 3) returns a type not listed in "Data Types for Masking" in "Operation Guide".

#### FUJITSU Enterprise Postgres 12 or later

Operator or function results of 3) in the following error.

```
ERROR: The output data type is incompatible with the confidential policy.

HINT: Data type of the result value(s) produced by expression/function using confidential columns is not supported by Data masking module. Consider removing confidential columns from such expressions/functions.
```

This is correct because the operator or function returns a data type not listed in "Data Types for Masking" in "Operation Guide".

#### **Action method**

Do one of the following to ensure that the results are identical to those of FUJITSU Enterprise Postgres 11 and earlier:

- 1) Modify the Masking policy to prevent masking from being performed for the user executing the SQL.
- 2) Modify SQL to not use operators or functions that return types not listed in "Data Types for Masking" in "Operation Guide".

## 2.2.2 Removing UTL\_FILE for Oracle Database Compatibility Features

#### Incompatibility

In FUJITSU Enterprise Postgres 12, remove the Oracle database compatibility UTL\_FILE.

#### **Action method**

None.

N: Incompatibility does not exist

# 2.2.3 Changing the Method of Specifying the Application Connection Switch Feature

#### Incompatibility

The target server specified using the application connection switch feature is changed to the same name as that of PostgreSQL.

#### **Action method**

Use the FUJITSU Enterprise Postgres 10 or later client to execute applications that have the target server specified as the application connection switch feature after changing the target server and specified value as listed below and recompiling.

#### Specified name of the target server

Usage target	Before correction	After correction
JDBC driver	No change	
ODBC driver	TargetServer	target_session_attrs
.NET Data Provider	target_server	TargetServerType
Connection service file	target_server	target_session_attrs
Library (libpq) for C	target_server	target_session_attrs
	Environment variable (PGXTARGETSERVER)	Environment variable (PGTARGETSESSIONATTRS)
Embedded SQL	target_server	target_session_attrs
psql	Environment variable (PGXTARGETSERVER)	Environment variable (PGTARGETSESSIONATTRS)

#### Specified value for the target server

Server	JDBC driver	.NET Data Provider		Other driver	
selection order		Before correction	After correction	Before correction	After correction
Primary server	No change	primary	master	primary	read-write
Standby server	No change	-	slave	-	-
Priority given to a standby server	No change	prefer_standby	preferSlave	prefer_standby	prefer-read
Any	No change	-	any	-	any

# 2.2.4 PostgreSQL Compatibility of Embedded SQL Applications in C and COBOL

#### Incompatibility

Embedded SQL applications in C and COBOL that were compiled with FUJITSU Enterprise Postgres 9.6 or earlier cannot be used in the FUJITSU Enterprise Postgres 10 client.

#### **Action method**

Use the FUJITSU Enterprise Postgres 10 or later client to execute the applications only after recompiling them with the FUJITSU Enterprise Postgres 10 or later client.

## 2.2.5 Changing Vertical Clustered Index (VCI)

#### Incompatibility

The functions below have been added to the functions for which VCI is not used.

Туре	Function
Date/time functions	age(timestamp)
	current_date
	current_timestamp
	current_time
	localtime
	localtimestamp
Session information	current_user
functions	current_role

#### **Action method**

To use VCI, specify the second argument when using the age function, and specify other functions as subqueries.

[Example]

Before: select age(column A), current\_date from table

After: select age(column A, now()), (select current\_date) from table

## 2.2.6 Changing how to Use the Features Compatible with Oracle Databases

#### Incompatibility

In FUJITSU Enterprise Postgres 9.6 or later, to use the features compatible with Oracle databases, create a new instance and execute the following command for the "postgres" and "template1" databases:

CREATE EXTENSION oracle\_compatible

#### **Action method**

None.

## 2.3 Operation Migration Incompatibility

Item		Pre-migrat	ion version	
item	9.5	9.6	10	11
Mirroring Controller no longer retries to monitor database processes when they are detected as down	Y	N	Y	Y
Changing the Name and Parameter Name of the Mirroring Controller Post-Promote Command	N	N	Y	Y
Changing Mirroring Controller User Command Input Values	N	N	Y	Y
Changing the maximum number of connections per server	Y	Y	N	N
Changing the encryption settings using the ALTER TABLESPACE statement	Y	Y	N	N
Changing the default configuration of the cluster system using database multiplexing	Y	N	N	N
Changing the default operation when mc_ctl command options are omitted	Y	N	N	N

Item	Pre-migration version				
nem	9.5	9.6	10	11	
Changing the connection settings when Mirroring Controller connects to an instance	Y	N	N	N	
Changing the status display of the Mirroring Controller server	Y	N	N	N	
Changing the operation when the synchronous_standby_names parameter is changed during database multiplexing operation	Y	N	N	N	
Changing masking policy definition for unsupported data types	Y	N	N	N	

Y: Incompatibility exists

# 2.3.1 Mirroring Controller no Longer Retries to Monitor Database Processes when they are Detected as Down

#### Incompatibility

For the FUJITSU Enterprise Postgres Mirroring Controller, change the heartbeat monitoring of the database process so that it does not retry monitoring when it detects down.

#### **Action method**

None.

# 2.3.2 Changing the Name and Parameter Name of the Mirroring Controller Post-Promote Command

#### Incompatibility

In the FUJITSU Enterprise Postgres 12 Mirroring Controller, change the name of the post-promote command, which is the state transition command, and the parameter name in the server configuration file that specifies the post-promote command.

#### FUJITSU Enterprise Postgres 11 or earlier

- Command Name
  - post-promote command
- The parameter name in the server configuration file that specifies the post-promote command post\_promote\_command

#### FUJITSU Enterprise Postgres 12 or later

- Command Name
  - post-switch command
- The parameter name in the server configuration file that specifies the post-promote command post\_switch\_command

#### **Action method**

The post\_promote\_command parameter in the server configuration file continues to be available in FUJITSU Enterprise Postgres 12 and later.If specified in the server configuration file, it acts as a post-switch command.You cannot specify the post\_promote\_command and post\_switch\_command parameters at the same time.

N: Incompatibility does not exist

### 2.3.3 Changing Mirroring Controller User Command Input Values

#### Incompatibility

Arguments (Fixed value: primarycenter) have been added to the following user commands:

- Fencing command of the database server
- Arbitration command
- Post-switch command
- Pre-detach command
- Post-attach command

#### **Action method**

If you are checking the number of arguments in a user command, increase the number of arguments by one.

### 2.3.4 Changing the Maximum Number of Connections per Server

#### Incompatibility

The maximum number of connections per server is changed from 262,143 to 65,535.

#### **Action method**

None.

# 2.3.5 Changing the Encryption Settings Using the ALTER TABLESPACE Statement

#### Incompatibility

The ALTER TABLESPACE statement can be used to change the tablespace encryption settings if data is not stored in the tablespace.

#### **Action method**

None.

# 2.3.6 Changing the Default Configuration of the Cluster System Using Database Multiplexing

#### Incompatibility

In FUJITSU Enterprise Postgres 10, a split brain will not occur if heartbeat monitoring using an admin network times out, so the default is changed to a system configuration requiring an arbitration server acting as a third party.

FUJITSU Enterprise Postgres 9.6 or earlier

A cluster system comprises two database servers.

FUJITSU Enterprise Postgres 10 or later

By default, a cluster system comprises two database servers and an arbitration server.

#### **Action method**

If selecting the same system configuration as that in FUJITSU Enterprise Postgres 9.6 or earlier, set the parameters below in the server configuration file to perform automatic degradation unconditionally when a heartbeat abnormality is detected during OS/server heartbeat monitoring.

- Parameter: heartbeat\_error\_action

- Value: fallback

### 2.3.7 Changing the Default Operation when mc\_ctl Command Options are Omitted

#### Incompatibility

In FUJITSU Enterprise Postgres 10, the default operation when the mc\_ctl command options below are omitted has been changed:

- The -f option during start mode
- The -w option during start mode

#### FUJITSU Enterprise Postgres 9.6 or earlier

- If the -f option is omitted, automatic switching and disconnection immediately after the startup of Mirroring Controller will not be enabled.
- If the -w option is omitted, the system will not wait for operations to finish.

#### FUJITSU Enterprise Postgres 10 or later

- Even if the -f option is omitted, automatic switching and disconnection immediately after the startup of Mirroring Controller will be enabled.
- Even if the -w option is omitted, the system will wait for operations to finish.

#### Action method

If selecting the same operation as that in FUJITSU Enterprise Postgres 9.6 or earlier, take the action below when executing start mode of the mc\_ctl command.

- If the -f option is omitted, specify the -F option.
- If the -w option is omitted, specify the -W option.

# 2.3.8 Changing the Connection Settings when Mirroring Controller Connects to an Instance

#### Incompatibility

In FUJITSU Enterprise Postgres 10, Mirroring Controller changes the connection setting below when connecting to an instance to detect failure of each database element.

- Application name

#### FUJITSU Enterprise Postgres 9.6 or earlier

- The application name is an empty string.

#### FUJITSU Enterprise Postgres 10 or later

- The application name is 'mc\_agent'.

#### **Action method**

There is no method for changing the application name.

Therefore, if there is an application that identifies a database connection session of Mirroring Controller with the application name being an empty string, modify the process so that identification is performed using 'mc\_agent'.

## 2.3.9 Changing the Status Display of the Mirroring Controller Server

#### Incompatibility

In FUJITSU Enterprise Postgres 10, the condition for displaying an abnormality for the process (WAL receive process) that receives transaction logs has been changed in regard to the server status display performed in status mode of the mc\_ctl command of Mirroring Controller.

- Role of the applicable server (host\_role): standby (standby)
- Display item: DBMS process status (db\_proc\_status)
- Display content: abnormal (abnormal process name)

  The applicability for incompatibility is determined based on whether 'wal\_receiver' is included in the abnormality process name.

#### FUJITSU Enterprise Postgres 9.6 or earlier

Displays an abnormality when Mirroring Controller detects that the WAL send process has stopped.

#### FUJITSU Enterprise Postgres 10 or later

Displays an abnormality when Mirroring Controller detects that a streaming replication connection has not been established with the primary server using the WAL receive process.

#### **Action method**

If there is an application that determines the WAL receive process status from the display content retrieved using status mode of the mc\_ctl command, modify as below to determine at the same level of accuracy as previously. However, this is not recommended due to incompatibilities accompanying the improvement in accuracy.

- If 'wal\_sender' is included in the abnormal process name displayed in the 'abnormal' db\_proc\_status item of the 'primary' host\_role, it is determined that an abnormality has occurred in the WAL receive process.
- If 'wal\_sender' is not included in the abnormal process name displayed in the 'abnormal' db\_proc\_status item of the 'primary' host\_role, it is determined that an abnormality has not occurred in the WAL receive process.

# 2.3.10 Changing the Operation when the synchronous\_standby\_names Parameter is Changed during Database Multiplexing Operation

#### Incompatibility

In FUJITSU Enterprise Postgres 9.6 or earlier, Mirroring Controller periodically accesses the database to retrieve the synchronous\_standby\_names parameter value using the SHOW command, and automatically recovers if changes are detected, in case the user accidentally changes the synchronous\_standby\_names parameter in postgresql.conf during database multiplexing mode. However, because this process uses the CPU of the database server for redundancy and SQL statements are executed with high frequency, these processes are not executed by default in FUJITSU Enterprise Postgres 10.

#### FUJITSU Enterprise Postgres 9.6 or earlier

Mirroring Controller checks if the synchronous standby names parameter in postgresql.conf has been mistakenly changed by the user.

#### FUJITSU Enterprise Postgres 10 or later

By default, Mirroring Controller does not check if the synchronous\_standby\_names parameter in postgresql.conf has been mistakenly changed by the user.

#### **Action method**

If selecting the same operation as that in FUJITSU Enterprise Postgres 9.6 or earlier, set the parameter below in the server definition file.

- Parameter: check\_synchronous\_standby\_names\_validation
- Value: 'on'

## 2.3.11 Changing Masking Policy Definition for Unsupported Data Types

#### Incompatibility

The data masking feature of FUJITSU Enterprise Postgres is updated so that an error occurs when a masking policy is defined for a column of an unsupported data type (array type or timestamp with timezone type).

#### FUJITSU Enterprise Postgres 9.5

The definition of a masking policy for a column of array type or timestamp with timezone type ends normally, but an error occurs when accessing a column of a table for which the policy is defined.

#### FUJITSU Enterprise Postgres 9.6 or later

An error occurs when defining a masking policy for a column of array type or timestamp with timezone type.

#### **Action method**

If a policy that was defined in FUJITSU Enterprise Postgres 9.5 includes masking of a column of an unsupported data type (array type or timestamp with timezone type), perform one of the following in FUJITSU Enterprise Postgres 9.6 or later:

- Ignore the error if the masking policy contains columns of only these data types, or delete the policy with the pgx\_drop\_confidential\_policy system administration function.
- If the masking policy target contains columns of these data types and a supported data type, delete the problematic columns with the pgx\_alter\_confidential\_policy system function.

## 2.4 .NET Data Provider Incompatibility

Item		Pre-migration version				
		9.6	10	11		
Changing the Server Explorer View in Visual Studio Integration with .NET Data Provider	Y	Y	N	N		

Y: Incompatibility exists

N: Incompatibility does not exist

# 2.4.1 Changing the Server Explorer View in Visual Studio Integration with .NET Data Provider

#### Incompatibility

FUJITSU Enterprise Postgres 10 and later, the following changes occur when you view database resources in Server Explorer.

- "Schema Name. Table Name" display under the Tables folder.
- The Indexes folder is not displayed.
- The Triggers folder does not displayed.
- The Sequences folder is not displayed.
- The Procedures folder does not displayed.

#### **Action method**

None.

# Chapter 3 Program Updates

For program fix information, see "Program Updates".

# <u>Index</u>

Compatibility Information	[C]5
, , , , , , , , , , , , , , , , , , ,	[F]
Features Added in 12	1
Program Updates	[P]16



# FUJITSU Enterprise Postgres 12



# Program Updates

Linux > Windows >



# FUJITSU Enterprise Postgres 12



Program Updates (Linux)

## Preface

This document explains the updates that have been fixed at this version. The contents of this document are subject to change without notice.

#### **Notations**

The status for each edition is shown in the following table.

P number	Update summary	AE	SE
Number that uniquely identifies the update	Summary of update details	Symbol in the fixed	_

ΑE

Indicates Advanced Edition.

SE

Indicates Standard Edition.

Symbol

Indicates the incorporated status for each

edition. Y: Fixed

-: Not relevant to this product

#### **Export restrictions**

Exportation/release of this document may require necessary procedures in accordance with the regulations of your resident country and/or US export control laws.

#### Issue date and version

Edition 1.0: July 2020

#### Copyright

Copyright 2020 FUJITSU LIMITED

# Chapter 1 Program Updates

This version incorporates the updates implemented in PostgreSQL 12, and 12.1.



Refer to the PostgreSQL Global Development Group website for information on the updates implemented in the following releases:

.....

#### [PostgreSQL 12]

https://www.postgresql.org/docs/12/release-12.html

#### [PostgreSQL 12.1]

https://www.postgresql.org/docs/12/release-12-1.html

In addition, issues that occurred in previous versions are also fixed.

Refer to the following for details of the program fixes included in this version and level.

- FUJITSU Enterprise Postgres 12 Program Updates

# **FUJITSU Enterprise Postgres 12 Program Updates**

PH14715 A is PH15139 W PH15542 Ev th PH15823 W m PH15957 W DI PH15959 W PH16007 Tr	When the ecobpg command is excuted, the segmentation violations may occur.  In memory leak occurs when an SQL statement containing the output parameter is executed from COBOL application.  When the Mirroring Controller fails to start, the instance may stay up.  Even in the Mirroring Controller the automatic switch/disconnection is enabled, the automatic switch/disconnection may not occur when the server goes down.  When fatal or panic is specified in the client_min_messages parameter, SQL may not work properly.  When using the data masking feature, if a SELECT statement containing DISTINCT or UNION is executed, the result may contain duplicate rows.	Y Y Y Y Y	Y Y - - Y
PH15139 W PH15542 Ev th PH15823 W m PH15957 W DI PH15959 W PH16007 Tr	When the Mirroring Controller fails to start, the instance may stay up. Even in the Mirroring Controller the automatic switch/disconnection is enabled, the automatic switch/disconnection may not occur when the server goes down.  When fatal or panic is specified in the client_min_messages parameter, SQL may not work properly.  When using the data masking feature, if a SELECT statement containing	Y Y Y	-
PH15542 Event the PH15823 W m PH15957 W DI PH15959 W PH16007 Tr	Even ih the Mirroring Controller the automatic switch/disconnection is enabled, the automatic switch/disconnection may not occur when the server goes down.  When fatal or panic is specified in the client_min_messages parameter, SQL may not work properly.  When using the data masking feature, if a SELECT statement containing	Y	- - Y
th PH15823 W m PH15957 W DI PH15959 W PH16007 Tr	The automatic switch/disconnection may not occur when the server goes down.  When fatal or panic is specified in the client_min_messages parameter, SQL may not work properly.  When using the data masking feature, if a SELECT statement containing	Υ	- Y
PH15823 W m PH15957 W DI PH15959 W PH16007 Th	When fatal or panic is specified in the client_min_messages parameter, SQL nay not work properly.  When using the data masking feature, if a SELECT statement containing	Υ	Y
MPH15957 WDI PH15959 WPH16007 Tr	nay not work properly.  When using the data masking feature, if a SELECT statement containing		Υ
PH15959 W PH16007 Tr			
PH16007 Tr		Υ	Υ
	When using the data masking feature, the results may not be masked.	Υ	Υ
DL16164 1VV	the transaction log mirroring stops.	Υ	Υ
	When the embedded SQL COBOL applications specifed the TYPE command is recompiled, the precompile may fail.	Υ	Υ
	he result of executing the SQL in the COBOL application may be incorrect.	Υ	Υ
PH16167 W	When using the transparent data encryption feature, if the disk that stores emporary files is exhausted, SQL may not respond.	Υ	Υ
PH16280 W	When the functions created with procedural languages is executed, the atabase server may goes down.	Υ	Y
	pplications that use the .NET Data Provider feature cannot use the latest lpgsql.	Υ	Y
	Vrong attribute value is set in Npgsql's machine.config.	Υ	Υ
	he .NET application may fail to connect to database when uninstalling Vindows client in multi-version installation environment of Windows client.	Υ	Υ
	When the password for the instance administrator user includes + or %, the direction of the instance administrator user includes + or %, the direction of the password for the instance administrator user includes + or %, the direction of the password for the instance administrator user includes + or %, the direction of the password for the instance administrator user includes + or %, the direction of the password for the instance administrator user includes + or %, the direction of the password for the instance administrator user includes + or %, the direction of the password for the instance administrator user includes + or %, the direction of the password for the instance administrator user includes + or %, the direction of the password for th	Υ	-
	he pgx_dmpall command may fail to execute.	Υ	Υ
	When the password for the instance administrator user includes &, #, ', or  the direction of the instance administrator user includes &, #, ', or  the direction of the instance administrator user includes &, #, ', or  the direction of the instance administrator user includes &, #, ', or  the direction of the instance administrator user includes &, #, ', or  the direction of the instance administrator user includes &, #, ', or  the direction of the instance administrator user includes &, #, ', or  the direction of the instance administrator user includes &, #, ', or  the direction of the instance administrator user includes &, #, ', or  the direction of the instance administrator user includes &, #, ', or  the direction of the instance administrator user includes &, #, ', or  the direction of the instance administrator user includes &, #, ', or  the direction of the instance administrator user includes &, #, ', or  the direction of the instance administrator user includes &, #, ', or  the direction of the instance administrator user includes &, #, ', or  the direction of the instance administrator user includes &, #, ', or  the direction of the instance administrator user includes &, #, ', or  the direction of the instance administrator user includes &, #, ', or  the direction of the instance administrator user includes &, #, ', or  the direction of the instance administrator user includes &, #, ', or  the direction of the instance administrator user includes &, #, ', or  the direction of the instance administrator user includes &, #, ', or  the direction of the instance administrator user includes &, #, ', or  the direction of the instance administrator user includes &, #, ', or  the direction of the instance administrator user includes &, #, ', or  the direction of the instance administrator user includes &, #, ', or  the direction of the instance administrator user includes &, #, ', or  the direction of the instance administrator user includes &, #, ', or	Υ	-
PH16755 W	When the output file of the pg_dumpall command, is restored, the encrypted ablespace data may not be encrypted.	Υ	Υ
	When a new instance is created in WebAdmin, the instance may fail to start.	Υ	Υ
PH16800 Tr	he Mirroring Controller process may go down, stopping monitoring, and failing instance anomaly detection or state checking.	Υ	-
PH16942 Co	Corresponds to the host variable definition using level number 77 in the pplication using embedded SQL in COBOL.	Υ	Υ
PH16965 EV	efined at level number 49 or higher, no error occurs at the precompilation.	Υ	Υ
PH17289 O	On the [Client authentication] window of the WebAdmin, [Method] is not pdated, or the presence of a check box may be incorrect.	Υ	Υ
	n the WebAdmin, PostgreSQL configuration changes and instance import may	Υ	Υ
PH17403 It	is not possible to refer to the FUJITSU Enterprise Postgres online manual age using the man command.	Υ	Y
	racle_fdw can not be used with OCI library versions other than 12 and 19.	Υ	Υ
PH17903 Ex	eature may bring down the database.	Y	Y
PH17980 Fu	rujitsu Enterprise Postgres is incorporated the fixes for security failures that were absorbed in PostgreSQL 12.2.	Υ	Y
PH18020 U <sub>I</sub>	Updating tables that use the vertical clustered Index may bring down the atabase.	Υ	-
PH18206 Ex	executing SQL using a VCI index after the SET ROLE statement was executed	Υ	-
PH18207 In	nay result in an error.  ncorrect values may appear in the backend_type of the pg_stat_activity view	Υ	-
	nd the logs when using VCI.  Executing the pgx_loader command may cause the database to crash.	Υ	_



# FUJITSU Enterprise Postgres 12



Program Updates (Windows)

## **Preface**

This document explains the updates that have been fixed at this version.

The contents of this document are subject to change without notice.

#### **Notations**

The status for each edition is shown in the following table.

P number	P number Update summary AE		SE
Number that uniquely identifies the update	Summary of update details	Symbol in the fixed	υ

ΑE

Indicates Advanced Edition.

SE

Indicates Standard Edition.

#### Symbol

Indicates the incorporated status for each edition.

Y: Fixed

-: Not relevant to this product

#### **Export restrictions**

Exportation/release of this document may require necessary procedures in accordance with the regulations of your resident country and/or US export control laws.

#### Issue date and version

Edition 1.0: July 2020

#### Copyright

Copyright 2020 FUJITSU LIMITED

# Chapter 1 Program Updates

This version incorporates the updates implemented in PostgreSQL 12, and 12.1.



Refer to the PostgreSQL Global Development Group website for information on the updates implemented in the following releases:

.....

#### [PostgreSQL 12]

https://www.postgresql.org/docs/12/release-12.html

#### [PostgreSQL 12.1]

https://www.postgresql.org/docs/12/release-12-1.html

In addition, issues that occurred in previous versions are also fixed.

Refer to the following for details of the program fixes included in this version and level.

- FUJITSU Enterprise Postgres 12 Program Updates

# **FUJITSU Enterprise Postgres 12 Program Updates**

P number	Update summary	AE	SE	
PH14054	When the ecobpg command is excuted, the segmentation violations may occur.	Υ	Y	
PH14715	A memory leak occurs when an SQL statement containing the output parameter is executed from COBOL application.	Υ	Y	
PH15139	When the Mirroring Controller fails to start, the instance may stay up.	Υ	-	
PH15542	Even ih the Mirroring Controller the automatic switch/disconnection is enabled, the automatic switch/disconnection may not occur when the server goes down.	Υ	-	
PH15823	When fatal or panic is specified in the client_min_messages parameter, SQL may not work properly.	Υ	Y	
PH15957	When using the data masking feature, if a SELECT statement containing DISTINCT or UNION is executed, the result may contain duplicate rows.	Υ	Y	
PH15959	When using the data masking feature, the results may not be masked.	Υ	Υ	
PH16007	The transaction log mirroring stops.	Υ	Υ	
PH16164	When the embedded SQL COBOL applications specifed the TYPE command is precompiled, the precompile may fail.	Υ	Y	
PH16166	The result of executing the SQL in the COBOL application may be incorrect.	Υ	Υ	
PH16280	When the functions created with procedural languages is executed, the database server may goes down.	Υ	Y	
PH16485	Applications that use the .NET Data Provider feature cannot use the latest Npgsql.	Υ	Y	
PH16486	Wrong attribute value is set in Npgsql's machine.config.	Υ	Υ	
PH16488	The .NET application may fail to connect to database when uninstalling Windows client in multi-version installation environment of Windows client.	Y	Y	
PH16511	When the password for the instance administrator user includes + or %, the Mirroring Controller will fail to start.	Υ	-	
PH16546	The pgx dmpall command may fail to execute.	Υ	Υ	
PH16549	When the password for the instance administrator user includes &, #, ', or  the Mirroring Controller will fail to start.	Υ	-	
PH16755	When the output file of the pg_dumpall command, is restored, the encrypted tablespace data may not be encrypted.	Υ	Y	
PH16791	When a new instance is created in WebAdmin, the instance may fail to start.	Υ	Y	
PH16800	The Mirroring Controller process may go down, stopping monitoring, and failing instance anomaly detection or state checking.	Y	-	
PH16942	Corresponds to the host variable definition using level number 77 in the application using embedded SQL in COBOL.	Υ	Y	
PH16965	Even If the host variables PIC X (n) VARYING and PIC N (n) VARYING are defined at level number 49 or higher, no error occurs at the precompilation.	Υ	Y	
PH17289	On the [Client authentication] window of the WebAdmin, [Method] is not updated, or the presence of a check box may be incorrect.	Υ	Y	
PH17313	In the WebAdmin, PostgreSQL configuration changes and instance import may fail.	Υ	Υ	
PH17700	oracle fdw can not be used with OCI library versions other than 12 and 19.	Υ	Υ	
PH17903	Executing an EXISTS subquery expression when using the data masking feature may bring down the database.	Υ	Υ	
PH17980	Fujitsu Enterprise Postgres is incorporated the fixes for security failures that were absorbed in PostgreSQL 12.2.	Y	Y	
PH18020	Updating tables that use the vertical clustered Index may bring down the database.	Υ	-	
PH18206	Executing SQL using a VCI index after the SET ROLE statement was executed may result in an error.	Y	-	
PH18207	Incorrect values may appear in the backend_type of the pg_stat_activity view and the logs when using VCI.	Υ	-	
PH18486	Executing the pgx_loader command may cause the database to crash.	Υ	-	



# FUJITSU Enterprise Postgres 12



# Software Release Notes

Linux > Windows >

This document lists special cautions and reference information relating to the use of this product.

#### [High Risk Activity]

The Customer acknowledges and agrees that the Product is designed, developed and manufactured as contemplated for general use, includfsing without limitation, general office use, personal use, household use, and ordinary industrial use, but is not designed, developed and manufactured as contemplated for use accompanying fatal risks or dangers that, unless extremely high safety is secured, could lead directly to death, personal injury, severe physical damage or other loss (hereinafter "High Safety Required Use"), including without limitation, nuclear reaction control in nuclear facilities, aircraft flight control, air traffic control, mass transport control, medical life support systems, and missile launch control in weapon systems. The Customer, shall not use the Product without securing the sufficient safety required for the High Safety Required Use. In addition, Fujitsu (or other affiliate's name) shall not be liable against the Customer and/or any third party for any claims or damages arising in connection with the High Safety Required Use of the Product.

#### [Export Restrictions]

Exportation/release of this document may require necessary procedures in accordance with the regulations of your resident country and/or US export control laws.

#### [Trademarks]

- Microsoft, Windows, Windows Server and other Microsoft product titles or names are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries.
- Linux(R) is a registered trademark of Linus Torvalds in the U.S. and other countries.
- Red Hat, RPM, and all Red Hat-based trademarks and logos are trademarks or registered trademarks of Red Hat, Inc. in the United States and other countries.
- SUSE and the SUSE logo are registered trademarks of SUSE LLC in the United States and other countries.
- Oracle and Java are registered trademarks of Oracle and/or its affiliates.
- Company names and product names used in this document are trademarks or registered trademarks of the respective companies.

The content of this manual are subject to change without notice.

#### [Abbreviations and Generic Terms Used]

Note that product names used in this document are shown as follows:

- Windows(R) 8.1 Windows (R) 8.1 Pro Windows (R) 8.1 Enterprise Windows (R) 10 Home Windows (R) 10 Education Windows(R) 10 Pro Windows (R) 10 Enterprise Microsoft(R) Windows Server(R) 2012 Datacenter Microsoft(R) Windows Server(R) 2012 Standard Microsoft(R) Windows Server(R) 2012 Essentials Microsoft(R) Windows Server(R) 2012 Foundation Microsoft(R) Windows Server(R) 2012 R2 Datacenter Microsoft(R) Windows Server(R) 2012 R2 Standard Microsoft(R) Windows Server(R) 2012 R2 Essentials Microsoft(R) Windows Server(R) 2012 R2 Foundation Microsoft(R) Windows Server(R) 2016 Datacenter Microsoft(R) Windows Server(R) 2016 Standard Microsoft(R) Windows Server(R) 2016 Essentials Microsoft(R) Windows Server(R) 2019 Datacenter Microsoft(R) Windows Server(R) 2019 Standard Microsoft(R) Windows Server(R) 2019 Essentials are abbreviated as "Windows (R)".

- Red Hat(R) Enterprise Linux(R) 7
Red Hat(R) Enterprise Linux(R) 8.1
SUSE Linux Enterprise Server 12
are abbreviated as "Linux".

Copyright FUJITSU LIMITED 2015-2020

Revision History

<Edition 1.0 (July 2020)>

This document is structured as follows:

- 1 Component Programs
- 2 Applicable Manuals
  - 2.1 Online Manuals
- 3 Software Environment
  - 3.1 Basic Software
  - J.I Dasic Solitware
  - 3.2 Required Software
  - 3.3 Update Information
  - 3.4 Mutually Exclusive Software
- 4 Hardware Environment
- 5 Restrictions
- 6 Notes on Migration
- 7 Other Notes
- 8 Installation Procedure
  - 8.1 Installation Method
  - 8.2 Uninstallation Method
- 9 FUJITSU Enterprise Postgres Client
- 10 About streaming replication

Appendix A License Terms and Exclusions of Liability

- A.1 License for Open Source Software
- A.2 Provision of the Source Code for Software Based on LGPL
- A.3 Programs for which Modification and Reverse Engineering are Permitted
- A.4 Apache License
- A.5 Python

#### 1 Component Programs

This product comes without any bundled optional products.

#### 2 Applicable Manuals

#### 2.1 Online Manuals

The manuals listed in the table below are included on the "Server program/Online manuals" DVD. You can read the manuals directly from the DVD or copy them to your computer. If you copy them to your computer, do not change the folder structure or file names because this can affect the correctness of how the manual is displayed and the accuracy of the PDF search function.

Note that Adobe(R) Reader(R) is required in order to read and search manuals in PDF format.

FUJITSU Enterprise Postgres manuals contain various notation conventions and rules.

When using FUJITSU Enterprise Postgres manuals, keep these rules in mind as you read them.

+----+

Also, refer to "FUJITSU Enterprise Postgres Documentation Road Map" for information on these rules.

		Manual T	itle							
	1	FUJITSU	Enterprise	_	Release Note					 
	2	FUJITSU	Enterprise	Postgres	Documentation	n Roa	ad Map			 
		FUJITSU	-	Postgres	General Desc	ripti	ion			 
				Postgres	Installation	and	Setup			
	5	FUJITSU	Enterprise	Postgres	Installation	and	Setup	Guide	for	Client
	6 l	FUJITSU	Enterprise		Installation					'

	Assistant
7	FUJITSU Enterprise Postgres Operation Guide
8	FUJITSU Enterprise Postgres Cluster Operation Guide (Database Multiplexing)
9	FUJITSU Enterprise Postgres Cluster Operation Guide (PRIMECLUSTER)
10	FUJITSU Enterprise Postgres Security Operation Guide
11	FUJITSU Enterprise Postgres Application Development Guide
12	FUJITSU Enterprise Postgres Connection Manager User's Guide
13	FUJITSU Enterprise Postgres Reference
14	FUJITSU Enterprise Postgres Java API Reference
15	FUJITSU Enterprise Postgres Npgsql API Reference
16	FUJITSU Enterprise Postgres Glossary
17	FUJITSU Enterprise Postgres Messages
18	FJQSS (Information Collection Tool) User's Guide
19	PostgreSQL 12.1 Documentation (*1)
*1) Thi	s is included as reference material.

\*1) This is included as reference material.

For the latest version, refer to the "The PostgreSQL Global Development Group" website (http://www.postgresql.org/).

A copy of this reference material, downloaded from the above website on December 19 2019, is included with this product.

#### 2.1.1 Reading Manuals Directly from the DVD

#### 3 Software Environment

#### 3.1 Basic Software

For information on the basic software that can be installed by this software, refer to the following chapter in the FUJITSU Enterprise Postgres Installation and Setup Guide for Server in the online manuals:

"Chapter 2 Operating Environment"

#### 3.2 Required Software

For information on the software required to install this software, refer to the following chapter in the FUJITSU Enterprise Postgres Installation and Setup Guide for Server in the online manuals:

"Chapter 2 Operating Environment"

#### 3.3 Update Information

For information on the required modifications to the operation of this software, refer to the following chapter in the FUJITSU Enterprise Postgres Installation and Setup Guide for Server in the online manuals:

"Chapter 2 Operating Environment"

#### 3.4 Mutually Exclusive Software

For information on the software with which this software is mutually exclusive, refer to the following chapter in the FUJITSU Enterprise Postgres Installation and Setup Guide for Server in the online manuals: "Chapter 2 Operating Environment"

#### 4 Hardware Environment

For information on the hardware required to install this software, refer to the following chapter in the FUJITSU Enterprise Postgres Installation and

Setup Guide for Server in the online manuals: "Chapter 2 Operating Environment"

#### 5 Restrictions

None.

#### 6. Notes on Migration

If an existing 9.4 or newer instance is going to be used in this version, perform pg\_upgrade in advance. Refer to "Notes on Upgrading Database Instances" in the Operation Guide for details.

Refer to "pg\_upgrade" in the PostgreSQL Documentation for information on using pg upgrade.

#### 7. Other Notes

None.

#### 8 Installation Procedure

#### 8.1 Installation Method

For information on installing this software, refer to the following chapters in the FUJITSU Enterprise Postgres Installation and Setup Guide for Server in the online manuals:

"Chapter 3 Installation"

#### 8.2 Uninstallation Method

For information on uninstalling this software, refer to the following chapter in the FUJITSU Enterprise Postgres Installation and Setup Guide for Server in the online manuals:

"Chapter 5 Uninstallation"

#### 9 FUJITSU Enterprise Postgres Client

This software includes FUJITSU Enterprise Postgres Client for the platforms listed below.

If an application is used in a remote environment, install FUJITSU Enterprise Postgres Client for each platform.

- Windows (R) FUJITSU Enterprise Postgres Client 12
- Linux FUJITSU Enterprise Postgres Client 12

#### 10. About streaming replication

Use streaming replication only with features compatible with PostgreSQL 12.

Appendix A License Terms and Exclusions of Liability

#### A.1. License for Open Source Software

This product uses the open-source software.

Refer to OSS Licence/OSS List pdf in this m

Refer to OSS\_Licence/OSS\_List.pdf in this media for the name, version and license conditions of the open source software included in this product. After providing this product and this media, the OSS version and license conditions may change due to the provision of the correction program. Refer to the following folder for the changed OSS version and license conditions.

```
+-----+
| <Installation destination of the product>/doc/Licence/ |
+------+
```

It is allowable to choose a license from one of the following for jQuery and jQuery UI Library; MIT License or GNU GENERAL PUBLIC LICENSE (GNU GPL). This product redistributes jQuery and jQuery UI Library based on the MIT License.

The ecj is covered by the Eclipse Public License.

We will provide, on request, the source code for software for which source code disclosure is a license requirement. If you require this source code, please contact our sales staff.

The jstl is covered by the COMMON DEVELOPMENT AND DISTRIBUTION LICENSE (CDDL). We will provide, on request, the source code for CDDL target software and for software for which source code disclosure is a license requirement. If you require this source code, please contact our sales staff.

Will Hartung - getline is license-free open source software available in the public domain.

A.2 Provision of the Source Code for Software Based on LGPL

This product comes with the following software licensed under the GNU LESSER GENERAL PUBLIC LICENSE (GNU LGPL) V2.1, GNU LESSER GENERAL PUBLIC LICENSE Version 3, and GNU LIBRARY GENERAL PUBLIC LICENSE Version 2. The source code of the software below will be provided on request at cost for a period of three years from the date of shipment of this product, or for the term of the maintenance support agreement entered into for this product, whichever is longer.

- gettext
- libiconv
- psqlODBCqt

Contact Fujitsu technical support if you require the source code.

A.3 Programs for which Modification and Reverse Engineering are Permitted

Modification and reverse engineering are permissible for the following programs included in this product to the extent permitted under the GNU LESSER GENERAL PUBLIC LICENSE V2.1, GNU LESSER GENERAL PUBLIC LICENSE Version 3, and the GNU LIBRARY GENERAL PUBLIC LICENSE Version 2, which are applied to the open source software used together with this product.

#### Folder names

- <FUJITSU Enterprise Postgres install dir>/bin
- <FUJITSU Enterprise Postgres install dir>/lib
- <Linux FUJITSU Enterprise Postgres Client install dir>/bin
- <Linux FUJITSU Enterprise Postgres Client install dir>/lib
- <Linux FUJITSU Enterprise Postgres Client install dir>/odbc

#### A.4 Apache License

This product includes the following software that are covered by Apache License, Version 2.0

- Modified source code of software Shaj
- apache commons io
- apache commons lang
- apache commons logging
- apache dbcp
- apache derby
- Apache Tomcat
- commons-codec
- commons-configuration
- commons-pool
- Jackson-annotations
- Jackson-core
- Jackson-databind
- javax-inject
- log4j
- spring framework
- spring-security-web

For the license statement of Apache License, Version 2.0, refer to the OSS Licence folder in this media.

#### A.5 Python

This product includes the following Python libraries.

- -ip - alembic 19.2.3 1.4.2

_	appdirs	1.4.3
_	Babel	2.8.0
-	bcrypt	3.1.7
-	blinker	1.4
-	certifi	2020.4.5.1
-	cffi	1.14.0
-	chardet	3.0.4
-	click	7.1.1
-	cryptography	2.9
-	distlib	0.3.0
-	filelock	3.0.12
-	Flask	1.0.2
-	Flask-BabelEx	0.9.4
-	Flask-Compress	1.4.0
-	Flask-Gravatar	0.5.0
-	Flask-Login	0.4.1
-	Flask-Mail	0.9.1
-	Flask-Migrate	2.4.0
-	Flask-Paranoid	0.2.0
-	Flask-Principal	0.4.0
-	Flask-Security	3.0.0
-	Flask-SQLAlchemy	2.3.2
-	Flask-WTF	0.14.3
_	idna	2.9 1.5.0
_	<pre>importlib-metadata itsdangerous</pre>	1.1.0
_	Jinja2	2.11.1
_	Mako	1.1.2
_	MarkupSafe	1.1.1
_	packaging	20.3
_	paramiko	2.7.1
_	passlib	1.7.1
_	pip	20.0.2
_	psutil	5.5.1
-	psycopg2	2.8.5
-	pycparser	2.20
-	PyNaCl	1.3.0
-	pyparsing	2.4.7
-	python-dateutil	2.8.1
-	python-editor	1.0.4
	pytz	2018.9
	setuptools	46.0.0
_	simplejson	3.16.0
_	six analklator	1.14.0
_	speaklater	
_	sphinxcontrib-applehelp sphinxcontrib-devhelp	1.0.2 1.0.2
_	sphinxcontrib-devnerp sphinxcontrib-htmlhelp	1.0.3
_	sphinxcontrib-jsmath	1.0.1
_	sphinxcontrib-qthelp	1.0.3
_	sphinxcontrib qualizinghtml	1.1.4
_	SQLAlchemy	1.3.15
_	sqlparse	0.2.4
	sshtunnel	0.1.5
	urllib3	1.25.8
	virtualenv	20.0.10
-	Werkzeug	1.0.1
-	wheel	0.34.2
-	WTForms	2.2.1
-	zipp	3.1.0

Refer to the files below for the license statement of each library.

<Windows FUJITSU Enterprise Postgres Client install dir>

\pgadmin4\venv\Lib\site-packages\bcrypt-3.1.7.dist-info\LICENSE - blinker <Windows FUJITSU Enterprise Postgres Client install dir> \pgadmin4\venv\Lib\site-packages\blinker-1.4.dist-info\LICENSE <Windows FUJITSU Enterprise Postgres Client install dir> \pgadmin4\venv\Lib\site-packages\certifi-2020.4.5.1.dist-info\LICENSE - cffi <Windows FUJITSU Enterprise Postgres Client install dir> \pgadmin4\venv\Lib\site-packages\cffi-1.14.0.dist-info\LICENSE - chardet <Windows FUJITSU Enterprise Postgres Client install dir>  $\verb|\pgadmin4| venv \perp b | site-packages \mid chardet-3.0.4. \\ dist-info \mid METADATA| \\$ click <Windows FUJITSU Enterprise Postgres Client install dir> \pgadmin4\venv\Lib\site-packages\click-7.1.1.dist-info\LICENSE.rst cryptography <Windows FUJITSU Enterprise Postgres Client install dir> \pgadmin4\venv\Lib\site-packages\cryptography-2.9.dist-info\LICENSE distlib <Windows FUJITSU Enterprise Postgres Client install dir> \pgadmin4\venv\Lib\site-packages\distlib-0.3.0-py3.7.egg-info\PKG-INFO - filelock <Windows FUJITSU Enterprise Postgres Client install dir> \pgadmin4\venv\Lib\site-packages\filelock-3.0.12.dist-info\LICENSE - Flask <Windows FUJITSU Enterprise Postgres Client install dir> \pgadmin4\venv\Lib\site-packages\Flask-1.0.2.dist-info\LICENSE.txt - Flask-BabelEx <Windows FUJITSU Enterprise Postgres Client install dir> \pgadmin4\venv\Lib\site-packages\Flask BabelEx-0.9.4.dist-info\LICENSE - Flask-Compress <Windows FUJITSU Enterprise Postgres Client install dir> \pgadmin4\venv\Lib\site-packages\Flask Compress-1.4.0.dist-info\LICENSE.txt - Flask-Gravatar <Windows FUJITSU Enterprise Postgres Client install dir> \pgadmin4\venv\Lib\site-packages\Flask Gravatar-0.5.0.dist-info\METADATA - Flask-Login <Windows FUJITSU Enterprise Postgres Client install dir> \pgadmin4\venv\Lib\site-packages\Flask Login-0.4.1.dist-info\LICENSE - Flask-Mail <Windows FUJITSU Enterprise Postgres Client install dir> \pgadmin4\venv\Lib\site-packages\Flask Mail-0.9.1.dist-info\LICENSE - Flask-Migrate <Windows FUJITSU Enterprise Postgres Client install dir> \pgadmin4\venv\Lib\site-packages\Flask Migrate-2.4.0.dist-info\LICENSE - Flask-Paranoid <Windows FUJITSU Enterprise Postgres Client install dir> \pgadmin4\venv\Lib\site-packages\Flask\_Paranoid-0.2.0.dist-info\METADATA - Flask-Principal <Windows FUJITSU Enterprise Postgres Client install dir> \pgadmin4\venv\Lib\site-packages\Flask Principal-0.4.0.dist-info\METADATA - Flask-Security <Windows FUJITSU Enterprise Postgres Client install dir> \pgadmin4\venv\Lib\site-packages\Flask Security-3.0.0.dist-info\METADATA - Flask-SQLAlchemy <Windows FUJITSU Enterprise Postgres Client install dir> \pgadmin4\venv\Lib\site-packages\Flask SQLAlchemy-2.3.2.dist-info \LICENSE.txt - Flask-WTF <Windows FUJITSU Enterprise Postgres Client install dir> \pgadmin4\venv\Lib\site-packages\Flask WTF-0.14.3.dist-info\LICENSE.txt - idna <Windows FUJITSU Enterprise Postgres Client install dir> \pgadmin4\venv\Lib\site-packages\idna-2.9.dist-info\LICENSE.rst importlib-metadata <Windows FUJITSU Enterprise Postgres Client install dir> \pgadmin4\venv\Lib\site-packages\importlib\_metadata-1.5.0.dist-info\LICENSE - itsdangerous <Windows FUJITSU Enterprise Postgres Client install dir> \pgadmin4\venv\Lib\site-packages\itsdangerous-1.1.0.dist-info\LICENSE.rst - Jinja2 <Windows FUJITSU Enterprise Postgres Client install dir> \pgadmin4\venv\Lib\site-packages\Jinja2-2.11.1.dist-info\LICENSE.rst <Windows FUJITSU Enterprise Postgres Client install dir> \pgadmin4\venv\Lib\site-packages\Mako-1.1.2.dist-info\LICENSE - MarkupSafe

<Windows FUJITSU Enterprise Postgres Client install dir> \pgadmin4\venv\Lib\site-packages\MarkupSafe-1.1.1.dist-info\LICENSE.rst packaging <Windows FUJITSU Enterprise Postgres Client install dir> \pgadmin4\venv\Lib\site-packages\packaging-20.3.dist-info\LICENSE - paramiko <Windows FUJITSU Enterprise Postgres Client install dir> \pgadmin4\venv\Lib\site-packages\paramiko-2.7.1.dist-info\LICENSE.txt <Windows FUJITSU Enterprise Postgres Client install dir> \pgadmin4\venv\Lib\site-packages\passlib-1.7.1.dist-info\METADATA <Windows FUJITSU Enterprise Postgres Client install dir> \pgadmin4\venv\Lib\site-packages\pip-20.0.2.dist-info\LICENSE.txt - psutil <Windows FUJITSU Enterprise Postgres Client install dir> \pgadmin4\venv\Lib\site-packages\psutil-5.5.1.dist-info\LICENSE psycopg2 <Windows FUJITSU Enterprise Postgres Client install dir> \pgadmin4\venv\Lib\site-packages\psycopg2-2.8.5.dist-info\LICENSE pycparser <Windows FUJITSU Enterprise Postgres Client install dir> \pgadmin4\venv\Lib\site-packages\pycparser-2.20.dist-info\LICENSE - PyNaCl <Windows FUJITSU Enterprise Postgres Client install dir> \pgadmin4\venv\Lib\site-packages\PyNaCl-1.3.0.dist-info\LICENSE.txt pyparsing <Windows FUJITSU Enterprise Postgres Client install dir> \pgadmin4\venv\Lib\site-packages\pyparsing-2.4.7.dist-info\LICENSE - python-dateutil <Windows FUJITSU Enterprise Postgres Client install dir>  $\verb|\pgadmin4| venv\Lib\site-packages\python_date util-2.8.1.dist-info\LICENSE| \\$ - python-editor <Windows FUJITSU Enterprise Postgres Client install dir> \pgadmin4\venv\Lib\site-packages\python editor-1.0.4.dist-info\METADATA - pytz <Windows FUJITSU Enterprise Postgres Client install dir> \pgadmin4\venv\Lib\site-packages\pytz-2018.9.dist-info\LICENSE.txt - setuptools <Windows FUJITSU Enterprise Postgres Client install dir> \pgadmin4\venv\Lib\site-packages\setuptools-46.0.0.dist-info\LICENSE - simplejson <Windows FUJITSU Enterprise Postgres Client install dir> \pgadmin4\venv\Lib\site-packages\simplejson-3.16.0.dist-info\LICENSE.txt - six <Windows FUJITSU Enterprise Postgres Client install dir> \pgadmin4\venv\Lib\site-packages\six-1.14.0.dist-info\LICENSE speaklater <Windows FUJITSU Enterprise Postgres Client install dir> \pgadmin4\venv\Lib\site-packages\speaklater-1.3.dist-info\LICENSE - sphinxcontrib-applehelp <Windows FUJITSU Enterprise Postgres Client install dir> \pgadmin4\venv\Lib\site-packages\sphinxcontrib applehelp-1.0.2.dist-info \LICENSE - sphinxcontrib-devhelp <Windows FUJITSU Enterprise Postgres Client install dir> \pgadmin4\venv\Lib\site-packages\sphinxcontrib\_devhelp-1.0.2.dist-info \LICENSE - sphinxcontrib-htmlhelp <Windows FUJITSU Enterprise Postgres Client install dir> \pgadmin4\venv\Lib\site-packages\sphinxcontrib\_htmlhelp-1.0.3.dist-info \LICENSE - sphinxcontrib-jsmath <Windows FUJITSU Enterprise Postgres Client install dir> \pgadmin4\venv\Lib\site-packages\sphinxcontrib jsmath-1.0.1.dist-info - sphinxcontrib-qthelp <Windows FUJITSU Enterprise Postgres Client install dir> \pgadmin4\venv\Lib\site-packages\sphinxcontrib qthelp-1.0.3.dist-info \LICENSE - sphinxcontrib-serializinghtml <Windows FUJITSU Enterprise Postgres Client install dir> \pgadmin4\venv\Lib\site-packages \sphinxcontrib serializinghtml-1.1.4.dist-info\LICENSE SQLAlchemy <Windows FUJITSU Enterprise Postgres Client install dir> \pgadmin4\venv\Lib\site-packages\SQLAlchemy-1.3.15.dist-info\LICENSE sqlparse

<Windows FUJITSU Enterprise Postgres Client install dir>
\pgadmin4\venv\Lib\site-packages\sqlparse-0.2.4.dist-info\METADATA

- sshtunnel

<Windows FUJITSU Enterprise Postgres Client install dir>
\pgadmin4\venv\Lib\site-packages\sshtunnel-0.1.5.dist-info\LICENSE

- urllib3

<Windows FUJITSU Enterprise Postgres Client install dir>
\pgadmin4\venv\Lib\site-packages\urllib3-1.25.8.dist-info\LICENSE.txt

- virtualenv

<Windows FUJITSU Enterprise Postgres Client install dir>
\pgadmin4\venv\Lib\site-packages\virtualenv-20.0.10.dist-info\LICENSE

- Werkzeug

<Windows FUJITSU Enterprise Postgres Client install dir>
\pgadmin4\venv\Lib\site-packages\Werkzeug-1.0.1.dist-info\LICENSE.rst

- wheel

<Windows FUJITSU Enterprise Postgres Client install dir>
\pgadmin4\venv\Lib\site-packages\wheel-0.34.2.dist-info\LICENSE.txt

- WTForms

<Windows FUJITSU Enterprise Postgres Client install dir>
\pgadmin4\venv\Lib\site-packages\WTForms-2.2.1.dist-info\METADATA

- zipp

<Windows FUJITSU Enterprise Postgres Client install dir>
\pgadmin4\venv\Lib\site-packages\zipp-3.1.0.dist-info\LICENSE

\*\*\*

\*\*

Software Release Notes

\*\*

\*\*

Windows FUJITSU Software Enterprise Postgres Advanced Edition (64bit)

\*\*

\*\*

12

\*\*\*\*\*\*\*\*\*\*\*\*\*\*\*\*\*

This document lists special cautions and reference information relating to the use of this product.

#### [High Risk Activity]

The Customer acknowledges and agrees that the Product is designed, developed and manufactured as contemplated for general use, including without limitation, general office use, personal use, household use, and ordinary industrial use, but is not designed, developed and manufactured as contemplated for use accompanying fatal risks or dangers that, unless extremely high safety is secured, could lead directly to death, personal injury, severe physical damage or other loss (hereinafter "High Safety Required Use"), including without limitation, nuclear reaction control in nuclear facilities, aircraft flight control, air traffic control, mass transport control, medical life support systems, and missile launch control in weapon systems. The Customer, shall not use the Product without securing the sufficient safety required for the High Safety Required Use. In addition, Fujitsu (or other affiliate's name) shall not be liable against the Customer and/or any third party for any claims or damages arising in connection with the High Safety Required Use of the Product.

#### [Export Restrictions]

Exportation/release of this document may require necessary procedures in accordance with the regulations of your resident country and/or US export control laws.

#### [Trademarks]

- Microsoft, Windows, Windows Server and other Microsoft product titles or names are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries.
- Linux(R) is a registered trademark of Linus Torvalds in the U.S. and other countries.
- Red Hat, RPM, and all Red Hat-based trademarks and logos are trademarks or registered trademarks of Red Hat, Inc. in the United States and other countries.
- SUSE and the SUSE logo are registered trademarks of SUSE LLC in the United States and other countries.
- Oracle and Java are registered trademarks of Oracle and/or its affiliates.
- Company names and product names used in this document are trademarks or registered trademarks of the respective companies.

The content of this manual are subject to change without notice.

#### [Abbreviations and Generic Terms Used]

Note that product names used in this document are shown as follows:

Microsoft(R) Windows Server(R) 2012 Datacenter Microsoft(R) Windows Server(R) 2012 Standard Microsoft(R) Windows Server(R) 2012 Essentials Microsoft(R) Windows Server(R) 2012 Foundation Microsoft(R) Windows Server(R) 2012 R2 Datacenter Microsoft(R) Windows Server(R) 2012 R2 Standard Microsoft(R) Windows Server(R) 2012 R2 Essentials Microsoft(R) Windows Server(R) 2012 R2 Foundation

are abbreviated as "Windows Server(R) 2012".
- Microsoft(R) Windows Server(R) 2016 Datacenter
Microsoft(R) Windows Server(R) 2016 Standard

Microsoft(R) Windows Server(R) 2016 Essentials are abbreviated as "Windows Server(R) 2016".

- Microsoft(R) Windows Server(R) 2019 Datacenter
Microsoft(R) Windows Server(R) 2019 Standard
Microsoft(R) Windows Server(R) 2019 Essentials
are abbreviated as "Windows Server(R) 2019".

- Windows(R) 8.1 Windows(R) 8.1 Pro Windows(R) 8.1 Enterprise Windows(R) 10 Home

Windows(R) 10 Education

```
Windows (R) 10 Pro
   Windows (R) 10 Enterprise
   Microsoft(R) Windows Server(R) 2012 Datacenter
   Microsoft(R) Windows Server(R) 2012 Standard
   Microsoft(R) Windows Server(R) 2012 Essentials
   Microsoft(R) Windows Server(R) 2012 Foundation
   Microsoft(R) Windows Server(R) 2012 R2 Datacenter
   Microsoft(R) Windows Server(R) 2012 R2 Standard Microsoft(R) Windows Server(R) 2012 R2 Essentials
   Microsoft(R) Windows Server(R) 2012 R2 Foundation
   Microsoft(R) Windows Server(R) 2016 Datacenter
   Microsoft(R) Windows Server(R) 2016 Standard
   Microsoft(R) Windows Server(R) 2016 Essentials
   Microsoft(R) Windows Server(R) 2019 Datacenter
   Microsoft(R) Windows Server(R) 2019 Standard
   Microsoft(R) Windows Server(R) 2019 Essentials
   are abbreviated as "Windows (R)".
  - Red Hat(R) Enterprise Linux(R) 7
    Red Hat(R) Enterprise Linux(R) 8.1
    SUSE Linux Enterprise Server 12
    are abbreviated as "Linux".
Copyright FUJITSU LIMITED 2015-2020
Revision History
<Edition 1.0 (July 2020)>
This document is structured as follows:
  1 Component Programs
  2 Applicable Manuals
    2.1 Online Manuals
  3 Software Environment
    3.1 Basic Software
    3.2 Required Software
    3.3 Update Information
    3.4 Mutually Exclusive Software
  4 Hardware Environment
  5 Restrictions
  6 Notes on Migration
  7 Other Notes
  8 Installation Procedure
    8.1 Installation Method
    8.2 Uninstallation Method
  9 FUJITSU Enterprise Postgres Client
  10 About streaming replication
  Appendix A License Terms and Exclusions of Liability
    A.1 License for Open Source Software
   A.2 Provision of the Source Code for Software Based on LGPL
   A.3 Provision of the Source Code for Software Based on
       Microsoft Reciprocal License
   A.4 Programs for which Modification and Reverse Engineering are Permitted
   A.5 Apache License
   A.6 Python
   A.7 Microsoft Visual C++ 2010 Redistributable Package
   A.8 Microsoft Visual C++ 2013 Redistributable Package
   A.9 Microsoft Visual C++ 2015-2019 Redistributable Package
  Appendix B Viewing FJQSS and FUJITSU Enterprise Postgres HTML Files
```

#### 1 Component Programs

This product comes without any bundled optional products.

#### 2 Applicable Manuals

#### 2.1 Online Manuals

The manuals listed in the table below are included on the "Server program/Online manuals" DVD. You can read the manuals directly from the DVD or copy them to your computer. If you copy them to your computer, do not change the folder structure or file names because this can affect the correctness of how the manual is displayed and the accuracy of the PDF search function.

Note that Adobe(R) Reader(R) is required in order to read and search manuals in PDF format.

FUJITSU Enterprise Postgres manuals contain various notation conventions and rules.

When using FUJITSU Enterprise Postgres manuals, keep these rules in mind as you read them.

Also, refer to "FUJITSU Enterprise Postgres Documentation Road Map" for information on these rules.

	·
	Manual Title
1	FUJITSU Enterprise Postgres Release Notes
2	FUJITSU Enterprise Postgres Documentation Road Map
3	FUJITSU Enterprise Postgres General Description
4	FUJITSU Enterprise Postgres Installation and Setup Guide for Server
5	FUJITSU Enterprise Postgres Installation and Setup Guide for Client
6	   FUJITSU Enterprise Postgres Installation and Setup Guide for Server   Assistant
7	FUJITSU Enterprise Postgres Operation Guide
8	FUJITSU Enterprise Postgres Cluster Operation Guide   (Database Multiplexing)
9	FUJITSU Enterprise Postgres Security Operation Guide
10	FUJITSU Enterprise Postgres Application Development Guide
11	FUJITSU Enterprise Postgres Reference
12	FUJITSU Enterprise Postgres Java API Reference
13	FUJITSU Enterprise Postgres Npgsql API Reference
14	FUJITSU Enterprise Postgres Glossary
1	FUJITSU Enterprise Postgres Messages
15	

For the latest version, refer to the "The PostgreSQL Global Development Group" website (http://www.postgresql.org/). A copy of this reference material, downloaded from the above website on December 19 2019, is included with this product.

#### 2.1.1 Reading Manuals Directly from the DVD

Insert the DVD in the DVD drive.
Open the file below in a Web browser.

| DVD Drive Name:\manual\en\manual.html |

#### 3 Software Environment

#### 3.1 Basic Software

For information on the basic software that can be installed by this software, refer to the following chapter in the FUJITSU Enterprise Postgres Installation and Setup Guide for Server in the online manuals:

"Chapter 2 Operating Environment"

#### 3.2 Required Software

For information on the software required to install this software, refer to the following chapter in the FUJITSU Enterprise Postgres Installation and Setup Guide for Server in the online manuals:

"Chapter 2 Operating Environment"

#### 3.3 Update Information

For information on the required modifications to the operation of this software, refer to the following chapter in the FUJITSU Enterprise Postgres Installation and Setup Guide for Server in the online manuals: "Chapter 2 Operating Environment"

#### 3.4 Mutually Exclusive Software

For information on the software with which this software is mutually exclusive, refer to the following chapter in the FUJITSU Enterprise Postgres Installation and Setup Guide for Server in the online manuals: "Chapter 2 Operating Environment"

#### 4 Hardware Environment

For information on the hardware required to install this software, refer to the following chapter in the FUJITSU Enterprise Postgres Installation and Setup Guide for Server in the online manuals: "Chapter 2 Operating Environment"

#### 5 Restrictions

None.

#### 6. Notes on Migration

If an existing 9.4 or newer instance is going to be used in this version, perform pg\_upgrade in advance. Refer to "Notes on Upgrading Database Instances" in the Operation Guide for details.

Refer to "pg\_upgrade" in the PostgreSQL Documentation for information on using pg upgrade.

#### 7. Other Notes

For notes on this software, refer to the following chapter in the FUJITSU Enterprise Postgres Installation and Setup Guide for Server in the online manuals:

"Chapter 3 Installation"

#### 8 Installation Procedure

#### 8.1 Installation Method

For information on installing this software, refer to the following chapters in the FUJITSU Enterprise Postgres Installation and Setup Guide for Server in the online manuals:

"Chapter 3 Installation"

#### 8.2 Uninstallation Method

For information on uninstalling this software, refer to the following chapter in the FUJITSU Enterprise Postgres Installation and Setup Guide for Server in the online manuals:
 "Chapter 5 Uninstallation"

#### 9 FUJITSU Enterprise Postgres Client

This software includes FUJITSU Enterprise Postgres Client for the platforms listed below.

If an application is used in a remote environment, install FUJITSU Enterprise Postgres Client for each platform.

- Windows (R) FUJITSU Enterprise Postgres Client 12
- Linux FUJITSU Enterprise Postgres Client 12

#### 10. About streaming replication

Use streaming replication only with features compatible with PostgreSQL 12.

Appendix A License Terms and Exclusions of Liability

#### A.1 License for Open Source Software

This product uses the open-source software.

Refer to OSS\_Licence/OSS\_List.pdf in this media for the name, version and license conditions of the open source software included in this product. After providing this product and this media, the OSS version and license conditions may change due to the provision of the correction program. Refer to the following folder for the changed OSS version and license conditions.

+-----+ | <Installation destination of the product>/doc/Licence/ | +------+

It is allowable to choose a license from one of the following for jQuery and jQuery UI Library; MIT License or GNU GENERAL PUBLIC LICENSE (GNU GPL). This product redistributes jQuery and jQuery UI Library based on the MIT License.

The ecj is covered by the Eclipse Public License.

We will provide, on request, the source code for software for which source code disclosure is a license requirement. If you require this source code, please contact our sales staff.

The jstl is covered by the COMMON DEVELOPMENT AND DISTRIBUTION LICENSE (CDDL). We will provide, on request, the source code for CDDL target software and for software for which source code disclosure is a license requirement. If you require this source code, please contact our sales staff.

Will Hartung - getline is license-free open source software available in the public domain.

#### A.2 Provision of the Source Code for Software Based on LGPL

This product comes with the following software licensed under the GNU LESSER GENERAL PUBLIC LICENSE (GNU LGPL) V2.1, GNU LESSER GENERAL PUBLIC LICENSE Version 3, and GNU LIBRARY GENERAL PUBLIC LICENSE Version 2. The source code of the software below will be provided on request at cost for a period of three years from the date of shipment of this product, or for the term of the maintenance support agreement entered into for this product, whichever is longer.

- gettext
- libiconv
- psqlODBC
- qt

Contact Fujitsu technical support if you require the source code.

# A.3 Provision of the Source Code for Software Based on Microsoft Reciprocal License

This product comes with the following software licensed under the Microsoft Reciprocal License.

The source code of the software below will be provided on request at cost for a period of three years from the date of shipment of this product, or for the term of the maintenance support agreement entered into for this product, whichever is longer.

- WiX Toolset

Contact Fujitsu technical support if you require the source code.

#### A.4 Programs for which Modification and Reverse Engineering are Permitted

Modification and reverse engineering are permissible for the following programs included in this product to the extent permitted under the GNU LESSER GENERAL PUBLIC LICENSE V2.1, GNU LESSER GENERAL PUBLIC LICENSE Version 3, and the GNU LIBRARY GENERAL PUBLIC LICENSE Version 2, which are applied to the open source software used together with this product.

#### Folder names

- <FUJITSU Enterprise Postgres install dir>\bin
- <FUJITSU Enterprise Postgres install dir>\lib
- <Windows FUJITSU Enterprise Postgres Client install dir>\bin
- <Windows FUJITSU Enterprise Postgres Client install dir>\lib
- <Windows FUJITSU Enterprise Postgres Client install dir>\odbc
- <Windows FUJITSU Enterprise Postgres Client install dir>

#### A.5 Apache License

This product includes the following software that are covered by Apache License, Version 2.0

- Modified source code of software Shaj
- apache commons io
- apache commons lang
- apache commons logging
- apache dbcp
- apache derbyApache Tomcat
- commons-codec
- commons-configuration
- commons-pool
- Jackson-annotationsJackson-core
- Jackson-databind
- javax-inject
- log4j
- spring framework
- spring-security-web

For the license statement of Apache License, Version 2.0, refer to the OSS\_Licence folder in this media.

#### A.6 Python

- -ip

This product includes the following Python libraries.

19.2.3

	±	
-	alembic	1.4.2
-	appdirs	1.4.3
	Babel	2.8.0
-	bcrypt	3.1.7
-	blinker	1.4
-	certifi	2020.4.5.3
-	cffi	1.14.0
-	chardet	3.0.4
-	click	7.1.1
-	cryptography	2.9
-	distlib	0.3.0
-	filelock	3.0.12
-	Flask	1.0.2
	Flask-BabelEx	0.9.4
-	Flask-Compress	1.4.0
-	Flask-Gravatar	0.5.0
-	Flask-Login	0.4.1
-	Flask-Mail	0.9.1
-	Flask-Migrate	2.4.0
-		0.2.0
-	Flask-Principal	0.4.0
-	Flask-Security	3.0.0
-	Flask-SQLAlchemy	2.3.2
-	Flask-WTF	0.14.3
	idna	2.9
-	importlib-metadata	1.5.0
-		1.1.0
	Jinja2	2.11.1
-	Mako	1.1.2
-	MarkupSafe	1.1.1
-	1 3 3	20.3
	paramiko	2.7.1
	passlib	1.7.1
-		20.0.2
	psutil	5.5.1
	psycopg2	2.8.5
	pycparser	2.20
	PyNaCl .	1.3.0
	pyparsing	2.4.7
-		2.8.1
-		1.0.4
	pytz	2018.9
-	setuptools	46.0.0

-	simplejson	3.16.0
-	six	1.14.0
-	speaklater	1.3
-	sphinxcontrib-applehelp	1.0.2
-	sphinxcontrib-devhelp	1.0.2
-	sphinxcontrib-htmlhelp	1.0.3
-	sphinxcontrib-jsmath	1.0.1
-	sphinxcontrib-qthelp	1.0.3
-	<pre>sphinxcontrib-serializinghtml</pre>	1.1.4
-	SQLAlchemy	1.3.15
-	sqlparse	0.2.4
-	sshtunnel	0.1.5
-	urllib3	1.25.8
-	virtualenv	20.0.10
-	Werkzeug	1.0.1
-	wheel	0.34.2
-	WTForms	2.2.1
-	zipp	3.1.0

Refer to the files below for the license statement of each library.

- -ip

<Windows FUJITSU Enterprise Postgres Client install dir>
\pgadmin4\venv\Lib\site-packages\~ip-19.2.3.dist-info\LICENSE.txt

- alembic

<Windows FUJITSU Enterprise Postgres Client install dir>
\pgadmin4\venv\Lib\site-packages\alembic-1.4.2.dist-info\LICENSE

- appdirs

<Windows FUJITSU Enterprise Postgres Client install dir>
\pgadmin4\venv\Lib\site-packages\appdirs-1.4.3.dist-info\METADATA

- Babel

<Windows FUJITSU Enterprise Postgres Client install dir>
\pgadmin4\venv\Lib\site-packages\Babel-2.8.0.dist-info\LICENSE

- bcrypt

<Windows FUJITSU Enterprise Postgres Client install dir>
\pgadmin4\venv\Lib\site-packages\bcrypt-3.1.7.dist-info\LICENSE

- blinker

<Windows FUJITSU Enterprise Postgres Client install dir>
\pgadmin4\venv\Lib\site-packages\blinker-1.4.dist-info\LICENSE

- certifi

<Windows FUJITSU Enterprise Postgres Client install dir>
\pgadmin4\venv\Lib\site-packages\certifi-2020.4.5.1.dist-info\LICENSE

- cffi

<Windows FUJITSU Enterprise Postgres Client install dir>
\pgadmin4\venv\Lib\site-packages\cffi-1.14.0.dist-info\LICENSE

- chardet

<Windows FUJITSU Enterprise Postgres Client install dir>
\pgadmin4\venv\Lib\site-packages\chardet-3.0.4.dist-info\METADATA

- click

<Windows FUJITSU Enterprise Postgres Client install dir>
\pgadmin4\venv\Lib\site-packages\click-7.1.1.dist-info\LICENSE.rst

- cryptography

<Windows FUJITSU Enterprise Postgres Client install dir>
\pgadmin4\venv\Lib\site-packages\cryptography-2.9.dist-info\LICENSE

- distlib

<Windows FUJITSU Enterprise Postgres Client install dir>

\pgadmin4\venv\Lib\site-packages\distlib-0.3.0-py3.7.egg-info\PKG-INFO

- filelock

<Windows FUJITSU Enterprise Postgres Client install dir>
\pgadmin4\venv\Lib\site-packages\filelock-3.0.12.dist-info\LICENSE

- Flask

<Windows FUJITSU Enterprise Postgres Client install dir>
\pgadmin4\venv\Lib\site-packages\Flask-1.0.2.dist-info\LICENSE.txt

- Flask-BabelEx

<Windows FUJITSU Enterprise Postgres Client install dir>

\pgadmin4\venv\Lib\site-packages\Flask BabelEx-0.9.4.dist-info\LICENSE

- Flask-Compress

<Windows FUJITSU Enterprise Postgres Client install dir>

\pgadmin4\venv\Lib\site-packages\Flask Compress-1.4.0.dist-info\LICENSE.txt

- Flask-Gravatar

<Windows FUJITSU Enterprise Postgres Client install dir>

\pgadmin4\venv\Lib\site-packages\Flask\_Gravatar-0.5.0.dist-info\METADATA

- Flask-Login

<Windows FUJITSU Enterprise Postgres Client install dir>

\pgadmin4\venv\Lib\site-packages\Flask Login-0.4.1.dist-info\LICENSE

- Flask-Mail

<Windows FUJITSU Enterprise Postgres Client install dir>

\pgadmin4\venv\Lib\site-packages\Flask Mail-0.9.1.dist-info\LICENSE - Flask-Migrate <Windows FUJITSU Enterprise Postgres Client install dir> \pgadmin4\venv\Lib\site-packages\Flask Migrate-2.4.0.dist-info\LICENSE - Flask-Paranoid <Windows FUJITSU Enterprise Postgres Client install dir> \pgadmin4\venv\Lib\site-packages\Flask Paranoid-0.2.0.dist-info\METADATA - Flask-Principal <Windows FUJITSU Enterprise Postgres Client install dir> \pgadmin4\venv\Lib\site-packages\Flask Principal-0.4.0.dist-info\METADATA - Flask-Security <Windows FUJITSU Enterprise Postgres Client install dir> \pgadmin4\venv\Lib\site-packages\Flask Security-3.0.0.dist-info\METADATA - Flask-SQLAlchemy <Windows FUJITSU Enterprise Postgres Client install dir> \pgadmin4\venv\Lib\site-packages\Flask\_SQLAlchemy-2.3.2.dist-info \LICENSE.txt Flask-WTF <Windows FUJITSU Enterprise Postgres Client install dir> \pgadmin4\venv\Lib\site-packages\Flask\_WTF-0.14.3.dist-info\LICENSE.txt - idna <Windows FUJITSU Enterprise Postgres Client install dir> \pgadmin4\venv\Lib\site-packages\idna-2.9.dist-info\LICENSE.rst - importlib-metadata <Windows FUJITSU Enterprise Postgres Client install dir> \pgadmin4\venv\Lib\site-packages\importlib\_metadata-1.5.0.dist-info\LICENSE - itsdangerous <Windows FUJITSU Enterprise Postgres Client install dir> \pgadmin4\venv\Lib\site-packages\itsdangerous-1.1.0.dist-info\LICENSE.rst - Jinja2 <Windows FUJITSU Enterprise Postgres Client install dir> \pgadmin4\venv\Lib\site-packages\Jinja2-2.11.1.dist-info\LICENSE.rst <Windows FUJITSU Enterprise Postgres Client install dir> \pgadmin4\venv\Lib\site-packages\Mako-1.1.2.dist-info\LICENSE - MarkupSafe <Windows FUJITSU Enterprise Postgres Client install dir> \pgadmin4\venv\Lib\site-packages\MarkupSafe-1.1.1.dist-info\LICENSE.rst <Windows FUJITSU Enterprise Postgres Client install dir> \pgadmin4\venv\Lib\site-packages\packaging-20.3.dist-info\LICENSE <Windows FUJITSU Enterprise Postgres Client install dir> \pgadmin4\venv\Lib\site-packages\paramiko-2.7.1.dist-info\LICENSE.txt - passlib <Windows FUJITSU Enterprise Postgres Client install dir> \pgadmin4\venv\Lib\site-packages\passlib-1.7.1.dist-info\METADATA - pip <Windows FUJITSU Enterprise Postgres Client install dir> \pgadmin4\venv\Lib\site-packages\pip-20.0.2.dist-info\LICENSE.txt - psutil <Windows FUJITSU Enterprise Postgres Client install dir> \pgadmin4\venv\Lib\site-packages\psutil-5.5.1.dist-info\LICENSE psycopq2 <Windows FUJITSU Enterprise Postgres Client install dir> \pgadmin4\venv\Lib\site-packages\psycopg2-2.8.5.dist-info\LICENSE - pycparser <Windows FUJITSU Enterprise Postgres Client install dir> \pgadmin4\venv\Lib\site-packages\pycparser-2.20.dist-info\LICENSE - PyNaCl <Windows FUJITSU Enterprise Postgres Client install dir> \pgadmin4\venv\Lib\site-packages\PyNaCl-1.3.0.dist-info\LICENSE.txt pvparsing <Windows FUJITSU Enterprise Postgres Client install dir> \pgadmin4\venv\Lib\site-packages\pyparsing-2.4.7.dist-info\LICENSE python-dateutil <Windows FUJITSU Enterprise Postgres Client install dir> \pgadmin4\venv\Lib\site-packages\python dateutil-2.8.1.dist-info\LICENSE - python-editor <Windows FUJITSU Enterprise Postgres Client install dir> \pgadmin4\venv\Lib\site-packages\python\_editor-1.0.4.dist-info\METADATA <Windows FUJITSU Enterprise Postgres Client install dir> \pgadmin4\venv\Lib\site-packages\pytz-2018.9.dist-info\LICENSE.txt <Windows FUJITSU Enterprise Postgres Client install dir> \pgadmin4\venv\Lib\site-packages\setuptools-46.0.0.dist-info\LICENSE - simplejson

<Windows FUJITSU Enterprise Postgres Client install dir> \pgadmin4\venv\Lib\site-packages\simplejson-3.16.0.dist-info\LICENSE.txt six <Windows FUJITSU Enterprise Postgres Client install dir> \pgadmin4\venv\Lib\site-packages\six-1.14.0.dist-info\LICENSE

- speaklater <Windows FUJITSU Enterprise Postgres Client install dir>

\pgadmin4\venv\Lib\site-packages\speaklater-1.3.dist-info\LICENSE

- sphinxcontrib-applehelp

<Windows FUJITSU Enterprise Postgres Client install dir>

\pgadmin4\venv\Lib\site-packages\sphinxcontrib applehelp-1.0.2.dist-info \LICENSE

- sphinxcontrib-devhelp

<Windows FUJITSU Enterprise Postgres Client install dir>

\pgadmin4\venv\Lib\site-packages\sphinxcontrib devhelp-1.0.2.dist-info \LICENSE

- sphinxcontrib-htmlhelp

<Windows FUJITSU Enterprise Postgres Client install dir>

\pgadmin4\venv\Lib\site-packages\sphinxcontrib htmlhelp-1.0.3.dist-info

sphinxcontrib-jsmath

<Windows FUJITSU Enterprise Postgres Client install dir>

\pgadmin4\venv\Lib\site-packages\sphinxcontrib jsmath-1.0.1.dist-info \LICENSE

- sphinxcontrib-qthelp

<Windows FUJITSU Enterprise Postgres Client install dir>

\pgadmin4\venv\Lib\site-packages\sphinxcontrib\_qthelp-1.0.3.dist-info \LICENSE

- sphinxcontrib-serializinghtml

<Windows FUJITSU Enterprise Postgres Client install dir>

\pgadmin4\venv\Lib\site-packages

\sphinxcontrib\_serializinghtml-1.1.4.dist-info\LICENSE

- SQLAlchemy

<Windows FUJITSU Enterprise Postgres Client install dir>

\pgadmin4\venv\Lib\site-packages\SQLAlchemy-1.3.15.dist-info\LICENSE

- sqlparse

<Windows FUJITSU Enterprise Postgres Client install dir>

\pgadmin4\venv\Lib\site-packages\sqlparse-0.2.4.dist-info\METADATA

<Windows FUJITSU Enterprise Postgres Client install dir>

\pgadmin4\venv\Lib\site-packages\sshtunnel-0.1.5.dist-info\LICENSE

urllib3

<Windows FUJITSU Enterprise Postgres Client install dir>

\pgadmin4\venv\Lib\site-packages\urllib3-1.25.8.dist-info\LICENSE.txt

virtualenv

<Windows FUJITSU Enterprise Postgres Client install dir>

\pgadmin4\venv\Lib\site-packages\virtualenv-20.0.10.dist-info\LICENSE

- Werkzeug

<Windows FUJITSU Enterprise Postgres Client install dir>

\pgadmin4\venv\Lib\site-packages\Werkzeug-1.0.1.dist-info\LICENSE.rst

- wheel

<Windows FUJITSU Enterprise Postgres Client install dir>

\pgadmin4\venv\Lib\site-packages\wheel-0.34.2.dist-info\LICENSE.txt

- WTForms

<Windows FUJITSU Enterprise Postgres Client install dir>

\pgadmin4\venv\Lib\site-packages\WTForms-2.2.1.dist-info\METADATA

- zipp

<Windows FUJITSU Enterprise Postgres Client install dir> \pgadmin4\venv\Lib\site-packages\zipp-3.1.0.dist-info\LICENSE

#### A.7 Microsoft Visual C++ 2010 Redistributable Package

Microsoft Visual C++ 2010 redistributable package are included in this

The customer cannot distribute Microsoft Visual C++ 2010 redistributable package.

#### A.8 Microsoft Visual C++ 2013 Redistributable Package

Microsoft Visual C++ 2013 redistributable package are included in this

The customer cannot distribute Microsoft Visual C++ 2013 redistributable package.

Microsoft Visual  $C++\ 2015-2019$  redistributable package are included in this software.

The customer cannot distribute Microsoft Visual C++ 2015-2019 redistributable package.

Appendix B Viewing FJQSS and FUJITSU Enterprise Postgres HTML Files

If you log in using the built-in Administrator account and attempt to view the FJQSS User's Guide(HTML) from the application menu or Manual(Java API) in the FUJITSU Enterprise Postgres group, the following error message will be output and you may not be able to view the file.

"This app can't open. Internet Explorer can't be opened using the Built-in Administrator account. Sign in with a different account and try again."

In this case, take one of the following actions to resolve the problem:

- As indicated in the message, log in again using an account other than the built-in Administrator account, and then view the above-mentioned document.
- Change the Internet Explorer settings for the built-in Administrator. By default, the Internet options are set to "Always in Internet Explorer". Change this setting to either of the following options:
  - -- Let Internet Explorer decide
  - -- Always in Internet Explorer on the desktop



# FUJITSU Enterprise Postgres 12



# Installation/Setup

Server >

Client >

Server Assistant >



# FUJITSU Enterprise Postgres 12



# Installation and Setup Guide for Server

Linux > Windows >



# FUJITSU Enterprise Postgres 12



# Installation and Setup Guide for Server (Linux)

# **Preface**

#### Purpose of this document

The FUJITSU Enterprise Postgres database system extends the PostgreSQL features and runs on the Linux platform. This document describes how to install and set up "FUJITSU Enterprise Postgres".

#### Intended readers

This document is intended for those who install and operate FUJITSU Enterprise Postgres.

Readers of this document are assumed to have general knowledge of:

- PostgreSQL
- SQL
- Linux

#### Structure of this document

This document is structured as follows:

#### Chapter 1 Overview of Installation

Describes the installation types and procedures

#### Chapter 2 Operating Environment

Describes the operating environment required to use FUJITSU Enterprise Postgres

#### Chapter 3 Installation

Describes how to perform a new installation of FUJITSU Enterprise Postgres

#### Chapter 4 Setup

Describes the setup to be performed after installation

#### Chapter 5 Uninstallation

Describes how to uninstall FUJITSU Enterprise Postgres

#### Appendix A Recommended WebAdmin Environments

Describes the recommended WebAdmin environment.

#### Appendix B Setting Up and Removing WebAdmin

Describes how to set up and remove WebAdmin

#### Appendix C WebAdmin Disallow User Inputs Containing Hazardous Characters

Describes characters that are not allowed in WebAdmin.

#### Appendix D Configuring Parameters

Describes FUJITSU Enterprise Postgres parameters.

#### Appendix E Estimating Database Disk Space Requirements

Describes how to estimate database disk space requirements

#### Appendix F Estimating Memory Requirements

Describes the formulas for estimating memory requirements

#### Appendix G Quantitative Limits

Describes the quantity range

#### Appendix H Configuring Kernel Parameters

Describes the settings for kernel parameters

#### Appendix I Determining the Preferred WebAdmin Configuration

Describes the two different configurations in which WebAdmin can be used and how to select the most suitable configuration

#### Appendix J System Configuration when using Pgpool-II

Describes the system configuration when using Pgpool-II.

#### Appendix K Supported contrib Modules and Extensions Provided by External Projects

Lists the PostgreSQL contrib modules and the extensions provided by external projects supported by FUJITSU Enterprise Postgres.

#### Appendix L Procedure when Modifying the Oracle JRE Installation

Describes the procedures to follow when modifying the Oracle JRE installation.

#### **Export restrictions**

Exportation/release of this document may require necessary procedures in accordance with the regulations of your resident country and/or US export control laws.

#### Issue date and version

Edition 1.0: July 2020

#### Copyright

Copyright 2015-2020 FUJITSU LIMITED

# Contents

Chapter 1 Overview of Installation	,
1.1 Features that can be Installed.	
1.2 Installation Types.	
1.2.1 New Installation.	
1.2.2 Reinstallation.	
1.2.3 Multi-Version Installation.	
1.3 Installation Procedure.	
1.3.1 Installation in Interactive Mode	
1.3.2 Installation in Silent Mode	
1.4 Uninstallation.	
1.4 Ullilistaliation	
Chapter 2 Operating Environment	
2.1 Required Operating System	
2.2 Related Software	
2.3 Excluded Software	
2.4 Required Patches.	
2.5 Hardware Environment.	
2.6 Disk Space Required for Installation.	
2.7 Supported System Environment	
2.7.1 TCP/IP Protocol.	
2.7.2 File System	
2.8 PostgreSQL Version Used for FUJITSU Enterprise Postgres	
2.9 Notes on Using Streaming Replication.	
2.5 Protes on Comg Streaming Repredation	
Chapter 3 Installation	
3.1 Pre-installation Tasks	
3.2 Installation in Interactive Mode	10
3.3 Installation in Silent Mode	13
Chapter 4 Setup	
4.1 Operating Method Types and Selection.	15
4.2 Preparations for Setup	15 16
4.2 Preparations for Setup	15 16
4.2 Preparations for Setup	
4.2 Preparations for Setup	15 10 17
4.2 Preparations for Setup	
4.2 Preparations for Setup 4.2.1 Creating an Instance Administrator 4.2.2 Preparing Directories for Resource Deployment 4.2.3 Estimating Resources 4.2.4 Configuring Corefile Names 4.3 Creating Instances	
4.2 Preparations for Setup 4.2.1 Creating an Instance Administrator 4.2.2 Preparing Directories for Resource Deployment 4.2.3 Estimating Resources 4.2.4 Configuring Corefile Names 4.3 Creating Instances 4.3.1 Using WebAdmin	
4.2 Preparations for Setup  4.2.1 Creating an Instance Administrator  4.2.2 Preparing Directories for Resource Deployment  4.2.3 Estimating Resources  4.2.4 Configuring Corefile Names  4.3 Creating Instances  4.3.1 Using WebAdmin  4.3.1.1 Logging in to WebAdmin	
4.2 Preparations for Setup	
4.2 Preparations for Setup  4.2.1 Creating an Instance Administrator  4.2.2 Preparing Directories for Resource Deployment  4.2.3 Estimating Resources  4.2.4 Configuring Corefile Names  4.3 Creating Instances  4.3.1 Using WebAdmin  4.3.1.1 Logging in to WebAdmin  4.3.1.2 Creating an Instance  4.3.1.3 Changing Instance Settings	
4.2 Preparations for Setup 4.2.1 Creating an Instance Administrator 4.2.2 Preparing Directories for Resource Deployment 4.2.3 Estimating Resources 4.2.4 Configuring Corefile Names 4.3 Creating Instances 4.3.1 Using WebAdmin 4.3.1.1 Logging in to WebAdmin 4.3.1.2 Creating an Instance 4.3.1.3 Changing Instance Settings 4.3.1.4 Importing Instances	
4.2 Preparations for Setup  4.2.1 Creating an Instance Administrator  4.2.2 Preparing Directories for Resource Deployment  4.2.3 Estimating Resources  4.2.4 Configuring Corefile Names  4.3 Creating Instances  4.3.1 Using WebAdmin  4.3.1.1 Logging in to WebAdmin  4.3.1.2 Creating an Instance  4.3.1.3 Changing Instance Settings	
4.2 Preparations for Setup 4.2.1 Creating an Instance Administrator 4.2.2 Preparing Directories for Resource Deployment 4.2.3 Estimating Resources 4.2.4 Configuring Corefile Names 4.3 Creating Instances 4.3.1 Using WebAdmin 4.3.1.1 Logging in to WebAdmin 4.3.1.2 Creating an Instance 4.3.1.3 Changing Instance Settings 4.3.1.4 Importing Instances	
4.2 Preparations for Setup 4.2.1 Creating an Instance Administrator 4.2.2 Preparing Directories for Resource Deployment 4.2.3 Estimating Resources 4.2.4 Configuring Corefile Names 4.3 Creating Instances 4.3.1 Using WebAdmin 4.3.1.1 Logging in to WebAdmin 4.3.1.2 Creating an Instance 4.3.1.3 Changing Instance Settings 4.3.1.4 Importing Instances 4.3.2 Using the initdb Command	15 16 16 17 17 18 19 19 19 19 19 19 19 19 19 19 19 19 19
4.2 Preparations for Setup 4.2.1 Creating an Instance Administrator 4.2.2 Preparing Directories for Resource Deployment 4.2.3 Estimating Resources 4.2.4 Configuring Corefile Names 4.3 Creating Instances 4.3.1 Using WebAdmin 4.3.1.1 Logging in to WebAdmin 4.3.1.2 Creating an Instance 4.3.1.3 Changing Instance Settings 4.3.1.4 Importing Instances 4.3.2 Using the initdb Command 4.3.2.1 Editing Kernel Parameters	15 16 16 17 17 18 19 19 19 19 19 19 19 19 19 19 19 19 19
4.2 Preparations for Setup 4.2.1 Creating an Instance Administrator. 4.2.2 Preparing Directories for Resource Deployment 4.2.3 Estimating Resources 4.2.4 Configuring Corefile Names 4.3 Creating Instances 4.3.1 Using WebAdmin 4.3.1.1 Logging in to WebAdmin 4.3.1.2 Creating an Instance 4.3.1.3 Changing Instance Settings 4.3.1.4 Importing Instances 4.3.2 Using the initdb Command 4.3.2.1 Editing Kernel Parameters 4.3.2.2 Creating an Instance	11 10 10 10 11 11 11 11 11 11 11 11 11 1
4.2 Preparations for Setup.  4.2.1 Creating an Instance Administrator.  4.2.2 Preparing Directories for Resource Deployment.  4.2.3 Estimating Resources.  4.2.4 Configuring Corefile Names.  4.3 Creating Instances.  4.3.1 Using WebAdmin.  4.3.1.1 Logging in to WebAdmin.  4.3.1.2 Creating an Instance.  4.3.1.3 Changing Instance Settings.  4.3.1.4 Importing Instances.  4.3.2 Using the initdb Command.  4.3.2.1 Editing Kernel Parameters.  4.3.2.2 Creating an Instance.  4.3.2.2 Creating an Instance.  4.4 Configuring Remote Connections.	16 10 10 10 11 11 11 11 11 11 11 11 11 11
4.2 Preparations for Setup 4.2.1 Creating an Instance Administrator 4.2.2 Preparing Directories for Resource Deployment 4.2.3 Estimating Resources 4.2.4 Configuring Corefile Names 4.3 Creating Instances 4.3.1 Using WebAdmin 4.3.1.1 Logging in to WebAdmin 4.3.1.2 Creating an Instance 4.3.1.3 Changing Instance Settings 4.3.1.4 Importing Instances 4.3.2 Using the initdb Command 4.3.2.1 Editing Kernel Parameters 4.3.2.2 Creating an Instance 4.4 Configuring Remote Connections 4.4.1 When an Instance was Created with WebAdmin	16 16 16 17 19 20 20 20 22 25 25 25 25 25 25 3 3 3 3 3 3 3 3 3
4.2 Preparations for Setup 4.2.1 Creating an Instance Administrator. 4.2.2 Preparing Directories for Resource Deployment. 4.2.3 Estimating Resources 4.2.4 Configuring Corefile Names 4.3 Creating Instances 4.3.1 Using WebAdmin 4.3.1.1 Logging in to WebAdmin 4.3.1.2 Creating an Instance 4.3.1.3 Changing Instance Settings 4.3.1.4 Importing Instances 4.3.2 Using the initdb Command 4.3.2.1 Editing Kernel Parameters 4.3.2.2 Creating an Instance 4.4 Configuring Remote Connections 4.4 Configuring Remote Connections 4.4.1 When an Instance was Created with WebAdmin 4.4.2 When an Instance was Created with the initdb Command	11 10 10 10 11 11 11 11 11 11 11 11 11 1
4.2 Preparations for Setup 4.2.1 Creating an Instance Administrator. 4.2.2 Preparing Directories for Resource Deployment. 4.2.3 Estimating Resources. 4.2.4 Configuring Corefile Names. 4.3 Creating Instances. 4.3.1 Using WebAdmin. 4.3.1.1 Logging in to WebAdmin. 4.3.1.2 Creating an Instance 4.3.1.3 Changing Instance Settings. 4.3.1.4 Importing Instances. 4.3.2 Using the initdb Command. 4.3.2.1 Editing Kernel Parameters. 4.3.2.2 Creating an Instance. 4.4 Configuring Remote Connections. 4.4.1 When an Instance was Created with WebAdmin. 4.4.2 When an Instance was Created with the initdb Command. 4.5 Other Settings.	11 10 10 10 11 11 11 11 11 11 11 11 11 1
4.2 Preparations for Setup 4.2.1 Creating an Instance Administrator. 4.2.2 Preparing Directories for Resource Deployment 4.2.3 Estimating Resources. 4.2.4 Configuring Corefile Names. 4.3 Creating Instances. 4.3.1 Using WebAdmin. 4.3.1.1 Logging in to WebAdmin. 4.3.1.2 Creating an Instance. 4.3.1.3 Changing Instance Settings. 4.3.1.4 Importing Instances. 4.3.2 Using the initdb Command 4.3.2.1 Editing Kernel Parameters. 4.3.2.2 Creating an Instance. 4.4 Configuring Remote Connections. 4.4.1 When an Instance was Created with WebAdmin. 4.4.2 When an Instance was Created with the initdb Command. 4.5 Other Settings. 4.5.1 Error Log Settings.	11 10 10 10 11 11 11 11 11 11 11 11 11 1
4.2 Preparations for Setup 4.2.1 Creating an Instance Administrator 4.2.2 Preparing Directories for Resource Deployment 4.2.3 Estimating Resources 4.2.4 Configuring Corefile Names 4.3 Creating Instances 4.3.1 Using WebAdmin 4.3.1.1 Logging in to WebAdmin 4.3.1.2 Creating an Instance 4.3.1.3 Changing Instance Settings 4.3.1.4 Importing Instances 4.3.2 Using the initdb Command 4.3.2.1 Editing Kernel Parameters 4.3.2.2 Creating an Instance 4.4 Configuring Remote Connections. 4.4.1 When an Instance was Created with WebAdmin 4.4.2 When an Instance was Created with the initdb Command. 4.5 Other Settings 4.5.1 Error Log Settings 4.5.2 Configuring Automatic Start and Stop of an Instance	11 10 10 10 11 11 11 11 11 11 11 11 11 1
4.2 Preparations for Setup.  4.2.1 Creating an Instance Administrator.  4.2.2 Preparing Directories for Resource Deployment.  4.2.3 Estimating Resources.  4.2.4 Configuring Corefile Names.  4.3 Creating Instances.  4.3.1 Using WebAdmin.  4.3.1.1 Logging in to WebAdmin.  4.3.1.2 Creating an Instance.  4.3.1.3 Changing Instance Settings.  4.3.1.4 Importing Instances.  4.3.2 Using the initdb Command.  4.3.2.1 Editing Kernel Parameters.  4.3.2.2 Creating an Instance.  4.4 Configuring Remote Connections.  4.4.1 When an Instance was Created with WebAdmin.  4.4.2 When an Instance was Created with the initdb Command.  4.5 Other Settings.  4.5.1 Error Log Settings.  4.5.2 Configuring Automatic Start and Stop of an Instance.  4.5.3 Settings when Using the Features Compatible with Oracle Databases.	11: 10: 10: 10: 11: 11: 11: 11: 11: 11:
4.2 Preparations for Setup 4.2.1 Creating an Instance Administrator 4.2.2 Preparing Directories for Resource Deployment 4.2.3 Estimating Resources 4.2.4 Configuring Corefile Names 4.3 Creating Instances 4.3.1 Using WebAdmin 4.3.1.1 Logging in to WebAdmin 4.3.1.2 Creating an Instance 4.3.1.3 Changing Instance Settings 4.3.1.4 Importing Instances 4.3.2 Using the initdb Command 4.3.2.1 Editing Kernel Parameters 4.3.2.2 Creating an Instance 4.4 Configuring Remote Connections 4.4.1 When an Instance was Created with WebAdmin 4.4.2 When an Instance was Created with the initdb Command 4.5 Other Settings 4.5.1 Error Log Settings 4.5.2 Configuring Automatic Start and Stop of an Instance 4.6 Integration with Message-Monitoring Software	16 16 16 17 19 19 20 20 20 22 25 25 25 25 3 3 3 3 3 3 3 3 3 3 3 3

4.7.1.2 Removing oracle_fdw	36
4.7.2 pg_bigm	
4.7.2.1 Setting Up pg_bigm	37
4.7.2.2 Removing pg_bigm	37
4.7.3 pg_hint_plan	
4.7.3.1 Setting Up pg_hint_plan	37
4.7.3.2 Removing pg_hint_plan	38
4.7.4 pg_dbms_stats	38
4.7.4.1 Setting Up pg_dbms_stats	
4.7.4.2 Removing pg_dbms_stats	39
4.7.5 pg_repack	39
4.7.5.1 Setting Up pg_repack	39
4.7.5.2 Removing pg_repack	40
4.7.6 pg_rman	40
4.7.6.1 Setting Up pg_rman	
4.7.6.2 Removing pg_rman	40
4.7.7 pg_statsinfo	41
4.7.7.1 Setting Up pg_statsinfo	41
4.7.7.2 Removing pg_statsinfo	41
4.7.8 pgBadger	42
4.7.8.1 Setting Up pgBadger	
4.7.8.2 Removing pgBadger	
4.7.9 Pgpool-II	
4.7.9.1 Setting Up Pgpool-II	
4.7.9.2 Removing Pgpool-II	
4.8 Deleting Instances.	43
4.8.1 Using WebAdmin	43
4.8.2 Using Server Commands	44
Chapter 5 Uninstallation	4.5
•	
5.1 Uninstallation in Interactive Mode	
5.2 Uninstanation in Shent Wode	4
Appendix A Recommended WebAdmin Environments	50
A.1 Recommended Browser Settings	50
A.2 How to Set Up the Pop-up Blocker	50
Appendix B Setting Up and Removing WebAdmin	
B.1 Setting Up WebAdmin	
B.1.1 Setting Up WebAdmin	
B.1.2 Starting the Web Server Feature of WebAdmin	
B.1.3 Stopping the Web Server Feature of WebAdmin	
B.2 Removing WebAdmin.	
B.3 Using an External Repository for WebAdmin	
B.4 Using the WebAdmin Auto-Refresh Feature	55
Appendix C WebAdmin Disallow User Inputs Containing Hazardous Characters	56
Appendix D Configuring Parameters	57
Apparedix E Estimation Database Diels On the Description	~
Appendix E Estimating Database Disk Space Requirements	
E.1 Estimating Table Size Requirements.	
E.2 Estimating Index Size Requirements.	
E.3 Sizes of Data Types.	
E.3.1 Sizes of Fixed-Length Data Types	
E.3.2 Sizes of Variable-Length Data Types	
E.3.3 Sizes of Array Data Types	
E.3.4 Number of Bytes per Character	
E.4 Estimating Transaction Log Space Requirements	

E.5 Estimating Archive Log Space Requirements	65
E.6 Estimating Backup Disk Space Requirements	65
E.7 Estimating VCI Disk Space Requirements	65
Appendix F Estimating Memory Requirements	67
F.1 FUJITSU Enterprise Postgres Memory Requirements	67
F.2 Database Multiplexing Memory Requirements	68
F.3 VCI Memory Requirements	68
F.4 High-Speed Data Load Memory Requirements	70
F.5 Global Meta Cache Memory Requirements	70
Appendix G Quantitative Limits	71
Appendix H Configuring Kernel Parameters	76
Appendix I Determining the Preferred WebAdmin Configuration	77
I.1 WebAdmin Configurations	77
I.1.1 Single-Server Configuration	77
I.1.2 Multiserver Configuration	
I.2 Installing WebAdmin in a Single-Server Configuration	78
I.3 Installing WebAdmin in a Multiserver Configuration.	79
Appendix J System Configuration when using Pgpool-II	
J.1 Pgpool-II Configuration	80
J.1.1 Single-Machine Configuration	80
J.1.2 Two-Machine Configuration.	81
J.1.3 Three-Machine Configuration	81
J.2 Installing Pgpool-II	81
Appendix K Supported contrib Modules and Extensions Provided by External Projects	83
Appendix L Procedure when Modifying the Oracle JRE Installation	
L.1 When Using WebAdmin.	
L.2 When Performing Database Multiplexing	84
Index	86

# Chapter 1 Overview of Installation

This chapter provides an overview of FUJITSU Enterprise Postgres installation.

#### 1.1 Features that can be Installed

Each FUJITSU Enterprise Postgres feature is installed on the machine that was used to build the database environment.

The following table shows the relationship between the product to be installed and the features that can be installed.

Feature that can be installed Product name		t name
	AE	SE
Basic feature (server feature, client feature)	Y	Y

Y: Can be installed

# 1.2 Installation Types

The following installation types are available for FUJITSU Enterprise Postgres:

- New installation
- Reinstallation
- Multi-version installation

#### 1.2.1 New Installation

In initial installation, FUJITSU Enterprise Postgres is installed for the first time.

#### 1.2.2 Reinstallation

Perform reinstallation to repair installed program files that have become unusable for any reason.

#### 1.2.3 Multi-Version Installation

FUJITSU Enterprise Postgres products can be installed on the same server if the product version (indicated by "x" in "x SPz") is different from that of any version of the product that is already installed.

#### 1.3 Installation Procedure

The following installation procedures are available for FUJITSU Enterprise Postgres:

- Installation in interactive mode
- Installation in silent mode

Select the installation procedure that corresponds to your environment.



If you have antivirus software installed, the server may crash, fail to start, or stop responding, during installation or when starting up after installation. Set scan exception settings for the installation directory and resource allocation directory so that the files in these directories are not scanned for viruses.

#### 1.3.1 Installation in Interactive Mode

Interactive mode enables installation to be performed while the required information is entered interactively.

In the interactive mode installation, the installation state of FUJITSU Enterprise Postgres is determined automatically. Install FUJITSU Enterprise Postgres using one of the following installation types in accordance with the installation state:

- New installation
- Reinstallation
- Multi-version installation

#### 1.3.2 Installation in Silent Mode

Silent mode enables installation to be performed without the need to enter any information interactively.

New installations and multi-version installations can be performed in silent mode.

# 1.4 Uninstallation

Uninstallation removes the system files of the installed FUJITSU Enterprise Postgres.

# Chapter 2 Operating Environment

This chapter describes the operating environment required to use FUJITSU Enterprise Postgres.



Refer to "Operating Environment" in the Installation and Setup Guide for Client when installing the FUJITSU Enterprise Postgres client feature at the same time.

# 2.1 Required Operating System

One of the operating systems shown below is required in order to use FUJITSU Enterprise Postgres.

- RHEL7.2 or later
- RHEL8.1 or later
- SLES 12 SP2 or later



SLES 12 cannot be used if performing failover operation integrated with PRIMECLUSTER.



- The following packages are required for operations on RHEL7.

Package name	Remarks
alsa-lib	-
audit-libs	-
dstat	Required when using parallel scan.
glibc	-
glibc.i686	-
libgcc	-
1.1 .	Provides collation support.
libicu	Install 50.x.
libmemcached	Required when using Pgpool-II.
libstdc++	-
libtool-ltdl	Required when using ODBC drivers.
llvm	Version 5.0.x of llvm is required to run SQL with runtime compilation (just-in-time compilation). Install the "llvm-toolset -7 llvm-libs" published in Red Hat Software Collections.
	FUJITSU Enterprise Postgres 12 uses runtime compilation by default. If you do not want to use runtime compilation, turn off the jit parameter in postgresql.conf. You do not need to install llvm if you turn off the jit parameter.
	Failure to install llvm without turning off the jit parameter may result in an error when executing SQL.For more information about

Package name	Remarks
	runtime compilation, see "Just-in-Time Compilation (JIT)" in "PostgreSQL Documentation".
ncurses-libs	-
net-tools	-
nss-softokn-freebl	-
pam	Required when using PAM authentication.
perl-libs	Required when using PL/Perl. Install 5.16.
python3	Required when using PL/Python based on Python 3. Install 3.7.x.
redhat-lsb	-
libselinux	Required for sepgsql.
tcl	Required when using PL/Tcl. Install 8.5.
unzip	-
xz-libs	-
zlib	-

- The following packages are required for operations on RHEL8.1.

Package name	Remarks
alsa-lib	
audit-libs	
cyrus-sasl-lib	
pcp-system-tools	Required when using parallel scan.
glibc	
glibc.i686	
libnsl2	
1.1 .	Provides collation support.
libicu	Install 60.x.
libgcc	
libmemcached	Required when using Pgpool-II.
libstdc++	
libtool-ltdl	Required when using ODBC drivers.
llvm	Version 8.0.x of llvm is required to run SQL with runtime compilation (just-in-time compilation). Install the "llvm-libs" exposed by Application Streams.
	FUJITSU Enterprise Postgres 12 uses runtime compilation by default. If you do not want to use runtime compilation, turn off the jit parameter in postgresql.conf. You do not need to install llvm if you turn off the jit parameter.
	Failure to install llvm without turning off the jit parameter may result in an error when executing SQL.For more information about

Package name	Remarks
	runtime compilation, see "Just-in-Time Compilation (JIT)" in "PostgreSQL Documentation".
ncurses-libs	
net-tools	
nss-softokn-freebl	
pam	Required when using PAM authentication.
perl-libs	Required when using PL/Perl. Install 5.26.
python3	Required when using PL/Python based on Python 3. Install 3.7.x.
redhat-lsb	
libselinux	Required for sepgsql.
tcl	Required when using PL/Tcl. Install 8.6.
unzip	
xz-libs	
zlib	

- The following packages are required for operations on SLES 12.

Package name	Remarks
alsa-lib	-
audit-libs	-
dstat	Required when using parallel scan.
glibc	-
glibc-32bit	-
libicu	Provides collation support. Install 52.
libgcc	-
libmemcached	Required when using Pgpool-II.
libstdc++	-
libtool-ltdl	Required when using ODBC drivers.
llvm	Install version 5.0.x of llvm to run SQL with runtime compilation (just-in-time compilation) and add the directory where the shared library libLLVM -5.0.so is located to the environment variable LD_LIBRARY_PATH.
	FUJITSU Enterprise Postgres 12 uses runtime compilation by default. If you do not want to use runtime compilation, turn off the jit parameter in postgresql.conf. You do not need to install llvm if you turn off the jit parameter.
	Failure to install llvm without turning off the jit parameter may result in an error when executing SQL.For more information about runtime compilation, see "Just-in-Time Compilation (JIT)" in "PostgreSQL Documentation".
LLVM-libs	Install version 5.0.2 or later.

Package name	Remarks
ncurses-libs	-
net-tools	-
nss-softokn-freebl	-
Oracle JRE 8	Required when using the installer or database multiplexing. Use update 31 or later of the 64-bit version.
pam	Required when using PAM authentication.
perl-libs	Required when using PL/Perl. Install 5.16.
python3	Required when using PL/Python based on Python 3. Install 3.7.x.
sysstat	Required when using FJQSS. Set up the sar command after installation.
tcl	Required when using PL/Tcl. Install 8.6.
unzip	-
xz-libs	-
zlib	-

# 2.2 Related Software

The following table lists the software required to use FUJITSU Enterprise Postgres.

Table 2.1 Related software

No.	Product name	Version	FUJITSU Enterprise Postgres product name		Remarks	
			AE	SE		
1	PRIMECLUSTER (*1)	4.5A10 (*1)	Y	Y	Mandatory when performing failover operation integrated with PRIMECLUSTER.	

<sup>\*1:</sup> The applicable products are shown below:

- PRIMECLUSTER Enterprise Edition
- PRIMECLUSTER HA Server

The following table lists servers that can be connected to the FUJITSU Enterprise Postgres client feature.

Table 2.2 Connectable servers

os	Product name
Windows Linux	- FUJITSU Software Enterprise Postgres Advanced Edition 9.5 or later
	- FUJITSU Software Enterprise Postgres Standard Edition 9.4 or later
Solaris	- FUJITSU Software Enterprise Postgres Standard Edition 9.6 or later

### 2.3 Excluded Software

This section describes excluded software.

#### **FUJITSU Enterprise Postgres**

FUJITSU Enterprise Postgres cannot be installed if all the following conditions are met:

- The product version (indicated by "x" in "x SPz") of the product to be installed is the same as that of the installed product
- The editions are different

#### Example

In the following cases, FUJITSU Enterprise Postgres cannot be installed as an exclusive product:

- The installed product is FUJITSU Software Enterprise Postgres Standard Edition (64bit) 12
- The product to be installed is FUJITSU Software Enterprise Postgres Advanced Edition (64bit) 12

#### Other products

There are no exclusive products.

# 2.4 Required Patches

There are no required patches.

# 2.5 Hardware Environment

The following hardware is required to use FUJITSU Enterprise Postgres.

#### Memory

At least 512 MB of memory is required.

# 2.6 Disk Space Required for Installation

The following table shows the disk space requirements for new installation of FUJITSU Enterprise Postgres. If necessary, increase the size of the file system.

Disk space required for installation

#### RHEL

Directory	Required disk space (Unit: MB)	
/etc	1 + 1 (*1)	
/var	1 + 2 (*1) +1 (*2)	
/opt	106 (*1) +1 (*2)	
Installation destination of the server	357	
Installation destination of WebAdmin	482	
Installation destination of the client (64-bit)	126	
Installation destination of Pgpool-II	13	

<sup>\*1:</sup> Required for the installation of the Uninstall (middleware) tool.

#### **SLES**

<sup>\*2:</sup> Required for the installation of FJQSS.

Directory	Required disk space (Unit: MB)	
/etc	1+1(*1)	
/var	1 + 2 (*1) +1 (*2)	
/opt	2 (*1) +1 (*2)	
Installation destination of the server	213	
Installation destination of WebAdmin	479	
Installation destination of the client (64-bit)	121	
Installation destination of Pgpool-II	10	

<sup>\*1:</sup> Required for the installation of the Uninstall (middleware) tool.

# 2.7 Supported System Environment

This section describes the supported system environment.

#### 2.7.1 TCP/IP Protocol

FUJITSU Enterprise Postgres supports version 4 and 6 (IPv4 and IPv6) of TCP/IP protocols.



Do not use link-local addresses if TCP/IP protocol version 6 addresses are used.

## 2.7.2 File System

All file systems with a POSIX-compliant interface are supported.

However, for stable system operation, the disk where the database is deployed must use a highly reliable file system. Consider this aspect when selecting the file system to be used.

The recommended file system is "ext4".

# 2.8 PostgreSQL Version Used for FUJITSU Enterprise Postgres

FUJITSU Enterprise Postgres is based on PostgreSQL 12.1.

# 2.9 Notes on Using Streaming Replication

To use streaming replication, build the primary server and all standby servers using the same FUJITSU Enterprise Postgres version (\*1).

\*1: The product version is indicated by "x" in the notation "x SPz".



Streaming replication cannot be used in combination with Open Source PostgreSQL.

<sup>\*2:</sup> Required for the installation of FJQSS.

# Chapter 3 Installation

This chapter explains each of the installation procedures of FUJITSU Enterprise Postgres.

#### 3.1 Pre-installation Tasks

Check the system environment below before installing FUJITSU Enterprise Postgres.

#### Check the disk space

Ensure that there is sufficient disk space to install FUJITSU Enterprise Postgres.

Refer to "2.6 Disk Space Required for Installation" for information on the required disk space.

Reconfigure the disk partition if disk space is insufficient.

#### Set JAVA\_HOME (SLES 12 only)

Ensure that Oracle JRE 8 is installed, and export the JAVA\_HOME environment variable.

```
#export JAVA_HOME="oracleJre8InstallDir"
```

Refer to "Appendix L Procedure when Modifying the Oracle JRE Installation" for information on modifying Oracle JRE after installation.

#### Check the installed products and determine the installation method

Using the operation shown below, start Uninstall (middleware), and check the installed products.

#### Example

```
# /opt/FJSVcir/cimanager.sh -c
Loading Uninstaller...

Currently installed products
1. productName
2. productName
3. productName
4. productName
5: productName
Type [number] to select the software you want to uninstall.
[number,q]
=>q

Exiting Uninstaller.
```

If FUJITSU Enterprise Postgres is already installed, determine which installation method to use:

- Reinstallation
- Multi-version installation

#### Remove applied updates

If you perform reinstallation as the installation method, remove applied updates using the procedure shown below.



If a product is installed without removing applied updates, the following problems will occur:

- Performing reinstallation

If an update with the same update and version number is applied, an error informing you that the update has already been applied is displayed.

Perform reinstallation after removing the update.

1. Display the applied updates

Execute the following command to display the applied updates:

# /opt/FJSVfupde/bin/uam showup

#### 2. Remove the updates

Execute the command below to remove the updates. If an update with the same update number was applied more than once, the updates are removed in order, starting from the highest version number.

#### Determine the preferred WebAdmin configuration

Starting with FUJITSU Enterprise Postgres 9.5, WebAdmin can be installed in two configurations:

- Single-server
- Multiserver



See

Refer to "Appendix I Determining the Preferred WebAdmin Configuration" for details.

#### **Determining the Pgpool-II System Configuration**

The system configuration when using Pgpool-II is as follows:

- Place on database server
- Place on application server
- Place on dedicated server



See

Refer to "Appendix J System Configuration when using Pgpool-II".

### 3.2 Installation in Interactive Mode

Install according to the following procedure:



- The following characters can be used as input values:

Alphanumeric characters, hyphens, commas and forward slashes

- When reinstalling the product, back up the following folder in which the WebAdmin instance management information is stored:

webAdminInstallFolder/data/fepwa

Follow the procedure below to perform the backup.

1. Stop the WebAdmin server. Refer to "B.1.3 Stopping the Web Server Feature of WebAdmin" for details.

#### 2. Back up the following folder:

webAdminInstallFolder/data/fepwa

Replace the above folder with the backed up folder when the reinstallation is complete.



For installation in interactive mode, default values are set for the installation information. The following settings can be changed for a new installation or a multi-version installation:

- Installation directory
- WebAdmin setup information, if WebAdmin is selected

To change the port number, confirm that it is an unused port number between 1024 and 32767.

#### 1. Stop applications and programs

If the installation method is the following, all applications and programs that use the product must be stopped:

- Reinstallation

Before starting the installation, stop the following:

- Applications that use the product
- Instance
- Web server feature of WebAdmin

Execute the WebAdminStop command to stop the Web server feature of WebAdmin.

#### Example

If WebAdmin is installed in /opt/fsepv<*x*>webadmin:

- # cd /opt/fsepv<x>webadmin/sbin
- # ./WebAdminStop
- Mirroring Controller

Execute the mc\_ctl command with the stop mode option specified and stop the Mirroring Controller.

#### Example

```
$ mc_ctl stop -M /mcdir/inst1
```

- pgBadger
- Pgpool-II

#### 2. Change to the superuser

Run the following command to switch to the superuser on the system.

```
$ su -
Password:*****
```

#### 3. Mount the DVD drive

Insert the server program DVD into the DVD drive, and run the command given below.

#### Example

```
# mount -t iso9660 -r -o loop /dev/dvd /media/dvd
```

Here /dev/dvd is the device name for the DVD drive (which may vary depending on your environment), and /media/dvd is the mount point (which may need to be created before calling the command).



If the DVD was mounted automatically using the automatic mount daemon (autofs), "noexec" is set as the mount option, so the installer may fail to start. In this case, use the mount command to remount the DVD correctly, and then run the installation. Note that the mount options of a mounted DVD can be checked by executing the mount command without any arguments.

#### 4. Run the installation

Execute the following command:

Example

# cd /media/dvd
# ./install.sh

In the example above, /media/dvd is the DVD mount point.

#### 5. Select the product for installation

The list of installation target products is displayed.

Type the number for the product to be installed, or "all", and press Enter.



- The FUJITSU Enterprise Postgres server component and WebAdmin can be installed on the same machine by selecting the "FUJITSU Enterprise Postgres server component" and the "WebAdmin component".

- Pgpool-II to be installed on the same server as the database server, you can install it on the same machine by selecting "FUJITSU Enterprise Postgres Server Component" and "Pgpool-II component". Pgpool-II to be installed on the same server as the application server, or on a dedicated server that is different from the database server and application server, see "Appendix J System Configuration when using Pgpool-II".
- If the selected product has been installed, a window for selecting reinstallation or multi-version installation is displayed for each product. Follow the on-screen instructions to select the installation method.

#### 6. Check the settings

The window for checking the installation information is displayed.

Type "y" and press Enter to start the installation.

To change the settings, type "c", press Enter, and follow the on-screen instructions. This option is not displayed if there is no information that can be modified.

If you have not set up WebAdmin during installation, refer to "Appendix B Setting Up and Removing WebAdmin" for details.

#### 7. Check the changed settings

If the installation information has been modified, the new installation information is displayed.

Type "y" and press Enter to start the installation.

To change the settings again, type "c" and press Enter.

#### 8. Display the installation status

When the installation is started, the installation status is displayed.

Note that even if WebAdmin setup fails, the process will continue.

If the following message is displayed, manually execute WebAdmin setup after installation is completed.

Refer to "B.1 Setting Up WebAdmin" for information on the WebAdmin setup procedure.

Starting setup.

Setting up WebAdmin.

ERROR: WebAdmin setup failed.

#### 9. Finish the installation

When the installation is complete, a message is displayed showing the status. If installation was successful, a message like the one shown below is displayed:

Installed successfully.



If an error occurs during the installation, read the error message and remove the cause of the error, and then reexecute the install.sh command.

#### 3.3 Installation in Silent Mode

Installation in silent mode can be performed only when the installation method is one of the following:

- New installation
- Multi-version installation



See

Refer to the FUJITSU Enterprise Postgres product website for information on installation in silent mode, such as the installation parameters and error messages.

The installation procedure is described below.

#### 1. Change to the superuser

Run the following command to switch to the superuser on the system.

\$ su Password:\*\*\*\*\*

#### 2. Mount the DVD drive

Insert the server program DVD into the DVD drive, and run the command given below.

Example

# mount -t iso9660 -r -o loop /dev/dvd /media/dvd

Here /dev/dvd is the device name for the DVD drive (which may vary depending on your environment), and /media/dvd is the mount point (which may need to be created before calling the command).



If the DVD was mounted automatically using the automatic mount daemon (autofs), "noexec" is set as the mount option, so the installer may fail to start. In this case, use the mount command to remount the DVD correctly, and then run the installation. Note that the mount options of a mounted DVD can be checked by executing the mount command without any arguments.

#### 3. Create an installation parameters CSV file

Consider the features that will be required for system operations, and then create an installation parameters CSV file that uses the following specification format.

```
sectionName, parameterName, value
sectionName, parameterName, value
:
```



The template for the installation parameters CSV file is "mountPoint/sample/sample.csv".

#### 4. Run the installation

Execute the following command:

#### Example

```
# cd /media/dvd
# ./silent.sh /home/work/inspara.csv
```

In the example above, /media/dvd is the DVD mount point, and /home/work/inspara.csv is the installation parameter CSV.

If the installer ends in an error, a message is output to the log file and return values are returned.

# Chapter 4 Setup

This chapter describes the setup procedures to be performed after installation completes.

# 4.1 Operating Method Types and Selection

This section describes how to operate FUJITSU Enterprise Postgres.

There are two methods of managing FUJITSU Enterprise Postgres operations - select one that suits your purposes:

The Operation Guide describes the operating method using WebAdmin, and the equivalent operating method using the server commands.

#### Simple operation management using a web-based GUI tool (WebAdmin)

Suitable when using frequently used basic settings and operations for operation management.

This method allows you to perform simple daily tasks such as starting the system before beginning business, and stopping the system when business is over, using an intuitive operation.

#### Usage method

Usage is started by using WebAdmin to create the instance.

By using an external scheduler and the pgx\_dmpall command, periodic backups can be performed, which can then be used in recovery using WebAdmin.



Do not use a server command other than pgx\_dmpall and pgx\_keystore or a server application. Operation modes that use server commands and server applications cannot be used in conjunction with WebAdmin. If used, WebAdmin will not be able to manage the instances correctly.

In addition, to perform a backup by copy command from the pgx\_dmpall command, select the operating method using the server commands.

Refer to Reference and the PostgreSQL Documentation for information on server commands and server applications.

#### Advanced operation management using server commands

When operating in a system that is automated by operation management middleware (Systemwalker Centric Manager, for example), this method allows you to use more detailed settings and operations and perform higher level operation management.

An overview of the operating method using the GUI, and its relationship with the operating method using the server commands, are shown below.

Refer to the Operation Guide for details.

Operation		Operation with the GUI	Operation with commands	
Setup	Creating an instance	WebAdmin is used. The server machine capacity, and the optimum parameter for operations using WebAdmin, are set automatically.	The configuration file is edited directly using the initdb command.	
	Creating a standby instance	WebAdmin is used.  WebAdmin performs a base backup of the source instance and creates a standby instance.	A standby instance is created using the pg_basebackup command.	
	Changing the configuration files	WebAdmin is used.	The configuration file is edited directly.	
Starting and stopping an instance		WebAdmin is used.	The pg_ctl command is used.	

Operation		Operation with the GUI	Operation with commands		
Creating a database			This is defined using pgAdmin of the GUI tool, or using the psql command or the application after specifying the DDL statement.		
Backing up the database		WebAdmin, or the pgx_dmpall command, is used.	It is recommended that the pgx_dmpall command be used. Recovery to the latest database can be performed.		
Database recovery		WebAdmin is used.	To use the backup that was performed using the pgx_dmpall command, the pgx_rcvall command is used.		
Monitoring	Database errors	The status in the WebAdmin window can be checked. (*1)	The messages that are output to the database server log are monitored (*1)		
	Disk space	The status in the WebAdmin window can be checked. A warning will be displayed if the free space falls below 20%. (*1)	This is monitored using the df command of the operating system, for example. (*1)		
	Connection status	This can be checked using pgAdmin of the GUI tool, or referencing pg_stat_activity of the standard statistics view from psql or the application.			

<sup>\*1:</sup> This can be used together with system log monitoring using operations management middleware (Systemwalker Centric Manager, for example).



Refer to "Periodic Operations" and "Actions when an Error Occurs" in the Operation Guide for information on monitoring and database recovery.

# 4.2 Preparations for Setup

This section describes the preparation required before setting up FUJITSU Enterprise Postgres.

# 4.2.1 Creating an Instance Administrator

Decide which OS user account will be assigned the instance administrator role. You can assign it to a new user or to an existing one, but you cannot assign it to the OS superuser (root).

The following example shows an OS user account with the name "fsepuser" being assigned the instance administrator role.

#### Example

# useradd fsepuser
# passwd fsepuser



The following note applies if using WebAdmin for operations:

- If the password is changed for the user account of the instance administrator, set the changed password using ALTER ROLE WITH ENCRYPTED PASSWORD. This will encrypt the password with the default MD5 algorithm. To encrypt the password with the SCRAM-SHA-256 algorithm, set the password encryption parameter to "scram-sha-256" using SET password\_encryption = 'scram-sha-256' prior to using the ALTER ROLE function.

### 4.2.2 Preparing Directories for Resource Deployment

Prepare the directories required when creating instances.

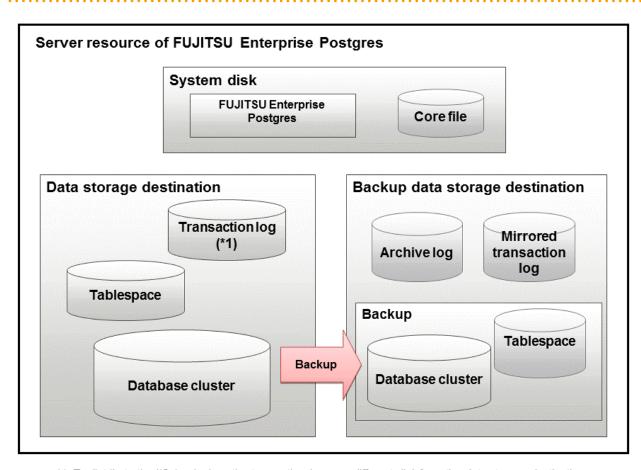
#### Considerations when deploying resources

The disk configuration on the resource deployment destination is important, because it affects not only recovery following disk corruption, but normal operation as well. The points for determining the disk configuration are as follows:

- 1. If the backup data storage destination and the data storage destination are both lost, it will not be possible to recover the data, so deploy them to separate disks.
- 2. To shorten the recovery time following a single disk fault, deploy the system disk and data storage destination to separate disks.
- 3. The backup data storage destination requires at least double the capacity of the data storage destination, so deploy it to the disk with the most space available.
- 4. When large amounts of data are updated, the write-to load for the data storage destination, transaction log storage destination, and backup data storage destination (mirrored transaction log) will also be great. For this reason, deploy them to separate disks, out of consideration for performance.



When using the volume manager provided by the operating system, be aware of which physical disk the file system has been created on, for example, by deploying the data storage destination and the backup data storage destination to separate disks.



<sup>\*1:</sup> To distribute the I/O load, place the transaction log on a different disk from the data storage destination.

Resource	Role	
Database cluster	The area where the database is stored. It is a collection of databases managed by an instance.	
Tablespace	Stores table files and index files in a separate area from the database cluster.	
	Specify a space other than that under the database cluster.	
Transaction log	Stores log information in preparation for a crash recovery or rollback.	
	This is the same as the WAL (Write Ahead Log).	
Archive log	Stores log information for recovery	
Mirrored transaction log (mirrored WAL)	Enables a database cluster to be restored to the state immediately before an error even if both the database cluster and transaction log fail when performing backup/recovery operations using the pgx_dmpall command or WebAdmin.	
Corefile	FUJITSU Enterprise Postgres process corefile output when an error occurs with a FUJITSU Enterprise Postgres process.	

#### Examples of disk deployment

The following are examples of disk deployment:

Number of disks	Disk	Deployment
3	System disk	FUJITSU Enterprise Postgres program
		Corefile
	Connected physical disk	Data storage destination, transaction log storage destination
	Connected physical disk	Backup data storage destination
2	System disk	FUJITSU Enterprise Postgres program
		Corefile
		Data storage destination, transaction log storage destination
	Connected physical disk	Backup data storage destination

#### Proposal for disk deployment using WebAdmin

To generate an instance using WebAdmin, we recommend an optimum deployment that takes into account the status of all disks at the time of instance generation, and items 1 to 3 in the "Considerations when deploying resources" subheading above, based on the limitations below (note that a different deployment can also be specified).

- The mount point does not include national characters
- The instance administrator has the proper permissions to read and write on the mount point

#### **Preparing directories**

The directories to be prepared depend on the way that you create the instances.

The following shows the directories that need to be prepared:

Directory to be prepared	Using WebAdmin	Using the initdb command
Data storage destination	Y (*1)	Y
Backup data storage destination	O (*1) (*4)	О

Directory to be prepared	Using WebAdmin	Using the initdb command
Transaction log storage destination	O (*1) (*2)	О
Corefile output destination	N (*3)	0

Y: Required

O: Optional

N: Not required

- \*1: WebAdmin automatically creates a directory
- \*2: The default is to create in a directory in the data storage destination. When it is necessary to distribute the I/O load for the database data and the transaction log, consider putting the transaction log storage destination on a different disk from the data storage destination
- \*3: The corefile path is as follows:

/var/tmp/fsep\_version/instanceAdminUser\_instanceNamePortNumber/core

version: product version\_WA\_architecture

Note: The product version is normally the version of WebAdmin used to create the instance. For example, WebAdmin 12 allows a user to create a FUJITSU Enterprise Postgres 9.6 instance on a database server having WebAdmin 9.6. In this case, because WebAdmin 9.6 is used to create the instance, the product version will be "96".

instanceAdminUser. operating system user name

PortNumber: port number specified when creating the instance

#### Example:

/var/tmp/fsep\_120\_WA\_64/naomi\_myinst27599/core

Note that resources placed in /var/tmp that have not been accessed for 30 days or more will be deleted by the default settings of the operating system. Consider excluding them from deletion targets or changing the output destination in the operating system settings.

To change the output destination, configure the core\_directory and the core\_contents parameters in postgresql.conf. Refer to "Parameters" in the Operation Guide for information on the settings for these parameters.

\*4: This directory is required when instance backup is enabled.



- The directories must meet the following conditions:
  - The directory owner must be the OS user account that you want to be the instance administrator
  - The directory must have write permission
  - The directory must be empty
- It is not possible to use a directory mounted by NFS (Network File System) when using WebAdmin.

#### Example

The following example shows the OS superuser creating /database/inst1 as the directory for storing the database data and changing the owner of the directory to the OS user account "fsepuser".

```
# mkdir -p /database/inst1
# chown -R fsepuser:fsepuser /database/inst1
# chmod 700 /database/inst1
```

# 4.2.3 Estimating Resources

Estimate the resources to be used on the FUJITSU Enterprise Postgres.

Refer to "Appendix E Estimating Database Disk Space Requirements" for information on estimating database disk space requirements.

Refer to "Parameters automatically set by WebAdmin according to the amount of memory" when creating multiple instances with WebAdmin.

Refer to "Appendix F Estimating Memory Requirements" when creating instances with the initdb command, to estimate memory usage.

### 4.2.4 Configuring Corefile Names

If a process crashes, a corefile for the process will be generated by the operating system. If a corefile is generated with the same name as an existing corefile generated for a different process, the newly-generated corefile will overwrite the previously dumped corefile. To prevent this, configure a unique corefile name for each crash by appending the process ID, program name, and datetime.

Corefile names can be configured using the "kernel.core\_pattern" and "kernel.core\_uses\_pid" kernel parameters.

Refer to the "man page" in "core(5)" for information on how to use these parameters.

Note that with regard to the location for storing corefiles, the operating system settings take precedence over the core\_directory parameter of postgresql.conf.

If you specify systemd-coredump as the core\_pattern, the core file is not placed in the location specified by the core\_directory parameter. See the systemd-coredump (8) man page for the location of core files.

Use coredumpctl to retrieve core files. For more information about using coredumpctl, see the coredumpctl (1) man page.

# 4.3 Creating Instances

There are two methods that can be used to create an instance:

- 4.3.1 Using WebAdmin
- 4.3.2 Using the initdb Command

#### Creating multiple instances

Multiple instances can be created.

The memory allocated needs to be adjusted when multiple instances are created with WebAdmin (refer to "Parameters automatically set by WebAdmin according to the amount of memory" for details).

#### Features that cannot be set up using WebAdmin

The "Storage data protection using transparent data encryption" feature cannot be set up using WebAdmin.

To set up this feature in an instance created by WebAdmin, perform the additional setup tasks detailed in "Storage Data Protection using Transparent Data Encryption" in the Operation Guide.



- Instances created using the initdb command (command line instances) can be managed using WebAdmin, however, they must first be imported into WebAdmin. Refer to "4.3.1.4 Importing Instances" for details.

- Always use WebAdmin to delete instances that were created or imported using WebAdmin. Because WebAdmin management information cannot be deleted, WebAdmin will determine that the instance is abnormal.
- Databases with the names 'template0' and 'template1' are automatically created when an instance is created.

  These databases are used as the templates for databases created later. Furthermore, a default database with the name 'postgres' is automatically created, which will be used with FUJITSU Enterprise Postgres commands. It is important that you do not delete these databases created by default.

# 4.3.1 Using WebAdmin

This section describes how to create an instance using WebAdmin.

WebAdmin must be set up correctly before it can be used. Refer to "B.1 Setting Up WebAdmin" for details. Additionally, if WebAdmin needs to be configured to use an external repository database, refer to "B.3 Using an External Repository for WebAdmin" for details.

It is recommended to use the following browsers with WebAdmin:

- Internet Explorer 11
- Microsoft Edge (Build41 or later)

WebAdmin will work with other browsers, such as Firefox and Chrome, however, the look and feel may be slightly different.

Configure your browser to allow cookies and pop-up requests from the server on which FUJITSU Enterprise Postgres is installed.

Refer to "Appendix A Recommended WebAdmin Environments" for information on how to change the pop-up request settings and other recommended settings.



- WebAdmin does not run in Windows(R) safe mode.
- If the same instance is operated from multiple WebAdmin windows, it will not work correctly.
- If the same instance is operated from multiple WebAdmin versions, it will not work correctly. Always use the latest version of WebAdmin for instance operations.
- For efficient use of WebAdmin, it is recommended not to use the browser [Back] and [Forward] navigation buttons, the [Refresh] button, and context-sensitive menus, including equivalent keyboard shortcuts.
- Copying and pasting the WebAdmin URLs are not supported. Additionally, bookmarking of WebAdmin URLs is not supported.
- It is recommended to match the language between the instance server locale and WebAdmin.
- WebAdmin supports only two languages: English and Japanese.
- It is recommended to change the WebAdmin language setting from the instance details page only.
- It is recommended to operate WebAdmin using the WebAdmin launcher.
- WebAdmin uses the labels "Data storage path", "Backup storage path" and "Transaction log path" to indicate "data storage destination", "backup data storage destination" and "transaction log storage destination" respectively. In this manual these terms are used interchangeably.
- If the browser was not operated for a fixed period (about 30 minutes), the session will time out and the login page will be displayed again for the next operation.
- Port access permissions
  - If a port is blocked (access permissions have not been granted) by a firewall, enable use of the port by granting access. Refer to the vendor document for information on how to grant port access permissions.

    Consider the security risks carefully when opening ports.
- When creating or importing an instance in WebAdmin, set the log\_directory parameter in postgresql.conf to '/var/tmp/fsep\_version/ instanceAdminUser\_instanceNamePortNumber/log'. Note that resources placed in /var/tmp that have not been accessed for 30 days or more will be deleted by the default settings of the operating system. Therefore, consider excluding instances created using WebAdmin from deletion targets in the operating system settings if you need to stop those instances for a long time.

#### 4.3.1.1 Logging in to WebAdmin

This section describes how to log in to WebAdmin.

#### Startup URL for WebAdmin

In the browser address bar, type the startup URL of the WebAdmin window in the following format:

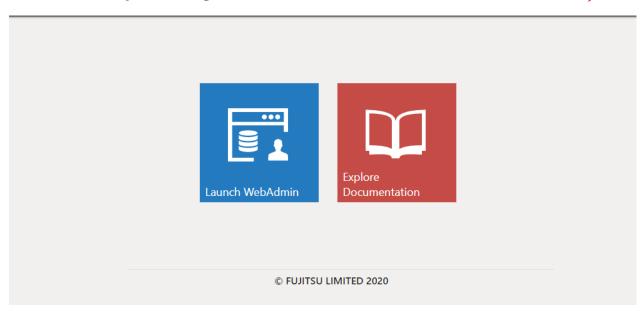
http://hostNameOrIpAddress:portNumber/

- hostNameOrIpAddress: Host name or IP address of the server where WebAdmin is installed.
- portNumber: Port number of WebAdmin. The default port number is 27515.

The startup URL window shown below is displayed. From this window you can log in to WebAdmin or access the product documentation.

# **FUJITSU Enterprise Postgres**





#### Logging in to WebAdmin

Click [Launch WebAdmin] in the startup URL window to start WebAdmin and display the login window. Enter the instance administrator user ID (operating system user account name) and password, and log in to WebAdmin. User credential (instance administrator user ID and password) should not contain hazardous characters. Refer to "Appendix C WebAdmin Disallow User Inputs Containing Hazardous Characters".

### 4.3.1.2 Creating an Instance

This section describes how to create an instance.

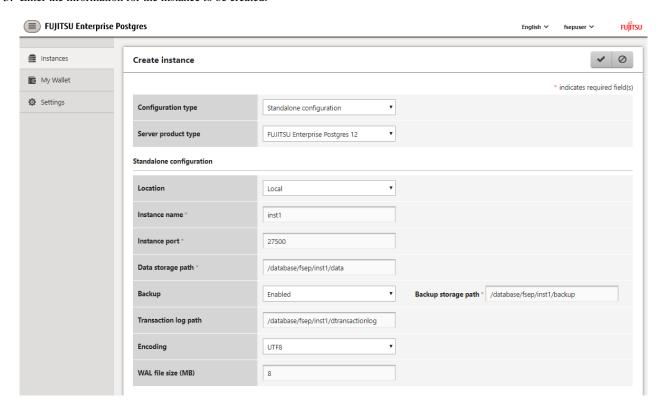


WebAdmin calculates values using the formula indicated in "Managing Kernel Resources" under "Server Administration" in the PostgreSQL Documentation, and configures these in the kernel parameters. Refer to "Appendix H Configuring Kernel Parameters" for information on configuring parameter names.

Refer to "Appendix D Configuring Parameters" for information on the postgresql.conf values required to derive the set values.

- 1. Start WebAdmin, and log in to the database server.
- 2. In the [Instances] tab, click -

3. Enter the information for the instance to be created.



#### Enter the following items:

- [Configuration type]: Whether to create a standalone instance or an instance that is part of a cluster.
- [Server product type]: Sets which of the following instances to create:
  - FUJITSU Enterprise Postgres 9.5 Instances
  - FUJITSU Enterprise Postgres 9.6 Instances
  - FUJITSU Enterprise Postgres 10 Instances
  - FUJITSU Enterprise Postgres 11 Instances
  - FUJITSU Enterprise Postgres 12 Instances

The default is "FUJITSU Enterprise Postgres 12".

WebAdmin can create and manage instances compatible with the following, but new functionality in FUJITSU Enterprise Postgres 12 may not support the instance or it may be disabled.

- FUJITSU Enterprise Postgres 9.5
- FUJITSU Enterprise Postgres 9.6
- FUJITSU Enterprise Postgres 10
- FUJITSU Enterprise Postgres 11
- [Location]: Whether to create the instance in the server that the current user is logged into, or in a remote server. The default is "Local", which will create the instance in the server machine where WebAdmin is currently running.
- [Instance name]: Name of the database instance to manage

The name must meet the conditions below:

- Maximum of 16 characters
- The first character must be an ASCII alphabetic character
- The other characters must be ASCII alphanumeric characters

- [Instance port]: Port number of the database server
- [Data storage path]: Directory where the database data will be stored
- [Backup]: Whether to enable or disable the WebAdmin backup feature. The default is "Enabled". Select "Disabled" to disable all backup and restore functionality for the instance. If "Enabled" is selected, enter the following item:
  - [Backup storage path]: Directory where the database backup will be stored
- [Transaction log path]: Directory where the transaction log will be stored
- [Encoding]: Database encoding system
- [WAL file size]: Allow the WAL file size to be set when creating an instance. The default is 16 MB if the field is blank. The size specified must be a power of 2 between 1 and 1024. This option is not available for standby instances.

If "Remote" is selected for [Location], enter the following additional items:

- [Host name]: Name of the host where the instance is to be created
- [Operating system credential]: Operating system user name and password for the remote machine where the instance is to be created
- [Remote WebAdmin port for standalone]: Port in which WebAdmin is accessible in the remote machine



- Refer to "4.2.2 Preparing Directories for Resource Deployment" - "Considerations when deploying resources" for information on points to consider when determining the data storage path, backup storage path, and transaction log path.

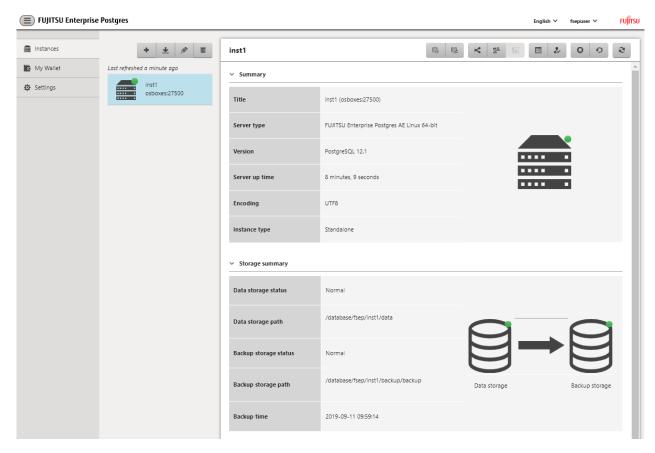
- The following items can be modified after the instance has been created. These items cannot be modified on instances that have compatibility with FUJITSU Enterprise Postgres 9.5.
  - Instance name
  - Port number
  - Backup storage path

Refer to "Editing instance information" for details.

- Do not specify directories that include symbolic link or multibyte characters when specifying the data storage destination or backup data storage destination.
- In the instance that is created using WebAdmin, the locale of the character set to be used in the database, and the locale of the collating sequence, are fixed using C.
- For enhanced security, WebAdmin encrypts the superuser password using SCRAM-SHA-256 authentication for all FUJITSU Enterprise Postgres 10 or later instances. The client/driver must therefore support SCRAM-SHA-256 authentication if they need to connect to FUJITSU Enterprise Postgres 10 or later instances created by WebAdmin with superuser credentials.
- Host name and Operating system credential (Operating system user name and password) should not contain hazardous characters. Refer to "Appendix C WebAdmin Disallow User Inputs Containing Hazardous Characters".
- 4. Click **v** to create the instance.

If the instance is created successfully, a message indicating the same will be displayed.

5. The instance will be started when it is created successfully.



6. Back up the basic information that was set

Back up the WebAdmin management information periodically to ensure operational continuity when a fault occurs on the system disk. Follow the procedure below to perform the backup.

- a. Stop the WebAdmin server. Refer to "B.1.3 Stopping the Web Server Feature of WebAdmin" for details.
- b. Back up the following directory:

webAdminInstallDir/data/fepwa

# 4.3.1.3 Changing Instance Settings

You can change the information that is set when an instance is created.

Change the following settings to suit the operating and management environment for FUJITSU Enterprise Postgres.

- Instance configuration
  - Character encoding
  - Communication
  - SQL options
  - Memory
  - Streaming replication
- Changing client authentication information
- Editing instance information

# **Information**

These settings are the same as the parameters that can be set in the files shown below. Refer to "Appendix D Configuring Parameters" for information on the equivalence relationship between the item name and the parameter.

- postgresql.conf
- pg\_hba.conf



The files shown below can also be modified directly, however if a parameter not described in "Appendix D Configuring Parameters" was edited by mistake, WebAdmin may not run correctly.

- postgresql.conf
- pg\_hba.conf

#### Instance configuration

- 1. Start WebAdmin and log in to the database server.
- 2. In the [Instances] tab, click
- 3. Click to change the configuration.
- 4. Click to save your changes.



#### See

Select a client-side encoding system that can be converted to/from the database encoding system. Refer to "Automatic Character Set Conversion Between Server and Client" in "Server Administration" in the PostgreSQL Documentation for information on the encoding system combinations that can be converted.

#### Changing client authentication information

- 1. Start WebAdmin and log in to the database server.
- 2. In the [Instances] tab, click ....

Click **\rightarrow** to register new authentication information.

To change authentication information, select the information, and then click



To delete authentication information, select the information, and then click iii.



When creating the instance, do not delete the entry below, because it is a connection required for WebAdmin to monitor the operational status of the database:

Type=local, Database=all, User=all, and Method=md5

#### Editing instance information

Use the [Edit instance] page to modify the following items for an instance:

- Instance name
- Port number

- Backup storage path
- 1. In the [Instances] tab, click . The [Edit instance] page is displayed.
- 2. Modify the relevant items.

If [Backup storage path] is changed, [Backup management] is enabled. Select the required option:

Retain existing backup: Create a backup in [Backup storage path] and retain the existing backup in its original location.

Copy existing backup to new path: Copy the existing backup to [Backup storage path]. A new backup will not be created. The existing backup will be retained in its original location.

Move existing backup to new path: Move the existing backup to [Backup storage path]. A new backup will not be created.

Remove existing backup: Create a backup in [Backup storage path]. The existing backup will be removed.

3. Click to save your changes.



- The [Edit instance] page is also displayed when the user selects 'Navigate to the "Edit instance" page' from the [Anomaly Error] dialog box. Refer to "Anomaly Detection and Resolution" in the Operation Guide for information on what takes place when an anomaly is detected.

.....

- When [Instance name] or [Instance port] is modified, the log\_directory and core\_directory parameters in postgresql.conf are updated. Also, the specified directories are created if they do not exist.

Refer to "4.3.1.4 Importing Instances" for information on the format of these directories.

### 4.3.1.4 Importing Instances

Instances can be created using WebAdmin, or via the command line using the initdb command. Instances created using the initdb command (command line instances) can be managed using WebAdmin, however, they must first be imported into WebAdmin.

This section explains how to import command line instances into WebAdmin.

- 1. In the [Instances] tab, click . The [Import instance] page is displayed.
- 2. Enter the information for the instance being imported. Refer to "4.3.1.2 Creating an Instance" for information on the items that need to be entered.
- 3. Click to import the instance.



- Importing neither starts nor stops the instance.
- The following restrictions apply to instance import:
  - Any instance already managed by WebAdmin cannot be imported again.
  - The postgresql.conf file must be located in the same directory as [Data storage path].
  - Read/write permissions are required for [Data storage path].
  - The location specified in postgresql.conf for the following files must not have been changed:
    - hba\_file
    - ident\_file
  - If the instance is part of a cluster that is monitored by Mirroring Controller, WebAdmin will be unable to detect the Mirroring Controller settings.
  - Instances making use of Mirroring Controller functionality should not be imported, because subsequent operations on those instances may cause unexpected and undesirable side-effects.

- It is not possible to import and operate an instance that uses a directory mounted by Network File System (NFS).
- You must make the following changes to the parameters in postgresql.conf prior to importing the instance in WebAdmin.

Parameter	Requirements
port	The port parameter should be uncommented.

The log\_directory and core\_directory parameters in postgresql.conf are updated during import. Also, the specified directories are created if they do not exist.

The format of these directories is as follows:

log\_directory: '/var/tmp/fsep\_version/instanceAdminUser\_instanceNamePortNumber/log'

core\_directory: '/var/tmp/fsep\_version/instanceAdminUser\_instanceNamePortNumber/core'

version: product version\_WA\_architecture

instanceAdminUser: operating system user name

PortNumber: port number specified when creating the instance

Examples:

log\_directory: '/var/tmp/fsep\_120\_WA\_64/naomi\_myinst27599/log'

core\_directory: '/var/tmp/fsep\_120\_WA\_64/naomi\_myinst27599/core'

- When a standby instance is imported, a valid entry, using the IP address of the standby instance, must exist in the pg\_hba.conf file of the master instance to allow the standby instance to connect to the master instance.
- When a standby instance is imported, the value for "host" in the primary\_conninfo parameter of postgresql.auto.conf should match the host name of the master instance.
- When a standby instance is imported, you cannot specify "passfile" in the primary\_conninfo parameter of postgresql.auto.conf. Be sure to specify "password".
- Instances created by other operating systems cannot be imported.
- If a FUJITSU Enterprise Postgres 10 or later instance is being imported while it is running, WebAdmin will encrypt the superuser password using SCRAM-SHA-256 authentication.

# 4.3.2 Using the initdb Command

This section describes the procedure to create an instance using the initdb command.



If a port is blocked (access permissions have not been granted) by a firewall, enable use of the port by granting access. Refer to the vendor document for information on how to grant port access permissions.

Consider the security risks carefully when opening ports.

### 4.3.2.1 Editing Kernel Parameters

Refer to "Appendix H Configuring Kernel Parameters" prior to editing these settings.

After the settings are complete, check the command specifications of the relevant operating system and restart the system if required.

### 4.3.2.2 Creating an Instance

Create an instance, with the database cluster storage destination specified in the PGDATA environment variable or in the -D option. Furthermore, the user that executed the initdb command becomes the instance administrator.



- Instances created using the initidb command (command line instances) can be managed using WebAdmin, however, they must first be imported into WebAdmin. Refer to "4.3.1.4 Importing Instances" for details.

- If creating multiple instances, ensure that there is no duplication of port numbers or the directories that store database clusters.



Refer to "initdb" in "Reference" in the PostgreSQL Documentation for information on the initdb command.

The procedure to create an instance is described below.

1. Use the OS user account that you want as the instance administrator.

Connect with the server using the OS user account that you want as the instance administrator.

You cannot use the OS superuser (root).

The following example shows the OS superuser connected to the server being changed to the OS user account "fsepuser".

#### Example

# su fsepuser

2. Configure the environment variables

Configure the environment variables in the server with the newly created instance.

Set the following environment variables:

- PATH environment variables

Add the installation directory "/bin".

- MANPATH environment variables

Add the installation directory "/share/man".

- LD\_LIBRARY\_PATH environment variables

Add the installation directory "/lib".

#### Example

The following example configures environment variables when the installation directory is "/opt/fsepv<*x*>server64".

Note that "<*x*>" indicates the product version.

#### sh, bash

- \$ PATH=/opt/fsepv<x>server64/bin:\$PATH ; export PATH
- \$ MANPATH=/opt/fsepv<x>server64/share/man:\$MANPATH ; export MANPATH
- \$ LD\_LIBRARY\_PATH=/opt/fsepv<x>server64/lib:\$LD\_LIBRARY\_PATH ; export LD\_LIBRARY\_PATH

#### csh, tcsh

- \$ setenv PATH /opt/fsepv<x>server64/bin:\$PATH
- \$ setenv MANPATH /opt/fsepv<x>server64/share/man:\$MANPATH
- setenv LD\_LIBRARY\_PATH /opt/fsepv<x>server64/lib:\$LD\_LIBRARY\_PATH

#### 3. Create a database cluster

Create the database cluster with the initdb command, specifying the storage destination directory.

Specify the transaction log storage destination and the locale setting option as required.

#### Example

```
$ initdb -D /database/instl --waldir=/transaction/instl --lc-collate="C" --lc-ctype="C" --
encoding=UTF8
```



In some features, instance names are requested, and those names are required to uniquely identify the instance within the system. These features allow names that conform to WebAdmin naming conventions, so refer to the following points when determining the names:

- Maximum of 16 characters
- The first character must be ASCII alphabetic character
- The other characters must be ASCII alphanumeric characters



- To balance I/O load, consider deploying the transaction log storage destination to a disk device other than the database cluster storage destination and the backup data storage destination.

.....

- Specify "C" or "POSIX" for collation and character category. Performance deteriorates if you specify a value other than "C" or "POSIX", although the behavior will follow the rules for particular languages, countries and regions. Furthermore, this may need to be revised when running applications on systems with different locales. For example, specify as follows:

```
initdb --locale="C" --lc-messages="C"
initdb --lc-collate="C" --lc-ctype="C"
```

- Specify the same string in the LANG environment variable of the terminal that starts FUJITSU Enterprise Postgres as was specified in lc-messages of initdb (lc\_messages of postgresql.conf). If the same string is not specified, messages displayed on the terminal that was started, as well as messages output to the log file specified in the -l option of the pg\_ctl command or the postgres command used for startup, may not be output correctly.
- Specify an encoding system other than SQL\_ASCII for the database. If SQL\_ASCII is used, there is no guarantee that the encryption system for data in the database will be consistent, depending on the application used to insert the data.



Refer to "Locale Support" in "Localization" in "Server Administration" in the PostgreSQL Documentation for information on locales.

#### 4. Set port number.

Specify a port number in the port parameter of postgresql.conf. Ensure that the specified port number is not already used for other software. If a port number is not specified, "27500" is selected.

Register the specified port numbers in the /etc/services file if WebAdmin is used to create other instances. WebAdmin uses the /etc/services file to check if port numbers specified as available candidates have been duplicated.

Register any name as the service name.

5. Set the corefile output destination.

Specify the output destination of the corefile, which can later be used to collect information for investigation, by setting the core\_directory and core\_contents parameters of postgresql.conf.



Refer to "Parameters" in the Operation Guide for information on the settings for these parameters.

6. Set the backup storage destination.

Specify the backup data storage destination and other backup settings when backup is to be performed as a provision against database errors



#### See

Refer to "Backup Methods" in the Operation Guide for information on specifying backup settings.

7. Start an instance.

Start with the start mode of the pg\_ctl command.

If either of the following conditions are met, the message "FATAL:the database system is starting up(11189)" may be output.

- An application, command, or process connects to the database while the instance is starting
- An instance was started without the -W option specified

This message is output by the pg\_ctl command to check if the instance has started successfully.

Therefore, ignore this message if there are no other applications, commands, or processes that connect to the database.

#### Example

\$ pg\_ctl start -D /database/inst1



#### See

Refer to "pg\_ctl" in "Reference" in the PostgreSQL Documentation for information on the pg\_ctl command.



If the -W option is specified, the command will return without waiting for the instance to start. Therefore, it may be unclear as to whether instance startup was successful or failed.

# 4.4 Configuring Remote Connections

This section describes the settings required when connecting remotely to FUJITSU Enterprise Postgres from a database application or a client command.

#### 4.4.1 When an Instance was Created with WebAdmin

#### Settings related to connection

The default is to accept connections from remote computers to the database.

Change "listen\_addresses" in postgresql.conf to modify the default behavior.

Refer to "Appendix D Configuring Parameters" for information on postgresql.conf.

#### **Client Authentication Information settings**

The following content is set by default when WebAdmin is used to create an instance.

- Authentication of remote connections from local machines is performed.

When changing Client Authentication Information, select [Client Authentication] from [Setting], and then change the settings.

### 4.4.2 When an Instance was Created with the initdb Command

#### **Connection settings**

The default setting only permits local connections from the client to the database. Remote connections are not accepted.

Change "listen\_addresses" in postgresql.conf to perform remote connection.

All remote connections will be allowed when changed as shown below.

#### Example

```
listen_addresses = '*'
```

Also, configure the parameters shown below in accordance with the applications and number of client command connections.

Parameter name	Parameter description
superuser_reserved_connections	Number of connections reserved for database maintenance, for example backup or index rebuilding. If you need to simultaneously perform a large number of processes that exceed the default value, change this value accordingly.
max_connections	Set the value as:  numberOfSimultaneousConnectionsToInstance + superuser_reserved_connections

#### Client authentication information settings

When trying to connect from a client to a database, settings are required to determine whether the instance permits connections from the client - if it does, then it is possible to make settings to determine if authentication is required.



Refer to "The pg\_hba.conf File" in "Server Administration" in the PostgreSQL Documentation for details.

# 4.5 Other Settings

This section describes settings that are useful for operations.

# 4.5.1 Error Log Settings

This section explains the settings necessary to monitor errors in applications and operations, and to make discovering the causes easier.

Make error log settings only when instances are created with the initdb command.

When creating instances with WebAdmin, these settings are already made and hence do not need to be set.

Furthermore, some parameters are used by WebAdmin, and if changed, may cause WebAdmin to no longer work properly. Refer to "Appendix D Configuring Parameters" for details.



Set the output destination for the system log to the server log so that it cannot be viewed by administrators of other instances.

Application errors are output to the system log or server log. The output destination directory for the system log and server log should have access permissions set so that they cannot be viewed by people other than the instance administrator.

Edit the following parameters in postgresql.conf:

Parameter name	Parameter description	How to enable the settings
syslog_ident	Used to specify labels to attach to messages, so that these can be identified when output to the system log if more than one FUJITSU Enterprise Postgres is used.	reload option of the pg_ctl mode
logging_collector	Specify "on" to ensure that messages are output by FUJITSU Enterprise Postgres to the server log file. The server log file is created in the pg_log directory in the database cluster.	restart option of the pg_ctl mode
log_destination	Specify "stderr,syslog" to output messages from FUJITSU Enterprise Postgres to the screen and either the system log or the event log.	reload option of the pg_ctl mode
log_line_prefix	Specify information to be added at the start of messages output by an instance. This information is useful for automatic monitoring of messages.	reload option of the pg_ctl mode
	You can output the SQLSTATE value, output time, executing host, application name, and user ID.	
	Refer to "What To Log" in the PostgreSQL Documentation for details.	
	Example: log_line_prefix = '%e: %t [%p]: [%l-1] user = %u,db = %d,remote = %r app = %a '	



- If you want fewer application errors being output to the system log, refer to "When To Log" and "What To Log" in the PostgreSQL Documentation for information on how to reduce the output messages.
- If you want to separate errors output from other software, refer to "Where To Log" in the PostgreSQL Documentation to change the output destination to the server log file rather than the system log.

# 4.5.2 Configuring Automatic Start and Stop of an Instance

You can automatically start or stop an instance when the operating system on the database server is started or stopped.

Use the following procedure to configure automatic start and stop of an instance.

Note that, if an instance is started in a failover operation, the cluster system will control the start or stop, therefore this feature should not be used. Also, when performing database multiplexing, refer to "Enabling Automatic Start and Stop of Mirroring Controller and Multiplexed Instances" in the Cluster Operation Guide (Database Multiplexing).

Note that "<*x*>" in paths indicates the product version.

If using an rc script

1. Set the OS user account of the instance administrator

If you logged in using the OS user account of the instance administrator, set the environment variables required for starting the instance.

2. Copy the shell script and set access permissions

Copy the sample shell script (stored in the location shown below) to /etc/init.d, then set and register the access rights. After that, revise the shell script according to the target instance.

fujitsuEnterprisePostgresInstallDir/share/rc\_fsepsvoi.sample

The content to be revised is as follows:

- In the "SYMDATA" variable, set the data storage destination.
- In the "SYMUSER" variable, set the OS user account of the instance administrator mentioned in step 1.

- Ensure that the file name in "/var/lock/subsys" of the "LOCKFILE" variable is the same as the shell script name.

#### Example

In the following example, the installation directory is "/opt/fsepv<x>server64", and the shell script has been created as "rc\_fsepsvoi\_inst1".

```
# cp /opt/fsepv<x>server64/share/rc_fsepsvoi.sample /etc/init.d/rc_fsepsvoi_inst1
# chmod 755 /etc/init.d/rc_fsepsvoi_inst1
```

#### 3. Register and enable automatic start and stop

As the OS superuser, execute the chkconfig command to register and enable the script.

Execute "chkconfig --add" to register the script, and execute "chkconfig --level" to set the run level and enable the script.

#### Example

```
# chkconfig --add rc_fsepsvoi_inst1
# chkconfig --level 35 rc_fsepsvoi_inst1 on
```

#### If using systemd



When you start an instance using systemd, you should wait for time correction, network setup, and so on.

#### 1. Create a unit file

Copy the unit file sample stored in the directory below, and revise it to match the target instance.

fujitsuEnterprisePostgresInstallDir/share/fsepsvoi.service.sample

#### Example

In the following example, the installation directory is "/opt/fsepv<x>server64", and the instance name is "inst1".

```
# cp /opt/fsepv<x>server64/share/fsepsvoi.service.sample /usr/lib/systemd/system/
fsepsvoi_inst1.service
```

Revise the underlined portions of the options below in the unit file.

Section	Option	Specified value	Description
Unit	Description	FUJITSU Enterprise Postgres  instanceName	Specifies the feature overview.  Specifies the name of the target instance. (*1)
Service	Service ExecStart '/bin/bash-c'installDir/bin/pgx_symstd start installDir dataStorageDestinationDir	Command to be executed when the service is started.	
	ExecStop	'/bin/bash -c ' <i>installDir</i> /bin/pgx_symstd stop <u>installDir dataStorageDestinationDir</u>	Command to be executed when the service is stopped.
	ExecReload	'/bin/bash -c ' <i>installDir</i> /bin/pgx_symstd reload <i>installDir dataStorageDestinationDir</i> '	Command to be executed when the service is reloaded
	User	<u>User</u>	OS user account of the instance administrator.
	Group	Group	Group to which the instance administrator user belongs.

<sup>\*1:</sup> The instance name should be as follows:

If WebAdmin is used to create the instance: instanceName

If the initdb command is used to create the instance: nameThatIdentifiesTheInstance

The naming conventions for the instance name or for identifying the instance are as follows:

- Up to 16 bytes
- The first character must be an ASCII alphabetic character
- The other characters must be ASCII alphanumeric characters
- 2. Enable automatic start and stop

As the OS superuser, use the systemctl command to enable automatic start and stop.

#### Example

# systemctl enable fsepsvoi instl.service

# 4.5.3 Settings when Using the Features Compatible with Oracle Databases

To use the features compatible with Oracle databases, create a new instance and execute the following command for the "postgres" and "template1" databases:

```
CREATE EXTENSION oracle_compatible;
```

Features compatible with Oracle databases are defined as user-defined functions in the "public" schema created by default when database clusters are created, so they can be available for all users without the need for special settings.

For this reason, ensure that "public" (without the double quotation marks) is included in the list of schema search paths specified in the search\_path parameter.

There are also considerations for use the features compatible with Oracle databases. Refer to "Precautions when Using the Features Compatible with Oracle Databases" in the Application Development Guide for details.

# 4.6 Integration with Message-Monitoring Software

To monitor messages output by FUJITSU Enterprise Postgres using software, configure the product to monitor SQLSTATE, instead of the message text - this is because the latter may change when FUJITSU Enterprise Postgres is upgraded.

Configure FUJITSU Enterprise Postgres to output messages in a format that can be read by the message-monitoring software by specifying "%e" in the log\_line\_prefix parameter of postgresql.conf to output the SQLSTATE value.

A setting example is shown below - it outputs the output time, executing host, application name, and user ID, in addition to the SQLSTATE value.

#### Example

```
log_line_prefix = '%e: %t [%p]: [%l-1] user = %u,db = %d,remote = %r app = %a '
```



#### See

Refer to "What To Log" in the PostgreSQL Documentation for information on how to configure the settings.

# 4.7 Setting Up and Removing OSS

This section explains how to set up OSS supported by FUJITSU Enterprise Postgres.

If you want to use OSS supported by FUJITSU Enterprise Postgres, follow the setup procedure.

If you decide not to use the OSS supported by FUJITSU Enterprise Postgres, follow the removing procedure.



- In this section, the applicable database that enables the features of each OSS is described as "postgres".

- Execute CREATE EXTENSION for the "template1" database also, so that each OSS can be used by default when creating a new database

Refer to "OSS Supported by FUJITSU Enterprise Postgres" in the General Description for information on OSS other than those described below.

### 4.7.1 oracle fdw

### 4.7.1.1 Setting Up oracle\_fdw

- 1. Add the path of the OCI library to the environment variable. The available version of the OCI library is 11.2 or later. Add the installation path of the OCI library to the LD\_LIBRARY\_PATH environment variable.
- 2. As superuser, run the following command:

```
$ su -
Password:*****
# cp -r /opt/fsepv<x>server64/OSS/oracle_fdw/* /opt/fsepv<x>server64
```

3. If a file named libclntsh.so.11.1 does not exist in your OCI library, create a symbolic link with the name libclntsh.so.11.1 to libclntsh.so.xx.1 (xx is the version of the OCI library).

```
# ln -s libclntsh.so.12.1 libclntsh.so.11.1
```

- 4. Restart FUJITSU Enterprise Postgres.
- 5. Execute CREATE EXTENSION for the database that will use this feature. Use the psql command to connect to the "postgres" database.

```
postgres=# CREATE EXTENSION oracle_fdw;
CREATE EXTENSION
```

# **Information**

- If the OCI library is not installed on the server, install it using the Oracle client or Oracle Instant Client. Refer to the relevant Oracle manual for information on the installation procedure.
- If the version of the OCI library is updated, change the path of the OCI library in the LD\_LIBRARY\_PATH environment variable to the updated path. Also, re-create the symbolic link named libclntsh.so.11.1 if necessary.



This feature cannot be used on instances created in WebAdmin. It can only be used via server commands.

### 4.7.1.2 Removing oracle\_fdw

1. Execute DROP EXTENSION for the database that will use this feature. Use the psql command to connect to the "postgres" database.

```
postgres=# DROP EXTENSION oracle_fdw CASCADE;
DROP EXTENSION
```

2. As superuser, run the following command:

```
$ su -
Password:*****
# rm -rf /opt/fsepv<x>server64/filesCopiedDuringSetup
```



The files copied during setup can be checked below.

# find /opt/fsepv<x>server64/OSS/oracle\_fdw

### 4.7.2 pg\_bigm

# 4.7.2.1 Setting Up pg\_bigm

Set the postgresql.conf file parameters.
 Add "pg\_bigm" to the shared\_preload\_libraries parameter.

2. As superuser, run the following command:

```
$ su -
Password:*****
# cp -r /opt/fsepv<x>server64/OSS/pg_bigm/* /opt/fsepv<x>server64
```

.....

- 3. Restart FUJITSU Enterprise Postgres.
- 4. Execute CREATE EXTENSION for the database that will use this feature. Use the psql command to connect to the "postgres" database.

```
postgres=# CREATE EXTENSION pg_bigm;
CREATE EXTENSION
```

### 4.7.2.2 Removing pg\_bigm

Execute DROP EXTENSION for the database that will use this feature.
 Use the psql command to connect to the "postgres" database.

```
postgres=# DROP EXTENSION pg_bigm CASCADE;
DROP EXTENSION
```

2. As superuser, run the following command:

```
$ su -
Password:*****
# rm -rf /opt/fsepv<x>server64/filesCopiedDuringSetup
```



The files copied during setup can be checked below.

```
# find /opt/fsepv<x>server64/OSS/pg_bigm
```

- Set the postgresql.conf file parameters.Delete "pg\_bigm" to the shared\_preload\_libraries parameter.
- 4. Restart FUJITSU Enterprise Postgres.

# 4.7.3 pg\_hint\_plan

### 4.7.3.1 Setting Up pg\_hint\_plan

Set the postgresql.conf file parameters.
 Add "pg\_hint\_plan" to the "shared\_preload\_libraries" parameter.

2. As superuser, run the following command:

```
$ su -
Password:*****
# cp -r /opt/fsepv<x>server64/OSS/pg_hint_plan/* /opt/fsepv<x>server64
```

- 3. Restart FUJITSU Enterprise Postgres.
- 4. Run CREATE EXTENSION for the database that uses this feature.

The target database is described as "postgres" here.

Use the psql command to connect to the "postgres" database.

```
postgres=# CREATE EXTENSION pg_hint_plan;
CREATE EXTENSION
```



#### See

Refer to "Optimizer Hints" in the Application Development Guide for details.

### 4.7.3.2 Removing pg\_hint\_plan

1. Execute DROP EXTENSION for the database that will use this feature. Use the psql command to connect to the "postgres" database.

```
postgres=# DROP EXTENSION pg_hint_plan CASCADE;
DROP EXTENSION
```

2. As superuser, run the following command:

```
$ su -
Password:*****
# rm -rf /opt/fsepv<x>server64/filesCopiedDuringSetup
```



The files copied during setup can be checked below.

# find /opt/fsepv<x>server64/OSS/pg\_hint\_plan

- Set the postgresql.conf file parameters.
   Delete "pg\_hint\_plan" to the shared\_preload\_libraries parameter.
- 4. Restart FUJITSU Enterprise Postgres.

# 4.7.4 pg\_dbms\_stats

### 4.7.4.1 Setting Up pg\_dbms\_stats

- Set the postgresql.conf file parameter.
   Add "pg\_dbms\_stats" to the "shared\_preload\_libraries" parameter.
- 2. As superuser, run the following command:

```
$ su -
Password:*****
# cp -r /opt/fsepv<x>server64/OSS/pg_dbms_stats/* /opt/fsepv<x>server64
```

.....

3. Restart FUJITSU Enterprise Postgres.

4. Run CREATE EXTENSION for the database that will use this feature.

The target database is described as "postgres" here.

Use the psql command to connect to the "postgres" database.

postgres=# CREATE EXTENSION pg\_dbms\_stats;
CREATE EXTENSION



Refer to "Optimizer Hints" in the Application Development Guide for details.

### 4.7.4.2 Removing pg\_dbms\_stats

1. Execute DROP EXTENSION for the database that will use this feature. Use the psql command to connect to the "postgres" database.

postgres=# DROP EXTENSION pg\_dbms\_stats CASCADE; DROP EXTENSION

2. As superuser, run the following command:

\$ su Password:\*\*\*\*\*
# rm -rf /opt/fsepv<x>server64/filesCopiedDuringSetup



The files copied during setup can be checked below.

# find /opt/fsepv<x>server64/OSS/pg\_dbms\_stats

- 3. Set the postgresql.conf file parameters.

  Delete "pg\_dbms\_stats" to the shared\_preload\_libraries parameter.
- 4. Restart FUJITSU Enterprise Postgres.

# 4.7.5 pg\_repack

### 4.7.5.1 Setting Up pg\_repack

1. As superuser, run the following command:

\$ su Password:\*\*\*\*\*
# cp -r /opt/fsepv<x>server64/OSS/pg\_repack/\* /opt/fsepv<x>server64

2. Execute CREATE EXTENSION for the database that will use this feature. Use the psql command to connect to the "postgres" database.

postgres=# CREATE EXTENSION pg\_repack;
CREATE EXTENSION

### 4.7.5.2 Removing pg\_repack

1. Execute DROP EXTENSION for the database that will use this feature. Use the psql command to connect to the "postgres" database.

```
postgres=# DROP EXTENSION pg_repack CASCADE;
DROP EXTENSION
```

2. As superuser, run the following command:

```
$ su -
Password:*****
# rm -rf /opt/fsepv<x>server64/filesCopiedDuringSetup
```



The files copied during setup can be checked below.

# find /opt/fsepv<x>server64/OSS/pg\_repack

### 4.7.6 pg\_rman

### 4.7.6.1 Setting Up pg\_rman

1. As superuser, run the following command:

```
$ su -
Password:*****
# cp -r /opt/fsepv<x>server64/OSS/pg_rman/* /opt/fsepv<x>server64
```

2. Restart FUJITSU Enterprise Postgres.

# **Information**

Before initialization of the backup catalog, it is recommended to set the parameters below in postgresql.conf. Refer to the  $pg_{man}$  manual (http://ossc-db.github.io/ $pg_{man}$ /index-ja.html) for details.

- log\_directory
- archive\_mode
- archive\_command



This feature cannot be used on instances created in WebAdmin. It can only be used via server commands.

### 4.7.6.2 Removing pg\_rman

1. As superuser, run the following command:

```
$ su -
Password:*****
# rm -rf /opt/fsepv<x>server64/filesCopiedDuringSetup
```



The files copied during setup can be checked below.

# find /opt/fsepv<x>server64/OSS/pg\_rman

2. Restart FUJITSU Enterprise Postgres.

# 4.7.7 pg\_statsinfo

# 4.7.7.1 Setting Up pg\_statsinfo

- 1. Set the postgresql.conf file parameters.
  - Add "pg\_statsinfo" to the shared\_preload\_libraries parameter.

- Specify the log file name for the log\_filename parameter.
- 2. As superuser, run the following command:

```
$ su -
Password:*****
# cp -r /opt/fsepv<x>server64/OSS/pg_statsinfo/* /opt/fsepv<x>server64
```

3. Restart FUJITSU Enterprise Postgres.



Note that pg\_statsinfo forcibly overwrites the settings below.

- log\_destination

"csvlog" is added and "stderr" is deleted.

- logging\_collector

"on" is set.



This feature cannot be used on instances created in WebAdmin. It can only be used via server commands.

### 4.7.7.2 Removing pg\_statsinfo

1. As superuser, run the following command:

```
$ su -
Password:*****
# rm -rf /opt/fsepv<x>server64/filesCopiedDuringSetup
```



The files copied during setup can be checked below.

# find /opt/fsepv<x>server64/OSS/pg\_statsinfo

- 2. Set the postgresql.conf file parameters.
  - Delete "pg\_statsinfo" to the shared\_preload\_libraries parameter.

- Delete the log file name for the log\_filename parameter.
- 3. Restart FUJITSU Enterprise Postgres.

### 4.7.8 pgBadger

### 4.7.8.1 Setting Up pgBadger

1. Set the postgresql.conf file parameters.

Set the parameters so that the information required for analysis is output to the server log.

Refer to "Documentation" in the pgBadger website (https://pgbadger.darold.net/) for details.

The pgBadger material is stored under /opt/fsepv<x>server64/OSS/pgbadger.

2. Restart FUJITSU Enterprise Postgres.

### 4.7.8.2 Removing pgBadger

- Set the postgresql.conf file parameters.
   Restores information you specified during Setup.
- 2. Restart FUJITSU Enterprise Postgres.

# 4.7.9 Pgpool-II

### 4.7.9.1 Setting Up Pgpool-II

1. As superuser, run the following command:

```
$ su -
Password:*****
# cp -r /opt/fsepv<x>server64/OSS/Pgpool-II/* /opt/fsepv<x>server64
```

2. Execute CREATE EXTENSION for the database that will use this feature.

Use the psql command to connect to the "postgres" database.

```
postgres=# CREATE EXTENSION pgpool_recovery;
CREATE EXTENSION
```

- 3. Set the postgresql.conf file parameters.

  Specify the path to pg\_ctl for the pgpool.pg\_ctl parameter.
- 4. Restart FUJITSU Enterprise Postgres.



The online recovery feature of Pgpool-II cannot be used on instances created in WebAdmin. It can only be used via server commands.

# 4.7.9.2 Removing Pgpool-II

1. Execute DROP EXTENSION for the database that will use this feature. Use the psql command to connect to the "postgres" database.

```
postgres=# DROP EXTENSION pgpool_recovery CASCADE;
DROP EXTENSION
```

2. As superuser, run the following command:

```
$ su -
Password:*****
# rm -rf /opt/fsepv<x>server64/filesCopiedDuringSetup
```



The files copied during setup can be checked below.

```
# find /opt/fsepv<x>server64/OSS/Pgpool-II
```

- 3. Set the postgresql.conf file parameters. Delete the pg\_ctl path for the pgpool.pg\_ctl parameter.
- 4. Restart FUJITSU Enterprise Postgres.

# 4.8 Deleting Instances

This section explains how to delete an instance.

- 4.8.1 Using WebAdmin
- 4.8.2 Using Server Commands



- Always use WebAdmin to delete instances that were created or imported using WebAdmin. Because WebAdmin management information cannot be deleted, WebAdmin will determine that the instance is abnormal.

- If you have set automatic start and stop of the instance, execute the following commands to disable the script and cancel registration.

#### If using rc

```
{\tt chkconfig}\ {\tt nameOfShellScriptForAutomaticStartAndStop}\ {\tt off}
chkconfig --del nameOfShellScriptForAutomaticStartAndStop
rm /etc/init.d/nameOfShellScriptForAutomaticStartAndStop
```

#### Example

```
# chkconfig rc_fsepsvoi_inst1 off
# chkconfig --del rc_fsepsvoi_inst1
# rm /etc/init.d/rc_fsepsvoi_inst1
```

#### If using systemd

 ${\tt systemctl\ disable\ } name Of {\tt UnitFileThatPerformsAutomaticStartAndStop}$ 

 $\verb|rm|/usr/lib/systemd/system|/nameOfUnitFileThatPerformsAutomaticStartAndStop||$ 

#### Example

```
# systemctl disable fsepsvoi_instl.service
# rm /usr/lib/systemd/system/fsepsvoi_instl.service
```

# 4.8.1 Using WebAdmin

This section explains how to delete an instance using WebAdmin.

Use the following procedure to delete an instance.

1. Stop the instance

In the [Instances] tab, select the instance to stop and click



2. Back up files.

Before deleting the instance, back up any required files under the data storage destination, the backup data storage destination, and the transaction log storage destination.

#### 3. Delete the instance

In the [Instances] tab, select the instance to delete and click iii.



Deleting an instance deletes only the following lowest-level directories. If they are not required, delete them manually.

- Data storage destination
- Backup data storage destination
- Transaction log storage destination (if different from the data storage destination)

# 4.8.2 Using Server Commands

This section explains how to delete an instance using server commands.

Use the following procedure to delete an instance.

1. Stop the instance

Execute the stop mode of the pg\_ctl command.

An example is shown below:

Example

\$ pg\_ctl stop -D /data/inst1

2. Back up files.

Before deleting the instance, back up any required files under the data storage destination, the backup data storage destination, and the transaction log storage destination.

3. Delete the instance

Use a standard UNIX tool (the rm command) to delete the following directories:

- Data storage destination
- Backup data storage destination
- Transaction log storage destination (if a directory different from the data storage directory was specified)

# Chapter 5 Uninstallation

This chapter describes the procedure for uninstalling FUJITSU Enterprise Postgres.

### 5.1 Uninstallation in Interactive Mode

Uninstall according to the following procedure:

Note that "xSPz" in sample windows indicates the version and level of products to uninstall and "<x>" in paths indicates the product version.



- All files and directories in the installation directory are deleted during uninstallation. If user files have been placed in the installation directory, back them up before uninstallation if necessary.
- On SLES 12, before uninstallation, ensure that Oracle JRE 8 is installed, and export the JAVA\_HOME environment variable.

#export JAVA\_HOME="oracleJre8InstallDir"

To reinstall FUJITSU Enterprise Postgres after it was uninstalled, and reuse an instance that was already created so that it can be
managed from WebAdmin, back up the directory shown below in which the WebAdmin instance management information had been
defined before uninstalling FUJITSU Enterprise Postgres, and then restore the backed up directory to its original location once
FUJITSU Enterprise Postgres has been reinstalled.

Follow the procedure below to perform the backup.

- 1. Stop the WebAdmin server. Refer to "B.1.3 Stopping the Web Server Feature of WebAdmin" for details.
- 2. Back up the following directory:

webAdminInstallDir/data/fepwa



If an error occurs while the product is being uninstalled, refer to "Uninstall (middleware) Messages" in the FUJITSU Enterprise Postgres product website, and take the required action.

#### 1. Delete the operation information

If the FUJITSU Enterprise Postgres operation information has been registered in the operating system or another middleware product, for example, then it must be deleted. Cases in which deletion is required are as follows:

- For a cluster system using failover operation integrated with PRIMECLUSTER, stop the cluster applications and delete the cluster application resources and FUJITSU Enterprise Postgres resources. Refer to the Cluster Operation Guide (PRIMECLUSTER) and PRIMECLUSTER Installation and Administration Guide for details.
- If you have set automatic start and stop of the instance, execute the following commands to disable the script and cancel registration.

If using an rc script

```
\label{localizero} {\tt chkconfig\ nameOfShellScriptForAutomaticStartAndStop\ off} $$ {\tt chkconfig\ --del\ nameOfShellScriptForAutomaticStartAndStop} $$ {\tt rm\ /etc/rc.d/init.d/nameOfShellScriptForAutomaticStartAndStop} $$
```

#### Example

```
# chkconfig rc_fsepsvoi_inst1 off
# chkconfig --del rc_fsepsvoi_inst1
# rm /etc/rc.d/init.d/rc_fsepsvoi_inst1
```

#### If using systemd

systemctl disable nameOfUnitFileThatPerformsAutomaticStartAndStop

 $\verb"rm" / usr/lib/systemd/system/nameOfUnitFileThatPerformsAutomaticStartAndStop" and a substitution of the content of the con$ 

#### Example

```
# systemctl disable fsepsvoi_instl.service
# rm /usr/lib/systemd/system/fsepsvoi_instl.service
```

#### 2. Stop applications and programs

Before starting the uninstallation, stop the following:

- Applications that use the product
- Instance

#### Using WebAdmin

Using server commands

Execute the pg\_ctl command in stop mode.

```
$ /opt/fsepv<x>server64/bin/pg_ctl stop -D /database/inst1
```

- Web server feature of WebAdmin

Execute the WebAdminStop command to stop the Web server feature of WebAdmin.

#### Example

If WebAdmin is installed in /opt/fsepv<*x*>webadmin:

```
# cd /opt/fsepv<x>webadmin/sbin
# ./WebAdminStop
```

- Mirroring Controller

Execute the mc\_ctl command with the stop mode option specified and stop the Mirroring Controller.

#### Example

```
$ mc_ctl stop -M /mcdir/instl
```

- pgBadger
- Pgpool-II

#### 3. Change to the superuser

Run the following command to switch to the superuser on the system.

```
$ su -
Password:****
```

#### 4. Start the Uninstall (middleware)

Execute the following command:

```
# /opt/FJSVcir/cimanager.sh -c
```

#### 5. Select the software

Type the number for the product to be uninstalled, and press Enter.

#### 6. Start the uninstallation

To start the uninstallation, type "y" and press Enter.

To display the list of products again, type "b" and press Enter.

#### 7. Finish the uninstallation

Upon successful completion, the window below is displayed.

The installation directory may remain after uninstallation. If it is not required, delete it.

# 5.2 Uninstallation in Silent Mode

Uninstall according to the following procedure:



- All files and directories in the installation directory are deleted during uninstallation. If user files have been placed in the installation directory, back them up before uninstallation if necessary.
- On SLES 12, before uninstallation, ensure that Oracle JRE 8 is installed, and export the JAVA\_HOME environment variable.

```
#export JAVA_HOME="oracleJre8InstallDir"
```

- To reinstall FUJITSU Enterprise Postgres after it was uninstalled, and reuse an instance that was already created so that it can be managed from WebAdmin, back up the directory shown below in which the WebAdmin instance management information had been defined before uninstalling FUJITSU Enterprise Postgres, and then restore the backed up directory to its original location once FUJITSU Enterprise Postgres has been reinstalled.

Follow the procedure below to perform the backup.

- $1. \ \, Stop \ the \ Web Admin \ server. \ Refer \ to \ "B.1.3 \ Stopping \ the \ Web \ Server \ Feature \ of \ Web Admin" \ for \ details.$
- 2. Back up the following directory:

webAdminInstallDir/data/fepwa



Refer to the FUJITSU Enterprise Postgres product website for information on uninstallation in silent mode, such as the error messages.

#### 1. Delete the operation information

If the FUJITSU Enterprise Postgres operation information has been registered in the operating system or another middleware product, for example, then it must be deleted. Cases in which deletion is required are as follows:

- For a cluster system using failover operation integrated with PRIMECLUSTER, stop the cluster applications and delete the cluster applications resources and FUJITSU Enterprise Postgres resources. Refer to the Cluster Operation Guide (PRIMECLUSTER) and PRIMECLUSTER Installation and Administration Guide for details.
- If you have set automatic start and stop of the instance, execute the following commands to disable the script and cancel registration.

#### If using an rc script

```
{\tt chkconfig\ nameOfShellScriptForAutomaticStartAndStop\ off}
chkconfig --del nameOfShellScriptForAutomaticStartAndStop
rm /etc/rc.d/init.d/nameOfShellScriptForAutomaticStartAndStop
```

#### Example

```
# chkconfig rc_fsepsvoi_inst1 off
# chkconfig --del rc_fsepsvoi_inst1
# rm /etc/rc.d/init.d/rc_fsepsvoi_inst1
```

#### If using systemd

```
systemctl disable nameOfUnitFileThatPerformsAutomaticStartAndStop
```

 $\verb|rm|/usr/lib/systemd/system|/nameOfUnitFileThatPerformsAutomaticStartAndStop||$ 

#### Example

```
systemctl disable fsepsvoi_instl.service
# rm /usr/lib/systemd/system/fsepsvoi_instl.service
```

#### 2. Stop applications and programs

Before starting the uninstallation, stop the following:

- Applications that use the product
- Instance

#### Using WebAdmin



#### Using server commands

Execute the pg\_ctl command with the stop mode option specified.

```
$ /opt/fsepv<x>server64/bin/pg_ctl stop -D /database/inst1
```

- Web server feature of WebAdmin

Execute the WebAdminStop command to stop the Web server feature of WebAdmin.

#### Example

If WebAdmin is installed in /opt/fsepv<*x*>webadmin:

```
# cd /opt/fsepv<x>webadmin/sbin
 ./WebAdminStop
```

#### - Mirroring Controller

Execute the mc\_ctl command with the stop mode option specified and stop the Mirroring Controller.

#### Example

```
$ mc_ctl stop -M /mcdir/inst1
```

- pgBadger
- Pgpool-II

### 3. Change to the superuser

Run the following command to switch to the superuser on the system.

```
$ su -
Password:*****
```

#### 4. Run the uninstallation

Execute the command below.

The installation directory may remain after uninstallation. If it is not required, delete it.

#### Example

```
# /opt/fsepv<x>server64/setup/suninst.sh
```

In the example above, /opt/fsepv<x>server64/setup is the name of the installation directory in which the server product is installed.

```
# /opt/fsepv<x>webadmin/setup/suninst.sh
```

In the example above, /opt/fsepv<x>webadmin/setup is the name of the installation directory in which the WebAdmin is installed.

```
# /opt/fsepv<x>pgpool-II/setup/suninst.sh
```

In the example above, /opt/fsepv<x>pgpool-II/setup is the name of the installation directory in which the Pgpool-II is installed.

# Appendix A Recommended WebAdmin Environments

This appendix describes the recommended WebAdmin environment. The following explanation is based on the assumption that Internet Explorer 11 or later is used unless otherwise stated.

# A.1 Recommended Browser Settings

- Use a display resolution of 1280 x 768 or higher, and 256 colors or more.
- Select [View] >> [Text size] >> [Medium].
- Select [View] >> [Zoom] >> [100%].
- Click [Tools] >> [Internet options] >> [General] >> [Fonts], and then:
  - Set [Webpage font] to [Times New Roman].
  - Set [Plain text font] to [Courier New].

# A.2 How to Set Up the Pop-up Blocker

If the Pop-up Blocker is enabled, use the procedure below to configure settings to allow pop-ups from the server where FUJITSU Enterprise Postgres is installed.

- 1. Click [Tools] >> [Internet options], and then select the [Privacy] tab.

  If [Turn on Pop-up Blocker] is not selected, the Pop-up Blocker feature will not operate, and therefore steps below are not required.
- 2. Click [Settings].
- 3. In the [Pop-up Blocker Settings] window, enter in the [Address of website to allow] the URL of the server where FUJITSU Enterprise Postgres is installed, and then click [Add].
- 4. Click [Close].
- 5. In the [Internet Options] window, click [OK].

# Appendix B Setting Up and Removing WebAdmin

This appendix describes how to set up and remove WebAdmin.

Note that "<*x*>" in paths indicates the product version.

# B.1 Setting Up WebAdmin

This section explains how to set up WebAdmin.

# **B.1.1 Setting Up WebAdmin**

Follow the procedure below to set up WebAdmin.

1. Change to the superuser

Acquire superuser privileges on the system.

#### Example

```
$ su -
Password:*****
```

#### 2. Set up WebAdmin

Set up WebAdmin.

#### Example

If WebAdmin is installed in /opt/fsepv<*x*>webadmin:

```
# cd /opt/fsepv<x>webadmin/sbin
# ./WebAdminSetup
```

#### 3. Specify the port number

Specify the following port numbers to be used in WebAdmin.

Refer to the "/etc/services" file and only change to a different port number if there is overlap with a port number from another service.

Make a note of the port number for the Web server, because it will be required for starting the WebAdmin window.

Item	Value (recommended value)
Web server port number enter port number of Web Server (default: 27515):	27515
WebAdmin internal port number enter Internal port number for WebAdmin (default: 27516):	27516
WebAdmin automatic start Start WebAdmin automatically when system starting? [y,n] (default: y)	у

#### Web server port number

Specify a numeric value from 1024 to 32767 for the port number to be used for communication between the Web browser and the Web server.

The Web server port number will be registered as a port number with the following service name in the "/etc/services" file.

x is the number part when the version is x.

fsep\_xSPz\_WA\_64\_WebAdmin\_Port1

#### WebAdmin internal port number

Specify a numeric value from 1024 to 32767 for the port number to be used for communication between the Web server and the WebAdmin runtime environment.

The WebAdmin internal port number will be registered as a port number with the following service name in the /etc/services file.

x is the number part when the version is x.

fsep\_xSPz\_WA\_64\_WebAdmin\_Port2

#### WebAdmin automatic start

Select whether or not to start WebAdmin when the machine is started.



#### - Unused port numbers

Irrespective of the information specified in the "/etc/services" file, unused port numbers in the OS and other products can sometimes be automatically numbered and then used, or port numbers specified in environment files within products may also be used. Check the port numbers used by the OS and other products, and ensure that these are not duplicated.

- Access restrictions

Prevent unauthorized access and maintain security by using a firewall product, or the packet filtering feature of a router device, to restrict access to the server IP address and the various specified port numbers.

- Port access permissions

If a port is blocked (access permissions have not been granted) by a firewall, enable use of the port by granting access. Refer to the vendor document for information on how to grant port access permissions.

Consider the security risks carefully when opening ports.

- Changing port numbers

When using WebAdmin in multiserver mode, it is recommended not to change WebAdmin ports after creating instances. Otherwise, the created instances may not be accessible through WebAdmin after the port is changed.

# **B.1.2 Starting the Web Server Feature of WebAdmin**

Follow the procedure below to start the Web server feature of WebAdmin.

1. Change to the superuser

Acquire superuser privileges on the system.

#### Example

```
$ su -
Password:*****
```

2. Start the Web server feature of WebAdmin

Execute the WebAdminStart command to start the Web server feature of WebAdmin.

Example

If WebAdmin is installed in /opt/fsepv<*x*>webadmin:

```
# cd /opt/fsepv<x>webadmin/sbin
```

# ./WebAdminStart

# **B.1.3 Stopping the Web Server Feature of WebAdmin**

Follow the procedure below to stop the Web server feature of WebAdmin.

1. Change to the superuser

Acquire superuser privileges on the system.

### Example

```
$ su -
Password:*****
```

2. Stop the Web server feature of WebAdmin

Execute the WebAdminStop command to stop the Web server feature of WebAdmin.

#### Example

If WebAdmin is installed in /opt/fsepv<*x*>webadmin:

```
# cd /opt/fsepv<x>webadmin/sbin
# ./WebAdminStop
```



- For efficient operation of WebAdmin, it is recommended that the Web server feature be stopped only during a scheduled maintenance period.
- When WebAdmin is used to create and manage instances in a multiserver configuration, the Web server feature must be started and running on all servers at the same time.

# B.2 Removing WebAdmin

This section explains how to remove WebAdmin.

This removal procedure stops WebAdmin and ensures that it no longer starts automatically when the machine is restarted.

1. Change to the superuser

Acquire superuser privileges on the system.

#### Example

```
$ su -
Password:*****
```

2. Remove WebAdmin setup

Execute the WebAdminSetup command to remove WebAdmin setup.

#### Example

If WebAdmin is installed in /opt/fsepv<*x*>webadmin:

```
# cd /opt/fsepv<x>webadmin/sbin
# ./WebAdminSetup -d
```

# **B.3 Using an External Repository for WebAdmin**

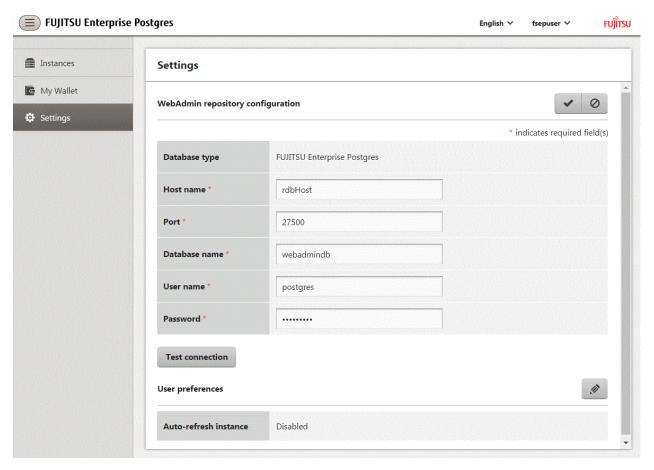
WebAdmin can be configured to use an external database, where it can store the various metadata information it uses. WebAdmin will use this database as a repository to store the information it uses to manage all the created instances. This can be a FUJITSU Enterprise Postgres database or an Open Source PostgreSQL V9.2 or later database.

Using an external database as a WebAdmin repository provides you with more flexibility in managing WebAdmin. This repository can be managed, backed up and restored as needed using pgAdmin or command line tools, allowing users to have greater flexibility and control.

Follow the procedure below to set up the repository.

1. Start WebAdmin, and log in to the database server.

2. Click the [Settings] tab, and then click in the [WebAdmin repository configuration] section.



- 3. Enter the following items:
  - [Host name]: Host name of the database server
  - [Port]: Port number of the database server
  - [Database name]: Name of the database
  - [User name]: User name to access the database
  - [Password]: Password of the database user



- Database type
  - It is recommended to use a FUJITSU Enterprise Postgres database as a repository. A compatible PostgreSQL database can also be used as an alternative.
- It is recommended to click [Test connection] to ensure that the details entered are valid and WebAdmin is able to connect to the target database.
- Host name, Database name, User name, Password should not contain hazardous characters. Refer to "Appendix C WebAdmin Disallow User Inputs Containing Hazardous Characters".
- 4. Click **v** to register the repository details.



- Once the repository is set up, it can be changed any number of times by the user logged into WebAdmin. When a repository is changed:
  - It is recommended to preload the backup into this database.
  - If the data is not preloaded, WebAdmin will create a new repository.
- The database repository can be set up even after WebAdmin was already used to create instances. In that scenario, the instances already created are retained and can continue to be operated on.
- If the instance used as a repository is stopped, WebAdmin will be unusable. For this reason, it is recommended to be familiar with starting an instance from the command line. If the instance is stopped for any reason, start it from the command line and WebAdmin will be usable again.

# **B.4** Using the WebAdmin Auto-Refresh Feature

The WebAdmin auto-refresh feature automatically refreshes the operating status of all instances in the Instance list at the specified interval. It also refreshes the details of the selected instance.

Follow the procedure below to configure the auto-refresh options.

- 1. Click the [Settings] tab, and then click in the [User preferences] section.
- 2. Enter the following items:
  - [Auto-refresh instance]: To use the auto-refresh feature, select "Enabled". The default is "Disabled".
  - [Refresh interval (seconds)]: Number of seconds between each refresh. This is a countdown timer, which is reset every time the instance status is refreshed by any operation. Specify a value from 30 to 3600 (seconds). The default is 30.

3. Click to save the auto-refresh settings.



- Auto-refresh will run only if the [Instances] page is displayed and no user-initiated operation is in progress.
- A text indicator, which is independent of auto-refresh, is displayed at the top of the Instance list. It is dynamically updated to display when the page was last refreshed.

# Appendix C WebAdmin Disallow User Inputs Containing Hazardous Characters

WebAdmin considers the following as hazardous characters, which are not allowed in user inputs

| (pipe sign)
& (ampersand sign)
; (semicolon sign)
\$ (dollar sign)
% (percent sign)
@ (at sign)
' (single apostrophe)
" (quotation mark)
\' (backslash-escaped apostrophe)
\" (backslash-escaped quotation mark)
<> (triangular parenthesis)
() (parenthesis)
+ (plus sign) [14] CR (Carriage return, ASCII 0x0d)
LF (Line feed, ASCII 0x0a)
, (comma sign)

\ (backslash)

# Appendix D Configuring Parameters

WebAdmin operates and manages databases according to the contents of the following configuration files:

- postgresql.conf

Contains various items of information that define the operating environment of FUJITSU Enterprise Postgres.

- pg\_hba.conf

Contains various items of information related to client authentication.

These configuration files are deployed to a data storage destination. Data is written to them when the instance is created by WebAdmin and when settings are changed, and data is read from them when the instance is started and when information from the [Setting] menu is displayed.

Direct editing of each configuration file is possible with a text editor.



#### See

Refer to "Server Configuration" and "Client Authentication" in "Server Administration" in the PostgreSQL Documentation for information on the parameters.

.....



WebAdmin checks for port number and backup storage path anomalies when various operations are performed. An anomaly occurs when the value of [Port number] and/or [Backup storage path] in WebAdmin is different from the value of the corresponding parameter in postgresql.conf. Refer to "Anomaly Detection and Resolution" in the Operation Guide for details.

#### postgresql.conf

Parameters that can be changed in WebAdmin

The postgresql.conf parameters that can be changed in WebAdmin are shown below:

Section	WebAdmin item	postgresql.conf file parameter			
Instance Configuration					
Character encoding	Character set	client_encoding			
	Message locale	lc_messages			
Communication	Max connections	max_connections			
SQL options	Transform NULL format	transform_null_equals			
	Date output format	DateStyle (*1)			
	Interval output format	IntervalStyle			
	Number of digits for floating values	extra_float_digits			
	Transaction isolation levels	default_transaction_isolation			
	Currency format	lc_monetary			
	Date and time format	lc_time			
	Numerical value format	lc_numeric			
Memory	Sort memory (KB)	work_mem			
	Shared buffers (KB)	shared_buffers			
Streaming replication	WAL level	wal_level			

Section	WebAdmin item	postgresql.conf file parameter
	Maximum WAL senders	max_wal_senders
	WAL keep segments	wal_keep_segments
	Hot standby	hot_standby
	Synchronous standby names	synchronous_standby_names
	WAL receiver timeout (ms)	wal_receiver_timeout
Edit instance		
	Instance name	n/a
	Instance port	port
	Backup storage path	backup_destination

<sup>\*1:</sup> If you specify "Postgres" as the output format, dates will be output in the "12-17-1997" format, not the "Wed Dec 17 1997" format used in the PostgreSQL Documentation.



- Calculate the maximum number of connections using the formula below:

```
maximumNumberOfConnections = maximumNumberOfConnectionsFromApplications + 3 (*1)
```

\*1: 3 is the default number of connections required by the system.

Calculate the maximum number of connections using the following formula when changing superuser\_reserved\_connections (connections reserved for use by the superuser) in postgresql.conf.

maximumNumberOfConnections = maximumNumberOfConnectionsFromApplications +
superuser\_reserved\_connections

- Also check if the memory used exceeds the memory installed (refer to "Parameters automatically set by WebAdmin according to the amount of memory").
- When modifying "Shared buffers" or "Max connections", edit the kernel parameter. Refer to "Appendix H Configuring Kernel Parameters", and "Managing Kernel Resources" in "Server Administration" in the PostgreSQL Documentation for details.

## Parameters set by WebAdmin

The following postgresql.conf parameters are set by WebAdmin during instance startup (they will be ignored even if specified in postgresql.conf):

Parameter	Value
listen_addresses	*
log_destination	'stderr,syslog'
logging_collector	on
log_line_prefix	'%e: %t [%p]: [%l-1] user = %u,db = %d,remote = %r app = %a '
log_filename (*1) (*2)	'logfile-%a.log'
log_file_mode	0600
log_truncate_on_rotation	on
log_rotation_age	1d

<sup>\*1:</sup> The server logs are split into files based on the day of the week, and are rotated after each week.

<sup>\*2:</sup> If the date changes while the instance is stopped, old logs are not deleted and continue to exist. Manually delete old logs that are no longer required to release disk space.

Parameters automatically set by WebAdmin according to the amount of memory

The postgresql.conf parameters automatically set according to the amount of installed memory, during the creation of instances by WebAdmin, are shown below:

Parameter	Value
shared_buffers	30% of the machine's installed memory
work_mem	30% of the machine's installed memory / max_connections / 2
effective_cache_size	75% of the machine's installed memory
maintenance_work_mem	10% of the machine's installed memory / (1 + autovacuum_max_workers) (*1)

<sup>\*1:</sup> The value will be capped at 2097151 KB.

When determining the values to be configured in the above parameters, you must take into account any anticipated increases in access volume or effects on performance during business operations, such as the number of applications and commands that will access the instance, and the content of processes. Also, note that in addition to FUJITSU Enterprise Postgres, other software may be running on the actual database server. You will need to determine the degree of priority for the database and other software, as well as the memory allocation size.

WebAdmin automatically configures complex parameter settings such as those mentioned above, based on the size of the internal memory of the machine. This enables maximum leverage of the machine memory to facilitate resistance against fluctuations during business operations.

Accordingly, the effects of the above-mentioned factors must be estimated and taken into account when determining and configuring parameter values, so that memory resources can be effectively allocated among other software or instances, and so that adverse effects can be mutually avoided. Refer to "Memory" in "Resource Consumption", and "Planner Cost Constants" in "Query Planning", under "Server Administration" in the PostgreSQL Documentation for information on parameter values and required considerations.

Parameter values can be modified using the WebAdmin [Setting] menu, or edited directly using a text editor.

If adding an instance, determine the parameter values, including for existing instances, and make changes accordingly.



#### See

Kernel parameters need to be tuned according to the parameters being changed. Refer to "Appendix H Configuring Kernel Parameters", and "Managing Kernel Resources" in "Server Administration" in the PostgreSQL Documentation for information on tuning kernel parameters.



- Do not directly edit the following postgresql.conf parameters with a text editor, otherwise WebAdmin may not work properly if you make a mistake):

- archive\_mode
- archive\_command
- wal\_level
- core\_directory
- log\_line\_prefix
- log destination
- logging\_collector
- log\_directory
- log\_file\_mode

- log\_filename
- log\_truncate\_on\_rotation
- log\_rotation\_age
- You must take care with the following parameter:
  - superuser\_reserved\_connections

Set it to a number that includes the 3 connections required in WebAdmin (the default is 3).

#### pg\_hba.conf

Refer to "Client Authentication" in "Server Administration" in the PostgreSQL Documentation for information on content that can be configured in pg\_hba.conf.



- Configure the instance administrator permissions in the "local" connection format settings. WebAdmin may not work properly if permissions are not configured.

......

- If you specify an item or value that cannot be set by WebAdmin when editing the pg\_hba.conf file with a text editor, it will not be possible to reference that line from WebAdmin.

# Appendix E Estimating Database Disk Space Requirements

This appendix describes how to estimate database disk space requirements.

# E.1 Estimating Table Size Requirements

The following tables provide the formulas for estimating table size requirements.

Table E.1 Estimation formula when the record length is 2032 bytes or less

Item	Estimation formula (bytes)	
(1) Record length	27(*1) + NULL map + OID + column data	
	NULL map: Number of columns / 8 (*2) OID: 4	
	Column data: Sum of column lengths	
	*1: Record header section	
	*2: Round the result up to the next integer.	
	<ul> <li>Because the column data is placed in boundaries of 8 bytes, you need to make an adjustment so that the sum of the record header section, NULL map and OID is a multiple of 8.</li> <li>For example, if the calculated length is 27 + 1 / 8 (rounded up) + 0 = 28 bytes, add</li> </ul>	
	4 to make the length 32 bytes.	
	- Because the data of each column is placed in boundaries of the defined data type, take the boundary of each data type into account for the length of the column data. For example, the length of the column data in the table below will not be the sum of the data types, which is 37 bytes, but will instead be 64 bytes following boundary adjustment.	
	Definition: create table tb1(c1 char(1), c2 long, c3 int, c4 box)  Estimation: CHAR type 1 byte + boundary adjustment of 7 bytes for LONG type 8  bytes + LONG type 8 bytes + INT type 4 bytes + boundary adjustment of 12 bytes  for BOX type 32 bytes + BOX type 32 bytes = 64 bytes	
	- Because each record is placed in boundaries of 8 bytes, you need to make an adjustment so that the length of the column data is a multiple of 8.	
	<ul> <li>If the calculated record length exceeds 2,032 bytes, the variable length data in the record might be compressed automatically. If so, use the estimation formulas in "Table E.2 Estimation formula when the record length exceeds 2032 bytes" to estimate the table size.</li> </ul>	
(2) Page size requirement	8192 (*1) × fillfactor (*2) - 24 (*3)	
	*1: Page length (8192)	
	*2: Value of the fillfactor specified in the table definitions (if omitted, 100%)	
	*3: Page header (24)	
	- The calculated (2) page size requirement will be rounded down to the nearest integer.	
(3) Number of records per page	(2) Page size requirement / ((1) record length + 4 (*1))	
	*1: Pointer length (4)	
	- The result will be rounded down to the nearest integer.	

Item	Estimation formula (bytes)	
(4) Number of pages required	Total number of records / (3) number of records per page	
for storing records	- The result will be rounded up to the next integer.	
(5) Amount of space	(4) Number of pages required for storing records x page length x safety factor (*1)	
	*1: Specify 2.0 or higher.	
	<ul> <li>This is the safety factor assumed if vacuuming is performed for garbage collection in tables and indexes.</li> </ul>	

Table E.2 Estimation formula when the record length exceeds 2032 bytes

Item	Estimation formula (bytes)
(5) Amount of space	Total number of records x (1) record length x safety factor (*1)
	*1: Specify 2.0 or higher.
	<ul> <li>This is the safety factor assumed if vacuuming is performed for garbage collection in tables and indexes.</li> </ul>

# **E.2 Estimating Index Size Requirements**

This section provides the formulas for estimating index size requirements.

FUJITSU Enterprise Postgres provides six index types: B-tree, Hash, GiST, GIN, SP-GiST, and VCI. If you do not specify the index type in the CREATE INDEX statement, a B-tree index is generated.

The following describes how to estimate a B-tree index. Refer to "E.7 Estimating VCI Disk Space Requirements" for information on how to estimate VCI.

A B-tree index is saved as a fixed-size page of 8 KB. The page types are meta, root, leaf, internal, deleted, and empty. Since leaf pages usually account for the highest proportion of space required, you need to calculate the requirements for these only.

Table E.3 Estimation formula when the key data length is 512 bytes or less

Item	Estimation formula (bytes)	
(1) Entry length	8 (*1) + key data length (*2)	
	*1: Entry header	
	*2: The key data length depends on its data type (refer to "E.3 Sizes of Data Types" for details).	
	Because each entry is placed in boundaries of 8 bytes, you need to make an adjustment so that the length of the key data is a multiple of 8.	
	For example, if the calculated length is 28 bytes, add 4 to make the length 32 bytes.	
	<ul> <li>If the key data length exceeds 512 bytes, key data may be automatically compressed.</li> <li>In this case, use the estimation formula given in "Table E.4 Estimation formula when the key data length exceeds 512 bytes" to estimate the key data length.</li> </ul>	
(2) Page size requirement	8192 (*1) × fillfactor (*2) - 24 (*3) - 16 (*4)	
	*1: Page length (8192)	
	*2: Value of the fillfactor specified in the index definitions (if omitted, 90%)	
	In the case of indexes of primary key constraints and unique constraints,	
	the value of the fill factor specified for each constraint in the table definitions (if omitted, 90%)	
	*3: Page header (24)	
	*4: Special data (16)	

Item	Estimation formula (bytes)	
	- The calculated (2) page size requirement will be rounded down to the nearest integer.	
(3) Number of entries per page	(2) Page size requirement / ((1) entry length + 4 (*1))	
	*1: Pointer length	
	- Result of (3) number of entries per page will be rounded down to the nearest integer.	
(4) Number of pages required	Total number of records / (3) number of entries per page	
for storing indexes	- Result of (4) number of pages required for storing indexes will be rounded up to the nearest integer.	
(5) Space requirement	(4) Number of pages required for storing indexes x 8192 (*1) / usage rate (*2)	
	*1: Page length	
	*2: Specify 0.7 or lower.	

Table E.4 Estimation formula when the key data length exceeds 512 bytes

Item	Estimation formula (bytes)	
(5) Space requirement	Total number of records x key data length x compression ratio (*1) / usage rate (*2)	
	*1: The compression ratio depends on the data value, so specify 1.	
	*2: Specify 0.7 or lower as the usage rate.	

# **E.3 Sizes of Data Types**

This section lists the sizes of the data types.

# **E.3.1 Sizes of Fixed-Length Data Types**

The following table lists the sizes of fixed-length data types.

Data type	Size (bytes)
SMALLINT (INT2)	2
INTEGER (INT4)	4
BIGINT (INT8)	8
REAL	4
DOUBLE PRECISION	8
SERIAL (SERIAL4)	4
BIGSERIAL (SERIAL8)	8
MONEY	8
FLOAT	8
FLOAT (1-24)	4
FLOAT (25-53)	8
TIMESTAMP WITHOUT TIME ZONE	8
TIMESTAMP WITH TIME ZONE	8
DATE	4
TIME WITHOUT TIME ZONE	8
TIME WITH TIME ZONE	12

Data type	Size (bytes)
INTERVAL	12
BOOLEAN	1
CIDR	IPv4: 7 IPv6: 19
INET	IPv4: 7 IPv6: 19
MACADDR	6
MACADDR8	8
POINT	16
LINE	32
LSEG	32
BOX	32
CIRCLE	24

# E.3.2 Sizes of Variable-Length Data Types

The following table lists the sizes of variable-length data types.

Data type	Size (bytes)	Remarks
path	Length of size portion + 12 + 16 x number of vertices	1) When carrying out division, round to the next integer.
polygon	Length of size portion + 36 + 16 x number of vertices	2) If the real data length is less than 127, then the length of the size portion is 1 byte, otherwise it is 4 bytes.
decimal	Length of size portion + 2 + (integer	3) The number of bytes per character depends on the
numeric	precision / 4 + decimal precision / 4) x 2	character set (refer to "E.3.4 Number of Bytes per Character" for details).
bytea	Length of size portion + real data length	Character for details).
character varying(n), varchar(n)	Length of size portion + number of characters x number of bytes per character	
character(n), char(n)	Length of size portion + n x number of bytes per character	
text	Length of size portion + number of characters x number of bytes per character	

# **E.3.3 Sizes of Array Data Types**

The following table lists the sizes of array data types.

Data type	Size (bytes)	Remarks
Array	Length of size portion + 12 + 8 x number of dimensions + data size of each item	If the real data length is less than 127, then the length of the size portion is 1 byte, otherwise it is 4 bytes.  - Example of estimation when array data is "ARRAY[[1,2,3], [1,2,3]]"  Number of dimensions: 2  INTEGER data size: 4  Total size = 1+12+8x2+6x4 = 53

# E.3.4 Number of Bytes per Character

The following table lists the number of bytes per character.

The given values relate to the common character sets EUC-JP and UTF8.

Character type	Character set	Number of bytes per character
ASCII	EUC_JP	1
Halfwidth katakana	EUC_JP	2
JIS X 0208 kanji characters	EUC_JP	2
JIS X 0212 kanji characters	EUC_JP	3
ASCII	UTF8	1
Halfwidth katakana	UTF8	3
JIS X 0208 kanji characters	UTF8	3
JIS X 0212 kanji characters	UTF8	3

# **E.4 Estimating Transaction Log Space Requirements**

This section provides the formula for estimating transaction log space requirements.

```
Transaction log space requirements = max_wal_size
```

However, if the update volume is extremely high (for example, due to a large data load and batch processing), disk writing at a checkpoint may not be able to keep up with the load, and a higher number of transaction logs than indicated here may temporarily be accumulated.

# E.5 Estimating Archive Log Space Requirements

This section explains how to estimate archive log space requirements.

The archive log is an archive of the transaction logs from the time of a previous backup to the present, so it fluctuates depending on the backup period and the content of update transactions.

The longer the backup period and the more update transactions, the greater the space required for the archive log.

Therefore, measure the actual archive log space by using a test environment to simulate backup scheduling and database update in a real operating environment.

# E.6 Estimating Backup Disk Space Requirements

This section provides the formula for estimating backup disk space requirements.

Backup disk space requirements = size of the database cluster x 2 + transaction log space requirements + archive log space requirements



If the pgx\_dmpall command performs a backup using a user exit, the backup disk size differs according to the database resources targeted for backup and the copy method.

# E.7 Estimating VCI Disk Space Requirements

This section provides the formula for estimating VCI disk space requirements.

Disk space = (number of rows in tables) x (number of bytes per row) x (compression ratio) + (WOS size)

#### Number of bytes per row

```
Number of bytes per row = (19 + (number of columns specified in CREATE INDEX) / 8 + (number of bytes per single column value)) x 1.1
```

Note: Round up the result to the nearest integer.

#### Compression ratio

Specify a value between 0 and 1. Since compression ratio depends on the data being compressed, use actual data or test data that simulates it, then compare the value with the estimation result. As a guide, the compression ratio measured with the Fujitsu sample data is shown below:

- Data with high degree of randomness (difficult to compress): Up to approximately 0.9 times.
- Data with high degree of similarity (easy to compress): Up to approximately 0.5 times.

#### WOS size

```
WOS size = (number of WOS rows) / 185 x 8096
```

One row is added to the number of WOS rows for each INSERT and DELETE, and two rows are added for UPDATE. On the other hand, the number decreases to 520,000 rows or less during conversion to ROS performed by the ROS control daemon.



VCI does not support retrieval of disk space usage using the database object size function pg\_indexes\_size. To find out the actual total VCI disk space, check the disk space of the storage directory using an OS command or other method.

# Appendix F Estimating Memory Requirements

This appendix explains how to estimate the memory.

# F.1 FUJITSU Enterprise Postgres Memory Requirements

This section describes the formulas for estimating FUJITSU Enterprise Postgres memory requirements.

Use the following formula to obtain a rough estimate of memory required for FUJITSU Enterprise Postgres:

fujitsuEnterprisePostgresRequiredMemory = sharedMemoryAmount + localMemoryAmount

#### **Shared memory amount**

Refer to "Shared Memory and Semaphores" under "Server Administration" in the PostgreSQL Documentation for information on shared memory. If you enable the Global Meta Cache feature, you must also add the value of pgx\_global\_metacache. Refer to "Parameters" in the Operation Guide for the setting values.

However, note that if instances have been created using WebAdmin, the parameters below will be configured automatically when the instances are created. Take this into account when calculating the shared memory size.

Parameter name	Set value
shared_buffers	30 percent of the internal memory of the machine.
max_connections	100
max_prepared_transactions	100

#### Local memory amount

localMemoryAmount = processStackArea

- + memoryUsedInDbSessionsThatUseTempTables
- $+ \ \textit{memoryUsedInDbSessionsThatPerformSortAndHashTableOperations}$
- + memoryUsedInMaintenanceOperations
- + baseMemoryUsedInEachProcess
- $+ \ {\it memoryUsedPreparingForDataAccess}$

#### Process stack area

processStackArea

= max\_stack\_depth x (max\_connections + autovacuum\_max\_workers + 9)

This formula evaluates to the maximum value.

Actually it is used according to the growth of the stack.

In the formula above, 9 is the number of processes that perform roles specific to servers.

Memory used in database sessions that use temporary tables

This formula evaluates to the maximum value.

Memory is gradually used as temporary buffers are used, and is released when the session ends.

Memory used in database sessions that perform sort and hash table operations

memoryUsedInDbSessionsThatPerformSortAndHashTableOperations
= work\_mem x max\_connections

This formula evaluates to the maximum value.

Memory is gradually used as operations such as sort are performed, and is released when the query ends.

#### Memory used in maintenance operations

```
memoryUsedInMaintenanceOperations
= maintenance_work_mem x (numOfSessionsPerformingMaintenance + autovacuum_max_workers)
```

Note that 'maintenance operations' are operations such as VACUUM, CREATE INDEX, and ALTER TABLE ADD FOREIGN KEY.

#### Base memory used in each process

Use the result of the following formula for memory consumed per process. This formula evaluates to the memory used when server processes are running.

In the formula above, 9 is the number of processes that perform roles specific to servers.

The amount of memory consumed per process is determined by the number of tables, indexes, and all columns of all tables that the process accesses. If your system has about 100 tables, you can estimate it to be 3 MB, but otherwise use the following estimate:

\*1) Safety Factor (1.5)

There are variable length information. This value takes that into account.

#### Memory used preparing for data access

This formula evaluates to the memory required to access the database cache in the shared memory.

In the formula above, among the processes that perform roles specific to servers, 4 is the number of processes that access the database.

# F.2 Database Multiplexing Memory Requirements

This section describes the formula for estimating database multiplexing memory requirements for the database server.

Use the following formula to obtain a rough estimate of memory required for database multiplexing:

```
Memory usage of the database multiplexing feature for the database server

= Peak memory usage of the Mirroring Controller processes
+ Peak memory usage of the Mirroring Controller commands

Peak memory usage of the Mirroring Controller processes=150 MB

Peak memory usage of the Mirroring Controller commands=50 MB x Number of commands executed simultaneously
```

# F.3 VCI Memory Requirements

This section describes the formula for estimating VCI memory requirements.

Use the following formula to obtain a rough estimate of memory requirements:

```
memUsedByVci = memForData + memForEachProcess
```

#### Memory required to store data in memory

Secure the space estimated using the formula below on the stable buffer (part of shared\_buffers).

```
memForData = (numOfRowsInTables) x (numOfBytesPerRow) + (wosSize)
```

#### Number of bytes per row

```
numOfBytesPerRow =
  (19 + (numOfColsInCreateIndexStatement) / 8 + (numOfBytesPerSingleColValue)) x 1.1
```

Note: Round up the result to the nearest integer.

#### WOS size

```
wosSize = (numOfWosRows) / 185 x 8096
```

One row is added to the number of WOS rows for each INSERT and DELETE, and two rows are added for UPDATE. On the other hand, the number decreases to 520,000 rows or less during conversion to ROS performed by the ROS control daemon.

# Memory required for each process

```
memForEachProcess =
  memUsedPerScanning +
  memUsedForVciMaintenace +
  memUsedByCreateIndexStatement
```

#### Memory used per scanning

- Parallel scan

```
memUsedPerScanning =
  vci.shared_work_mem + (numOfParallelWorkers + 1) x vci.maintenance_work_mem
```

Note: The number of parallel workers used by VCI simultaneously in the entire instance is equal to or less than vci.max\_parallel\_degree.

- Non-parallel scan

```
memUsedPerScanning = vci.max_local_ros + vci.maintenance_work_mem
```



- vci.shared\_work\_mem, and vci.max\_local\_ros are used to create local ROS. If local ROS exceeds these sizes, execute a query without using VCI according to the conventional plan.
- vci.maintenance\_work\_mem specifies the memory size to be secured dynamically. If it exceeds the specified value, a disk temporary file is used for operation.

## Memory used for VCI maintenance

```
memUsedForVciMaintenace = vci.maintenance_work_mem x vci.control_max_workers
```

#### Memory used by CREATE INDEX

```
memUsedByCreateIndexStatement = vci.maintenance_work_mem
```



vci.maintenance\_work\_mem specifies the memory to be secured dynamically. If it exceeds the specified value, a disk temporary file is used for operation.

# F.4 High-Speed Data Load Memory Requirements

This section describes the formula for estimating memory requirements for the high-speed data load feature.

Use the following formula to obtain a rough estimate of memory requirements:

```
Memory usage of high speed data load

= (Peak memory usage of pgx_loader processes + Peak memory usage of the pgx_loader commands)

x Number of commands executed simultaneously

Peak memory usage of pgx_loader processes

= Peak memory usage of the backend process (6 MB)

+ Peak memory usage of parallel workers (6 MB x number of parallel workers)

+ Peak memory usage of dynamic shared memory (80 MB x number of parallel workers)

Peak memory usage of the pgx_loader commands=9 MB
```



In addition to the size calculated using the formula above, the database cache on the shared memory estimated using the shared\_buffers parameter is consumed according to the size of the data (table and index keys) loaded using this feature. Refer to "E.1 Estimating Table Size Requirements" and "E.2 Estimating Index Size Requirements" for information on estimating an appropriate shared buffers value.

# F.5 Global Meta Cache Memory Requirements

This section describes the formula for estimating Global Meta Cache memory requirements.

The memory calculated by "Size of the GMC area" is allocated to the shared memory. The memory calculated by the per-process meta cache management information is allocated to the local memory. Refer to the graphic in "Architecture of Global Meta Cache Feature" in the "Memory usage reduction by Global Meta Cache" in the General Description for more information.

Use the following formula to obtain a rough estimate of memory requirements:

```
Amount of memory used by the Global Meta Cache feature

= Size of GMC area + Per-process meta cache management information

Size of GMC area = (All user tables x 0.4 KB

+ All user indexes x 0.3 KB

+ All user columns x 0.8 KB) x 1.5 (*1)

Per-process meta cache management information

= (All user tables + All user indexes + All user columns) x 1.0KB x max_connections x 1.5 (*1)
```

This value takes into account the case where both GMC before and after the change temporarily exist at the same time in shared memory when the table definition is changed or the row of the system catalog is changed.

<sup>\*1)</sup> Safety Factor (1.5)

# Appendix G Quantitative Limits

This appendix lists the quantitative limits of FUJITSU Enterprise Postgres.

Table G.1 Length of identifier

Table G.1 Length of identifier  Item	Limit
Database name	Up to 63 bytes (*1) (*2)
Schema name	Up to 63 bytes (*1) (*2)
Table name	Up to 63 bytes (*1) (*2)
View name	Up to 63 bytes (*1) (*2)
Index name	Up to 63 bytes (*1) (*2)
Tablespace name	Up to 63 bytes (*1) (*2)
Cursor name	Up to 63 bytes (*1) (*2)
Function name	Up to 63 bytes (*1) (*2)
Aggregate function name	Up to 63 bytes (*1) (*2)
Trigger name	Up to 63 bytes (*1) (*2)
Constraint name	Up to 63 bytes (*1) (*2)
Conversion name	Up to 63 bytes (*1) (*2)
Role name	Up to 63 bytes (*1) (*2)
Cast name	Up to 63 bytes (*1) (*2)
Collation sequence name	Up to 63 bytes (*1) (*2)
Encoding method conversion name	Up to 63 bytes (*1) (*2)
Domain name	Up to 63 bytes (*1) (*2)
Extension name	Up to 63 bytes (*1) (*2)
Operator name	Up to 63 bytes (*1) (*2)
Operator class name	Up to 63 bytes (*1) (*2)
Operator family name	Up to 63 bytes (*1) (*2)
Rewrite rule name	Up to 63 bytes (*1) (*2)
Sequence name	Up to 63 bytes (*1) (*2)
Text search settings name	Up to 63 bytes (*1) (*2)
Text search dictionary name	Up to 63 bytes (*1) (*2)
Text search parser name	Up to 63 bytes (*1) (*2)
Text search template name	Up to 63 bytes (*1) (*2)
Data type name	Up to 63 bytes (*1) (*2)
Enumerator type label	Up to 63 bytes (*1) (*2)

<sup>\*1:</sup> This is the character string byte length when converted by the server character set character code.

Table G.2 Database object

Item	Limit
Number of databases	Less than 4,294,967,296 (*1)

<sup>\*2:</sup> If an identifier that exceeds 63 bytes in length is specified, the excess characters are truncated and it is processed.

Item	Limit
Number of schemas	Less than 4,294,967,296 (*1)
Number of tables	Less than 4,294,967,296 (*1)
Number of views	Less than 4,294,967,296 (*1)
Number of indexes	Less than 4,294,967,296 (*1)
Number of tablespaces	Less than 4,294,967,296 (*1)
Number of functions	Less than 4,294,967,296 (*1)
Number of aggregate functions	Less than 4,294,967,296 (*1)
Number of triggers	Less than 4,294,967,296 (*1)
Number of constraints	Less than 4,294,967,296 (*1)
Number of conversion	Less than 4,294,967,296 (*1)
Number of roles	Less than 4,294,967,296 (*1)
Number of casts	Less than 4,294,967,296 (*1)
Number of collation sequences	Less than 4,294,967,296 (*1)
Number of encoding method conversions	Less than 4,294,967,296 (*1)
Number of domains	Less than 4,294,967,296 (*1)
Number of extensions	Less than 4,294,967,296 (*1)
Number of operators	Less than 4,294,967,296 (*1)
Number of operator classes	Less than 4,294,967,296 (*1)
Number of operator families	Less than 4,294,967,296 (*1)
Number of rewrite rules	Less than 4,294,967,296 (*1)
Number of sequences	Less than 4,294,967,296 (*1)
Number of text search settings	Less than 4,294,967,296 (*1)
Number of text search dictionaries	Less than 4,294,967,296 (*1)
Number of text search parsers	Less than 4,294,967,296 (*1)
Number of text search templates	Less than 4,294,967,296 (*1)
Number of data types	Less than 4,294,967,296 (*1)
Number of enumerator type labels	Less than 4,294,967,296 (*1)
Number of default access privileges defined in the ALTER DEFAULT PRIVILEGES statement	Less than 4,294,967,296 (*1)
Number of large objects	Less than 4,294,967,296 (*1)
Number of index access methods	Less than 4,294,967,296 (*1)

<sup>\*1:</sup> The total number of all database objects must be less than 4,294,967,296.

# Table G.3 Schema element

Item	Limit
Number of columns that can be defined in one table	From 250 to 1600 (according to the data type)
Table row length	Up to 400 gigabytes
Number of columns comprising a unique constraint	Up to 32 columns
Data length comprising a unique constraint	Less than 2,000 bytes (*1) (*2)

Item	Limit
Table size	Up to 32 terabyte
Search condition character string length in a trigger definition statement	Up to 800 megabytes (*1) (*2)
Item size	Up to 1 gigabyte

<sup>\*1:</sup> Operation might proceed correctly even if operations are performed with a quantity outside the limits.

# Table G.4 Index

Item	Limit
Number of columns comprising a key (including VCI)	Up to 32 columns
Key length (other than VCI)	Less than 2,000 bytes (*1)

<sup>\*1:</sup> This is the character string byte length when converted by the server character set character code.

Table G.5 Data types and attributes that can be handled

,	Item		Limit
Character	Data length  Specification length (n)		Data types and attributes that can be handled (*1)
			Up to 10,485,760 characters (*1)
Numeric	External decimal expres	ssion	Up to 131,072 digits before the decimal point, and up to 16,383 digits after the decimal point
	Internal binary	2 bytes	From -32,768 to 32,767
	expression	4 bytes	From -2,147,483,648 to 2,147,483,647
		8 bytes	From -9,223,372,036,854,775,808 to 9,223,372,036,854,775,807
	Internal decimal expression		Up to 13,1072 digits before the decimal point, and up to 16,383 digits after the decimal point
	Floating point expression	4 bytes	From -3.4E+38 to -7.1E-46, 0, or from 7.1E-46 to 3.4E+38
		8 bytes	From -1.7E+308 to -2.5E-324, 0, or from 2.5E-324 to 1.7E+308
bytea		Up to one gigabyte minus 53 bytes	
Large object		Up to 4 terabyte	

<sup>\*1:</sup> This is the character string byte length when converted by the server character set character code.

#### Table G.6 Function definition

Item	Limit
Number of arguments that can be specified	Up to 100
Number of variable names that can be specified in the declarations section	No limit
Number of SQL statements or control statements that can be specified in a function processing implementation	No limit

<sup>\*2:</sup> This is the character string byte length when converted by the server character set character code.

Table G.7 Data operation statement

ltem	Limit
Maximum number of connections for one process in an application (remote access)	4,000 connections
Number of expressions that can be specified in a selection list	Up to 1,664
Number of tables that can be specified in a FROM clause	No limit
Number of unique expressions that can be specified in a selection list/ DISTINCT clause/ORDER BY clause/GROUP BY clause within one SELECT statement	Up to 1,664
Number of expressions that can be specified in a GROUP BY clause	No limit
Number of expressions that can be specified in an ORDER BY clause	No limit
Number of SELECT statements that can be specified in a UNION clause/INTERSECT clause/EXCEPT clause	Up to 4,000 (*1)
Number of nestings in joined tables that can be specified in one view	Up to 4,000 (*1)
Number of functions or operator expressions that can be specified in one expression	Up to 4,000 (*1)
Number of expressions that can be specified in one row constructor	Up to 1,664
Number of expressions that can be specified in an UPDATE statement SET clause	Up to 1,664
Number of expressions that can be specified in one row of a VALUES list	Up to 1,664
Number of expressions that can be specified in a RETURNING clause	Up to 1,664
Total expression length that can be specified in the argument list of one function specification	Up to 800 megabytes (*2)
Number of cursors that can be processed simultaneously by one session	No limit
Character string length of one SQL statement	Up to 800 megabytes (*1) (*3)
Number of input parameter specifications that can be specified in one dynamic SQL statement	No limit
Number of tokens that can be specified in one SQL statement	Up to 10,000
Number of values that can be specified as a list in a WHERE clause IN syntax	No limit
Number of expressions that can be specified in a USING clause	No limit
Number of JOINs that can be specified in a joined table	Up to 4,000 (*1)
Number of expressions that can be specified in COALESCE	No limit
Number of WHEN clauses that can be specified for CASE in a simple format or a searched format	No limit
Data size per record that can be updated or inserted by one SQL statement	Up to one gigabyte minus 53 bytes
Number of objects that can share a lock simultaneously	Up to 256,000 (*1)

<sup>\*1:</sup> Operation might proceed correctly even if operations are performed with a quantity outside the limits.

<sup>\*2:</sup> The total number of all database objects must be less than 4,294,967,296.

<sup>\*3:</sup> This is the character string byte length when converted by the server character set character code.

# Table G.8 Data size

Item	Limit
Data size per record for input data files (COPY statement, psql command \copy meta command)	Up to 800 megabytes (*1)
Data size per record for output data files (COPY statement, psql command \copy meta command)	Up to 800 megabytes (*1)

<sup>\*1:</sup> Operation might proceed correctly even if operations are performed with a quantity outside the limits.

# Appendix H Configuring Kernel Parameters

Use the "System V IPC Parameters" table in "Managing Kernel Resources" in the PostgreSQL Documentation for the relationship between configuration parameters and kernel parameters, as well as calculation formulas.

Refer to the "Managing Kernel Resources" in the PostgreSQL Documentation to calculate shared memory usage.

For multiple instances, the kernel parameters should be evaluated for all instances. For example, in the case of the maximum number of shared memory segments for the entire system (SHMMNI), the total number of segments obtained by all instances should be added to the kernel parameters. In the case of the maximum number of semaphores for each process (SEMMSL), the largest of all sizes obtained by all instances should be compared to the current value prior to configuring the settings.



If there is insufficient shared memory due to miscalculation of SHMMAX, a message will be output indicating that the shmget system call failed at "errno=22 (EINVAL)". Review the calculation, and reconfigure.

The relationship between System V IPC parameters and kernel parameters in various operating systems is shown below.

System	V IPC parameter	Kernel parameter action
SHMMAX	kernel.shmmax	If <i>currentValue</i> < <i>calculatedValue</i> , configure the calculated value
SHMMIN	No compatible parameter	
SHMALL	kernel.shmall	Specify currentValue + calculatedValue
SHMSEG	No compatible parameter	
SHMMNI	kernel.shmmni	Specify currentValue + calculatedValue
SEMMNI	Fourth parameter of kernel.sem	Specify currentValue + calculatedValue
SEMMNS	Second parameter of kernel.sem	Specify currentValue + calculatedValue
SEMMSL	First parameter of kernel.sem	If <i>currentValue</i> < <i>calculatedValue</i> , configure the calculated value
SEMMAP	No compatible parameter	
SEMVMX	No compatible parameter	

Remark 1: kernel.shmall specifies the number of pages.

Remark 2: Specify all four parameters for kernel.sem. At this time, the value specified in the third parameter should be the same value as before configuration.

# Appendix I Determining the Preferred WebAdmin Configuration

This appendix describes the two different configurations in which WebAdmin can be used and how to select the most suitable configuration.

# I.1 WebAdmin Configurations

WebAdmin can be installed in two configurations:

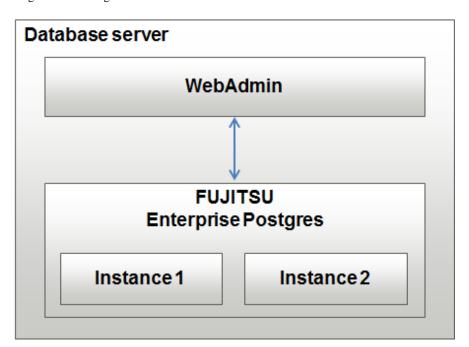
- Single-server
- Multiserver

WebAdmin does not support encrypted communication between browser and server or between servers. Therefore, when using WebAdmin in either configuration, build the communication path with the browser or each server on a network that cannot be accessed externally.

# I.1.1 Single-Server Configuration

A single-server configuration enables you to create and operate instances on a single server. In this configuration, WebAdmin must be installed on the same database server as the FUJITSU Enterprise Postgres Server component.

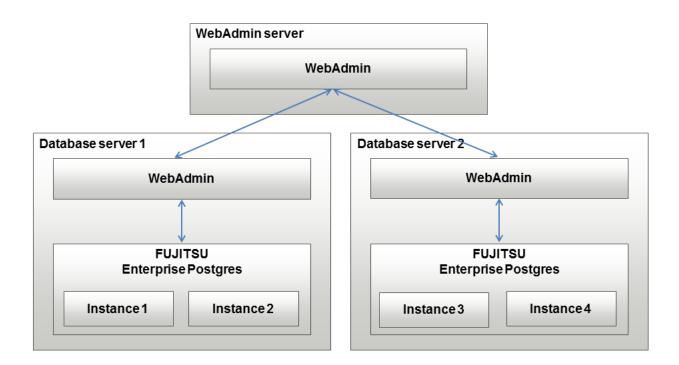
Single-server configuration



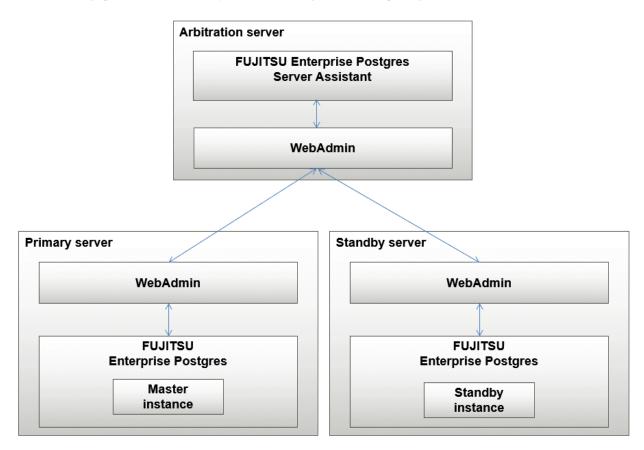
# I.1.2 Multiserver Configuration

A multiserver configuration enables you to create and operate instances stored on multiple database servers. As shown in the figure below, WebAdmin can be installed on a dedicated WebAdmin server and used to collectively manage the instances stored on the database servers.

Multiserver configuration



Also, when setting up the arbitration server by WebAdmin during database multiplexing mode, install WebAdmin on the arbitration server.



# I.2 Installing WebAdmin in a Single-Server Configuration

To install WebAdmin in a single-server configuration, the FUJITSU Enterprise Postgres Server component and WebAdmin must be installed on the same machine.

Select the following items when installing FUJITSU Enterprise Postgres in a single-server configuration:

- FUJITSU Enterprise Postgres Advanced Edition or FUJITSU Enterprise Postgres Standard Edition
- WebAdmin

# I.3 Installing WebAdmin in a Multiserver Configuration

In a multiserver configuration, install WebAdmin on one server, and both WebAdmin and the FUJITSU Enterprise Postgres Server component on any number of database servers.

Select the following items when installing FUJITSU Enterprise Postgres in a multiserver configuration:

- WebAdmin server:
  - WebAdmin
- Database server:
  - FUJITSU Enterprise Postgres Advanced Edition or FUJITSU Enterprise Postgres Standard Edition
  - WebAdmin

Also, when setting up the arbitration server by WebAdmin during database multiplexing mode, select the following when installing FUJITSU Enterprise Postgres.

......

- Arbitration server
  - FUJITSU Enterprise Postgres Server Assistant
  - WebAdmin



Refer to the Installation and Setup Guide for Server Assistant for details on how to install the Server Assistant.

# Appendix J System Configuration when using Pgpool-II

Describes the system configuration when using Pgpool-II.

The system configuration when using Pgpool-II is as follows:

#### Place on database server

System configuration to coexist the database server with Pgpool-II.

#### Place on application server

System configuration to coexist the application server with Pgpool-II.

#### Place on dedicated server

System configuration in which Pgpool-II resides on a dedicated server (Pgpool-II Server) that is separate from the database and application servers.

Select the system configuration that best meets your operational requirements.

# J.1 Pgpool-II Configuration

In this example, Pgpool-II is deployed on a different Pgpool-II server than the database and application servers.

There are three configurations of Pgpool-II:

- Single-machine configuration
- Two-machine configuration
- Three-machine configuration

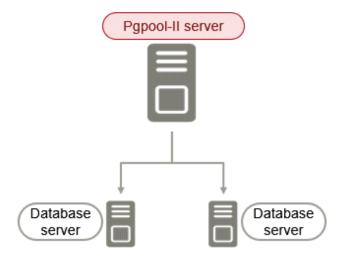
Although the Pgpool-II server can be operated on a single machine, to ensure business continuity, it is recommended to operate the Pgpool-II server using a three-machine configuration in FUJITSU Enterprise Postgres.

If employing a configuration of three or more machines, use an odd number of machines in the configuration.

# J.1.1 Single-Machine Configuration

This is the basic configuration when running Pgpool-II.

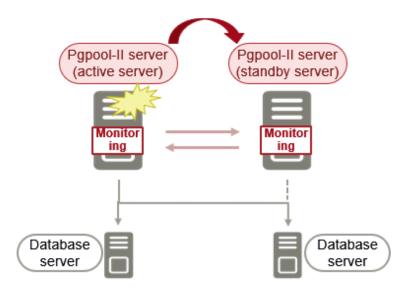
Although the database server has redundancy, if an error occurs on the Pgpool-II server that accesses the database server, the job will stop.



# J.1.2 Two-Machine Configuration

When an error occurs on the active server, the Pgpool-II monitoring feature that mutually monitors the status of the Pgpool-II servers enables jobs to continue uninterrupted by switching to the standby server.

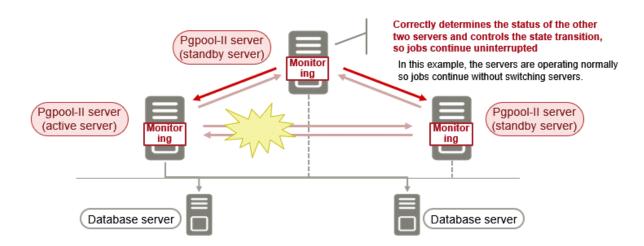
If the network between Pgpool-II servers is disconnected, even if the Pgpool-II servers are running correctly, which may lead to stoppage of jobs.



# J.1.3 Three-Machine Configuration

The Pgpool-II monitoring feature enables a Pgpool-II server to monitor the other two Pgpool-II servers.

Even if any of the networks monitoring the Pgpool-II servers are disconnected, the status of servers on a network that is operating normally can be checked correctly, enabling accurate continuation of jobs.



# J.2 Installing Pgpool-II

Pgpool-II is bundled with the server program and the client program. To use Pgpool-II, use the server program or the client program to install and set up Pgpool-II.

Depending on where Pgpool-II is installed, select the appropriate DVD for deployment:

# **Installing on Database Server (coexist)**

Install the Pgpool-II program along with the server program from the server program DVD.

# **Installing on Application Server (coexist)**

Install the Pgpool-II program along with the client program from the client program DVD.

# Installing on Dedicated server different from the above (Pgpool-II server)

Install only the Pgpool-II program from the client program DVD.

# Appendix K Supported contrib Modules and Extensions Provided by External Projects

FUJITSU Enterprise Postgres supports PostgreSQL contrib modules, and extensions provided by external projects.

Refer to the following for details on the supported contrib modules:

- "Additional Supplied Modules" in the PostgreSQL Documentation
- "Additional Supplied Programs" in the PostgreSQL Documentation



You can also check the list of available extensions using the pg\_available\_extensions view.

Refer to "OSS Supported by FUJITSU Enterprise Postgres" in the General Description for information on supported extensions provided by external projects.

# Appendix L Procedure when Modifying the Oracle JRE Installation

This appendix describes the procedures to follow when modifying the Oracle JRE installation.

The Oracle JRE, of which the installation destination is specified in the JAVA\_HOME environment variable when installing FUJITSU Enterprise Postgres, is used by features such as WebAdmin and database multiplexing.

Therefore, when updating or reinstalling Oracle JRE after installing FUJITSU Enterprise Postgres, the procedures below must be performed.

# L.1 When Using WebAdmin

WebAdmin must be set up again.

Follow the procedure below to modify the Oracle JRE installation:

1. Stop the Web server feature of WebAdmin

Refer to "B.1.3 Stopping the Web Server Feature of WebAdmin" for details.

2. Remove WebAdmin

Refer to "B.2 Removing WebAdmin" for details.

- 3. Modify the Oracle JRE installation
- 4. Set the JAVA\_HOME environment variable

Set the JAVA\_HOME environment variable to the installation destination of Oracle JRE 8.

Example

```
# export JAVA_HOME="oracleJre8InstallDir"
```

5. Set up WebAdmin

Refer to "B.1.1 Setting Up WebAdmin" for details.

6. Start the Web server feature of WebAdmin

Refer to "B.1.2 Starting the Web Server Feature of WebAdmin" for details.

# L.2 When Performing Database Multiplexing

Mirroring Controller must be restarted.

Follow the procedure below to modify the Oracle JRE installation:

1. Stop Mirroring Controller

Refer to the Cluster Operation Guide (Database Multiplexing) for details.

- 2. Modify the Oracle JRE installation
- 3. Change the installation environment to be used by Mirroring Controller

Set the JAVA\_HOME environment variable to the installation destination of Oracle JRE 8, and use the mc\_update\_jre\_env command to change the installation environment to be used by Mirroring Controller.

This procedure must be executed by the superuser.

#### Example

/opt/fsepv<x>server64/bin is the installation directory where the server product is installed.

```
$ su -
Password:*****
```

# export JAVA\_HOME="oracleJre8InstallDir"
# /opt/fsepv<x>server64/bin/mc\_update\_jre\_env

#### 4. Start Mirroring Controller

Refer to the Cluster Operation Guide (Database Multiplexing) for details.



- If database multiplexing is performed using WebAdmin, perform the procedure described in step 3 "Change the installation environment to be used by Mirroring Controller" after performing step 4 "Set the JAVA\_HOME environment variable" in "L.1 When Using WebAdmin".

# Index

Required Patches	7
[S]	
Settings related to connection	31
Starting the Web Server Feature of WebAdmin	
Startup URL for WebAdmin	
Stopping the Web Server Feature of WebAdmin	
Supported contrib Modules and Extensions Provided by Ex	
Projects	
110,000	00
[T]	
TCP/IP Protocol	8
[U]	
Uninstallation	
Uninstallation in Interactive Mode	
Uninstallation in Silent Mode	
Using the initdb Command	
Using WebAdmin	20
DAZ	
[W]	
WebAdmin automatic start	
Web server port number	
When an Instance was Created with the initdb Command	
When an Instance was Created with WebAdmin	
When Performing Database Multiplexing	
When Using WehAdmin	84



# FUJITSU Enterprise Postgres 12



# Installation and Setup Guide for Server (Windows)

# **Preface**

## Purpose of this document

The FUJITSU Enterprise Postgres database system extends the PostgreSQL features and runs on the Windows platform.

This document describes how to install and set up "FUJITSU Enterprise Postgres".

#### Intended readers

This document is intended for those who install and operate FUJITSU Enterprise Postgres.

Readers of this document are assumed to have general knowledge of:

- PostgreSQL
- SQL
- Windows

#### Structure of this document

This document is structured as follows:

#### Chapter 1 Overview of Installation

Describes the installation types and procedures

#### Chapter 2 Operating Environment

Describes the operating environment required to use FUJITSU Enterprise Postgres

#### Chapter 3 Installation

Describes how to perform a new installation of FUJITSU Enterprise Postgres

#### Chapter 4 Setup

Describes the setup to be performed after installation

#### Chapter 5 Uninstallation

Describes how to uninstall FUJITSU Enterprise Postgres

#### Appendix A Recommended WebAdmin Environments

Describes the recommended WebAdmin environment

### Appendix B Setting Up and Removing WebAdmin

Describes how to set up and remove WebAdmin

#### Appendix C WebAdmin Disallow User Inputs Containing Hazardous Characters

Describes characters that are not allowed in WebAdmin.

## Appendix D Configuring Parameters

Describes FUJITSU Enterprise Postgres parameters

# Appendix E Estimating Database Disk Space Requirements

Describes how to estimate database disk space requirements

#### Appendix F Estimating Memory Requirements

Describes the formulas for estimating memory requirements

#### Appendix G Quantitative Limits

Describes the quantity range

## Appendix H Determining the Preferred WebAdmin Configuration

Describes the two different configurations in which WebAdmin can be used and how to select the most suitable configuration

# Appendix I Supported contrib Modules and Extensions Provided by External Projects

Lists the PostgreSQL contrib modules and the extensions provided by external projects supported by FUJITSU Enterprise Postgres.

# **Export restrictions**

Exportation/release of this document may require necessary procedures in accordance with the regulations of your resident country and/or US export control laws.

#### Issue date and version

Edition 1.0: July 2020

# Copyright

Copyright 2015-2020 FUJITSU LIMITED

# **Contents**

Chapter 1 Overview of Installation	······
1.1 Features that can be Installed.	
1.2 Installation Types	
1.2.1 New Installation.	
1.2.2 Reinstallation.	
1.2.3 Multi-Version Installation.	
1.3 Installation Procedure	
1.3.1 Installation in Interactive Mode	2
1.3.2 Installation in Silent Mode	2
1.4 Uninstallation	2
Objection Comparation Equipment	,
Chapter 2 Operating Environment.	
2.1 Required Operating System.	
2.2 Related Software.	
2.3 Excluded Software.	
2.4 Required Patches.	
2.5 Hardware Environment.	
2.6 Disk Space Required for Installation	
2.7 Supported System Environment.	
2.7.1 TCP/IP Protocol.	
2.7.2 File System	
2.9 Notes on Using Streaming Replication	
Chapter 3 Installation.	<del>.</del>
3.1 Pre-installation Tasks.	
3.2 Installation in Interactive Mode	
3.3 Installation in Silent Mode	
Chapter 4 Setup	
4.1 Operating Method Types and Selection	13
4.1 Operating Method Types and Selection	13
4.1 Operating Method Types and Selection 4.2 Preparations for Setup	12
4.1 Operating Method Types and Selection 4.2 Preparations for Setup 4.2.1 Creating an Instance Administrator 4.2.1.1 Security policy settings	
4.1 Operating Method Types and Selection 4.2 Preparations for Setup 4.2.1 Creating an Instance Administrator 4.2.1.1 Security policy settings 4.2.2 Preparing Directories for Resource Deployment	12 14 15
4.1 Operating Method Types and Selection 4.2 Preparations for Setup 4.2.1 Creating an Instance Administrator 4.2.1.1 Security policy settings 4.2.2 Preparing Directories for Resource Deployment 4.2.3 Estimating Resources	
4.1 Operating Method Types and Selection 4.2 Preparations for Setup 4.2.1 Creating an Instance Administrator 4.2.1.1 Security policy settings 4.2.2 Preparing Directories for Resource Deployment 4.2.3 Estimating Resources 4.2.4 Windows Firewall Settings	
4.1 Operating Method Types and Selection 4.2 Preparations for Setup 4.2.1 Creating an Instance Administrator 4.2.1.1 Security policy settings 4.2.2 Preparing Directories for Resource Deployment 4.2.3 Estimating Resources	
4.1 Operating Method Types and Selection 4.2 Preparations for Setup 4.2.1 Creating an Instance Administrator 4.2.1.1 Security policy settings 4.2.2 Preparing Directories for Resource Deployment 4.2.3 Estimating Resources 4.2.4 Windows Firewall Settings 4.2.5 Preparing for Output to the Event Log 4.3 Creating Instances	
4.1 Operating Method Types and Selection.  4.2 Preparations for Setup	
4.1 Operating Method Types and Selection.  4.2 Preparations for Setup	
4.1 Operating Method Types and Selection.  4.2 Preparations for Setup.  4.2.1 Creating an Instance Administrator.  4.2.1.1 Security policy settings.  4.2.2 Preparing Directories for Resource Deployment.  4.2.3 Estimating Resources.  4.2.4 Windows Firewall Settings.  4.2.5 Preparing for Output to the Event Log.  4.3 Creating Instances.  4.3.1 Using WebAdmin.  4.3.1.1 Logging in to WebAdmin.  4.3.1.2 Creating an Instance.	
4.1 Operating Method Types and Selection 4.2 Preparations for Setup 4.2.1 Creating an Instance Administrator 4.2.1.1 Security policy settings 4.2.2 Preparing Directories for Resource Deployment 4.2.3 Estimating Resources 4.2.4 Windows Firewall Settings 4.2.5 Preparing for Output to the Event Log 4.3 Creating Instances 4.3.1 Using WebAdmin 4.3.1.1 Logging in to WebAdmin 4.3.1.2 Creating an Instance 4.3.1.3 Changing Instance Settings	
4.1 Operating Method Types and Selection.  4.2 Preparations for Setup.  4.2.1 Creating an Instance Administrator.  4.2.1.1 Security policy settings.  4.2.2 Preparing Directories for Resource Deployment.  4.2.3 Estimating Resources.  4.2.4 Windows Firewall Settings.  4.2.5 Preparing for Output to the Event Log.  4.3 Creating Instances.  4.3.1 Using WebAdmin.  4.3.1.1 Logging in to WebAdmin.  4.3.1.2 Creating an Instance.  4.3.1.3 Changing Instance Settings.  4.3.1.4 Importing Instances.	
4.1 Operating Method Types and Selection.  4.2 Preparations for Setup.  4.2.1 Creating an Instance Administrator.  4.2.1.1 Security policy settings.  4.2.2 Preparing Directories for Resource Deployment.  4.2.3 Estimating Resources.  4.2.4 Windows Firewall Settings.  4.2.5 Preparing for Output to the Event Log.  4.3 Creating Instances.  4.3.1 Using WebAdmin.  4.3.1.1 Logging in to WebAdmin.  4.3.1.2 Creating an Instance.  4.3.1.3 Changing Instance Settings.  4.3.1.4 Importing Instances.  4.3.1.4 Importing Instances.	11 14 14 15 15 15 15 15 15 15 15 15 15 15 15 15
4.1 Operating Method Types and Selection.  4.2 Preparations for Setup  4.2.1 Creating an Instance Administrator  4.2.1.1 Security policy settings  4.2.2 Preparing Directories for Resource Deployment.  4.2.3 Estimating Resources  4.2.4 Windows Firewall Settings  4.2.5 Preparing for Output to the Event Log  4.3 Creating Instances  4.3.1 Using WebAdmin  4.3.1.1 Logging in to WebAdmin  4.3.1.2 Creating an Instance  4.3.1.3 Changing Instance Settings  4.3.1.4 Importing Instances  4.3.2 Using the initdb Command  4.3.2.1 Creating an Instance	11 14 14 15 16 16 17 17 18 17 18 18 18 18 18 18 18 18 18 18 18 18 18
4.1 Operating Method Types and Selection.  4.2 Preparations for Setup.  4.2.1 Creating an Instance Administrator.  4.2.1.1 Security policy settings.  4.2.2 Preparing Directories for Resource Deployment.  4.2.3 Estimating Resources.  4.2.4 Windows Firewall Settings.  4.2.5 Preparing for Output to the Event Log.  4.3 Creating Instances.  4.3.1 Using WebAdmin.  4.3.1.1 Logging in to WebAdmin.  4.3.1.2 Creating an Instance.  4.3.1.3 Changing Instance Settings.  4.3.1 Importing Instances.  4.3.1 Using the initdb Command.  4.3.2 Using the initdb Command.  4.3.2.1 Creating an Instance.	113
4.1 Operating Method Types and Selection.  4.2 Preparations for Setup  4.2.1 Creating an Instance Administrator.  4.2.1.1 Security policy settings  4.2.2 Preparing Directories for Resource Deployment.  4.2.3 Estimating Resources.  4.2.4 Windows Firewall Settings.  4.2.5 Preparing for Output to the Event Log  4.3 Creating Instances  4.3.1 Using WebAdmin.  4.3.1.1 Logging in to WebAdmin.  4.3.1.2 Creating an Instance.  4.3.1.3 Changing Instance Settings.  4.3.1 Importing Instances.  4.3.2 Using the initdb Command.  4.3.2 Using the initdb Command.  4.3.1 Creating an Instance  4.4 Configuring Remote Connections.  4.4.1 When an Instance was Created with WebAdmin.	
4.1 Operating Method Types and Selection.  4.2 Preparations for Setup.  4.2.1 Creating an Instance Administrator.  4.2.1.1 Security policy settings  4.2.2 Preparing Directories for Resource Deployment.  4.2.3 Estimating Resources.  4.2.4 Windows Firewall Settings  4.2.5 Preparing for Output to the Event Log.  4.3 Creating Instances.  4.3.1 Using WebAdmin.  4.3.1.1 Logging in to WebAdmin.  4.3.1.2 Creating an Instance.  4.3.1.3 Changing Instance Settings.  4.3.1 Importing Instance Settings.  4.3.1 Using WebAdmin Instance.  4.3.1.2 Creating an Instance.  4.3.1.3 Changing Instance Settings.  4.3.1 Using Remote Connections.  4.3.2 Using the initdb Command.  4.3.3 Using Remote Connections.  4.4.1 When an Instance was Created with WebAdmin.  4.4.2 When an Instance was Created with the initdb Command.	
4.1 Operating Method Types and Selection 4.2 Preparations for Setup 4.2.1 Creating an Instance Administrator 4.2.1.1 Security policy settings 4.2.2 Preparing Directories for Resource Deployment 4.2.3 Estimating Resources 4.2.4 Windows Firewall Settings 4.2.5 Preparing for Output to the Event Log 4.3 Creating Instances 4.3.1 Using WebAdmin 4.3.1.1 Logging in to WebAdmin 4.3.1.2 Creating an Instance 4.3.1.3 Changing Instance Settings 4.3.1 Importing Instances 4.3.1 Using the initdb Command 4.3.2 Using the initdb Command 4.3.2 Using the initdb Command 4.3.4 Configuring Remote Connections 4.4 Configuring Remote Connections 4.4.1 When an Instance was Created with WebAdmin 4.4.2 When an Instance was Created with the initdb Command 4.4.3 Windows Firewall Settings	11 14 14 15 16 17 18 18 18 18 18 18 18 18 18 18 18 18 18
4.1 Operating Method Types and Selection 4.2 Preparations for Setup 4.2.1 Creating an Instance Administrator 4.2.1.1 Security policy settings 4.2.2 Preparing Directories for Resource Deployment 4.2.3 Estimating Resources 4.2.4 Windows Firewall Settings 4.2.5 Preparing for Output to the Event Log 4.3 Creating Instances 4.3.1 Using WebAdmin 4.3.1.1 Logging in to WebAdmin 4.3.1.2 Creating an Instance 4.3.1.3 Changing Instance Settings 4.3.1.4 Importing Instances 4.3.2 Using the initdb Command 4.3.2 Using the initdb Command 4.3.2 Using Remote Connections 4.4 Configuring Remote Connections 4.4.1 When an Instance was Created with WebAdmin 4.4.2 When an Instance was Created with the initdb Command 4.4.3 Windows Firewall Settings 4.5 Other Settings	11 14 14 15 16 16 16 16 16 16 16 16 16 16 16 16 16
4.1 Operating Method Types and Selection  4.2 Preparations for Setup  4.2.1 Creating an Instance Administrator  4.2.1.1 Security policy settings  4.2.2 Preparing Directories for Resource Deployment  4.2.3 Estimating Resources  4.2.4 Windows Firewall Settings  4.2.5 Preparing for Output to the Event Log  4.3 Creating Instances  4.3.1 Using WebAdmin  4.3.1.1 Logging in to WebAdmin  4.3.1.2 Creating an Instance  4.3.1.3 Changing Instance Settings  4.3.1.4 Importing Instances  4.3.2 Using the initdb Command  4.3.2.1 Creating an Instance  4.4 Configuring Remote Connections  4.4.1 When an Instance was Created with WebAdmin  4.4.2 When an Instance was Created with the initdb Command  4.4.3 Windows Firewall Settings  4.5 Other Settings  4.5.1 Error Log Settings	113
4.1 Operating Method Types and Selection  4.2 Preparations for Setup  4.2.1 Creating an Instance Administrator.  4.2.1.1 Security policy settings  4.2.2 Preparing Directories for Resource Deployment  4.2.3 Estimating Resources  4.2.4 Windows Firewall Settings  4.2.5 Preparing for Output to the Event Log  4.3 Creating Instances  4.3.1 Using WebAdmin  4.3.1.1 Logging in to WebAdmin  4.3.1.2 Creating an Instance  4.3.1.3 Changing Instance Settings  4.3.1 Using the initidb Command  4.3.2 Using the initidb Command  4.3.2.1 Creating an Instance  4.4 Configuring Remote Connections  4.4.1 When an Instance was Created with WebAdmin  4.4.2 When an Instance was Created with the initidb Command  4.4.3 Windows Firewall Settings  4.5 Other Settings  4.5.1 Error Log Settings  4.5.2 Configuring Automatic Start and Stop of an Instance	113
4.1 Operating Method Types and Selection  4.2 Preparations for Setup  4.2.1 Creating an Instance Administrator 4.2.1.1 Security policy settings 4.2.2 Preparing Directories for Resource Deployment 4.2.3 Estimating Resources 4.2.4 Windows Firewall Settings 4.2.5 Preparing for Output to the Event Log 4.3 Creating Instances 4.3.1 Using WebAdmin 4.3.1.1 Logging in to WebAdmin 4.3.1.2 Creating an Instance 4.3.1.3 Changing Instance Settings 4.3.1 Using the initdb Command 4.3.2 Using the initdb Command 4.3.2.1 Creating an Instance 4.4 Configuring Remote Connections 4.4.1 When an Instance was Created with WebAdmin 4.4.2 When an Instance was Created with the initdb Command 4.4.3 Windows Firewall Settings 4.5 Other Settings 4.5.1 Error Log Settings 4.5.2 Configuring Automatic Start and Stop of an Instance 4.5.3 Settings when Using the Features Compatible with Oracle Databases.	11 14 14 15 15 16 16 16 16 16 16 16 16 16 16 16 16 16
4.1 Operating Method Types and Selection  4.2 Preparations for Setup  4.2.1 Creating an Instance Administrator.  4.2.1.1 Security policy settings  4.2.2 Preparing Directories for Resource Deployment  4.2.3 Estimating Resources  4.2.4 Windows Firewall Settings  4.2.5 Preparing for Output to the Event Log  4.3 Creating Instances  4.3.1 Using WebAdmin  4.3.1.1 Logging in to WebAdmin  4.3.1.2 Creating an Instance  4.3.1.3 Changing Instance Settings  4.3.1 Using the initidb Command  4.3.2 Using the initidb Command  4.3.2.1 Creating an Instance  4.4 Configuring Remote Connections  4.4.1 When an Instance was Created with WebAdmin  4.4.2 When an Instance was Created with the initidb Command  4.4.3 Windows Firewall Settings  4.5 Other Settings  4.5.1 Error Log Settings  4.5.2 Configuring Automatic Start and Stop of an Instance	11 14 14 14 15 15 15 15 15 15 15 15 15 15 15 15 15

4.7.1 oracle_fdw	
4.7.1.1 Setting Up oracle_fdw	37
4.7.1.2 Removing oracle_fdw	37
4.7.2 pg_hint_plan	
4.7.2.1 Setting Up pg_hint_plan	
4.7.2.2 Removing pg_hint_plan	
4.7.3 pg_dbms_stats	38
4.7.3.1 Setting Up pg_dbms_stats	
4.7.3.2 Removing pg_dbms_stats	39
4.8 Deleting Instances	
4.8.1 Using WebAdmin	
4.8.2 Using Server Commands	40
Chapter 5 Uninstallation	41
5.1 Uninstallation in Interactive Mode	41
5.2 Uninstallation in Silent Mode	43
Appendix A Recommended WebAdmin Environments	47
A.1 Recommended Browser Settings	
A.2 How to Set Up the Pop-up Blocker.	
Appendix B Setting Up and Removing WebAdmin	
B.1 Setting Up WebAdmin	
B.1.1 Setting Up WebAdmin	
B.1.2 Starting the Web Server Feature of WebAdmin	
B.1.3 Stopping the Web Server Feature of WebAdmin	
B.2 Removing WebAdmin	
B.3 Using an External Repository for WebAdmin	
B.4 Using the WebAdmin Auto-Refresh Feature	52
Appendix C WebAdmin Disallow User Inputs Containing Hazardous Characters	53
Appendix D Configuring Parameters	5.4
Appendix E Estimating Database Disk Space Requirements	
E.1 Estimating Table Size Requirements	58
E.2 Estimating Index Size Requirements.	
E.3 Sizes of Data Types	60
E.3.1 Sizes of Fixed-Length Data Types	60
E.3.2 Sizes of Variable-Length Data Types	61
E.3.3 Sizes of Array Data Types	61
E.3.4 Number of Bytes per Character	
E.4 Estimating Transaction Log Space Requirements	62
E.5 Estimating Archive Log Space Requirements	
E.6 Estimating Backup Disk Space Requirements	62
E.7 Estimating VCI Disk Space Requirements	62
Appendix F Estimating Memory Requirements	64
F.1 FUJITSU Enterprise Postgres Memory Requirements	
F.2 Database Multiplexing Memory Requirements	
F.3 VCI Memory Requirements	
F.4 High-Speed Data Load Memory Requirements	
F.5 Global Meta Cache Memory Requirements	
Appendix G Quantitative Limits	68
Appendix H Determining the Preferred WebAdmin Configuration	73
H.1 WebAdmin Configurations.	
H.1.1 Single-Server Configuration.	

H.1.2 Multiserver Configuration.	73
H.2 Installing WebAdmin in a Single-Server Configuration.	74
H.3 Installing WebAdmin in a Multiserver Configuration	75
Appendix I Supported contrib Modules and Extensions Provided by External Projects	76
Index	77

# Chapter 1 Overview of Installation

This chapter provides an overview of FUJITSU Enterprise Postgres installation.

# 1.1 Features that can be Installed

Each FUJITSU Enterprise Postgres feature is installed on the machine that was used to build the database environment.

The following table shows the relationship between the product to be installed and the features that can be installed.

Feature that can be installed	Product name	
	AE	SE
Basic feature (server feature, client feature)	Y	Y

Y: Can be installed

# 1.2 Installation Types

The following installation types are available for FUJITSU Enterprise Postgres:

- New installation
- Reinstallation
- Multi-version installation

# 1.2.1 New Installation

In initial installation, FUJITSU Enterprise Postgres is installed for the first time.

# 1.2.2 Reinstallation

Perform reinstallation to repair installed program files that have become unusable for any reason.

## 1.2.3 Multi-Version Installation

FUJITSU Enterprise Postgres products can be installed on the same server if the product version (indicated by "x" in "x SPz") is different from that of any version of the product that is already installed.

# 1.3 Installation Procedure

The following installation procedures are available for FUJITSU Enterprise Postgres:

- Installation in interactive mode
- Installation in silent mode

Select the installation procedure that corresponds to your environment.



If you have antivirus software installed, the server may crash, fail to start, or stop responding, during installation or when starting up after installation. Set scan exception settings for the installation directory and resource allocation directory so that the files in these directories are not scanned for viruses.

### 1.3.1 Installation in Interactive Mode

Interactive mode enables installation to be performed while the required information is entered interactively.

In the interactive mode installation, the installation state of FUJITSU Enterprise Postgres is determined automatically. Install FUJITSU Enterprise Postgres using one of the following installation types in accordance with the installation state:

- New installation
- Reinstallation
- Multi-version installation

### 1.3.2 Installation in Silent Mode

Silent mode enables installation to be performed without the need to enter any information interactively.

New installations and multi-version installations can be performed in silent mode.

## 1.4 Uninstallation

Uninstallation removes the system files of the installed FUJITSU Enterprise Postgres.

# Chapter 2 Operating Environment

This chapter describes the operating environment required to use FUJITSU Enterprise Postgres.



Refer to "Operating Environment" in the Installation and Setup Guide for Client when installing the FUJITSU Enterprise Postgres client feature at the same time.

## 2.1 Required Operating System

One of the operating systems shown below is required in order to use FUJITSU Enterprise Postgres.

Table 2.1 Operating systems

#### Operating system name

- Microsoft(R) Windows Server(R) 2012 Datacenter
- Microsoft(R) Windows Server(R) 2012 Standard
- Microsoft(R) Windows Server(R) 2012 Essentials
- Microsoft(R) Windows Server(R) 2012 Foundation
- Microsoft(R) Windows Server(R) 2012 R2 Datacenter
- Microsoft(R) Windows Server(R) 2012 R2 Standard
- Microsoft(R) Windows Server(R) 2012 R2 Essentials
- Microsoft(R) Windows Server(R) 2012 R2 Foundation
- Microsoft(R) Windows Server(R) 2016 Datacenter
- Microsoft(R) Windows Server(R) 2016 Standard
- Microsoft(R) Windows Server(R) 2016 Essentials
- Microsoft(R) Windows Server(R) 2019 Datacenter
- Microsoft(R) Windows Server(R) 2019 Standard
- Microsoft(R) Windows Server(R) 2019 Essentials



- The following components of Windows Server(R) 2016 and Windows Server(R) 2019 are not supported:
  - Server Core
  - Nano Server
  - Windows Server Container
- The TCP/IP protocol must be installed.

## 2.2 Related Software

The following table lists the software required to use FUJITSU Enterprise Postgres.

Table 2.2 Related software

No. Software		Version	Product name of FUJITSU Enterprise Postgres		Remarks
	name		AE	SE	
1	Visual Studio	2015 2017 2019	Y	Y	Required when using applications that are developed in Visual Studio.
2	.NET Framework	4.5.1 or later 4.6/4.6.x 4.7/4.7.x 4.8	Y	Y	
3	Perl	5.16	Y	Y	Required when using PL/Perl.
4	Python	3.7.x	Y	Y	Required when using PL/Python based on Python 3.
5	Tcl	8.6	Y	Y	Required when using PL/Tcl.

Y: Can be used



- The following programs are installed during installation of FUJITSU Enterprise Postgres:
  - Microsoft Visual C++ 2010 Redistributable version 10.0.30319.01
  - Microsoft Visual C++ 2013 Redistributable version 12.0.40649.5
  - Microsoft Visual C++ 2015-2019 Redistributable version 14.24.28127.4

Do not uninstall the above programs as they are required for running FUJITSU Enterprise Postgres.

The following table lists servers that can be connected to the FUJITSU Enterprise Postgres client feature.

Table 2.3 Connectable servers

os	Product name	
Windows Linux	- FUJITSU Software Enterprise Postgres Advanced Edition 9.3 or later	
	- FUJITSU Software Enterprise Postgres Standard Edition 9.4 or later	
Solaris	- FUJITSU Software Enterprise Postgres Standard Edition 9.6 or later	

### 2.3 Excluded Software

This section describes excluded software.

### **FUJITSU Enterprise Postgres**

FUJITSU Enterprise Postgres cannot be installed if all the following conditions are met:

- The product version (indicated by "x" in "x SPz") of the product to be installed is the same as that of the installed product
- The editions are different

### Example

In the following cases, FUJITSU Enterprise Postgres cannot be installed as an exclusive product:

- The installed product is FUJITSU Software Enterprise Postgres Standard Edition (64bit) 12
- The product to be installed is FUJITSU Software Enterprise Postgres Advanced Edition (64bit) 12

Other products

There are no exclusive products.

## 2.4 Required Patches

There are no required patches.

### 2.5 Hardware Environment

The following hardware is required to use FUJITSU Enterprise Postgres.

#### Memory

256 MB or more is recommended (at least 128 MB is required).

## 2.6 Disk Space Required for Installation

The following table shows the disk space requirements for new installation of FUJITSU Enterprise Postgres. If necessary, increase the size of the file system.

Table 2.4 Disk space required for installation

Directory	Required disk space (Unit: MB)
Windows system drive	7 + 113 (*1) + 18 (*2)
Installation destination of the Server	596
Installation destination of the WebAdmin	600
Installation destination of the Client (32bit)	414
Installation destination of the Client (64 bit)	432

<sup>\*1:</sup> Required for the installation of the Uninstall (middleware) tool.

## 2.7 Supported System Environment

This section describes the supported system environment.

### 2.7.1 TCP/IP Protocol

FUJITSU Enterprise Postgres supports version 4 and 6 (IPv4 and IPv6) of TCP/IP protocols.



Do not use link-local addresses if TCP/IP protocol version 6 addresses are used.

## 2.7.2 File System

You can install FUJITSU Enterprise Postgres only if the system folder is an NTFS volume.

<sup>\*2:</sup> Required for the installation of FJQSS.

# 2.8 PostgreSQL Version Used for FUJITSU Enterprise Postgres

FUJITSU Enterprise Postgres is based on PostgreSQL 12.1.

# 2.9 Notes on Using Streaming Replication

To use streaming replication, build the primary server and all standby servers using the same FUJITSU Enterprise Postgres version (\*1).

\*1: The product version is indicated by "x" in the notation "x SPz".



Streaming replication cannot be used in combination with Open Source PostgreSQL.

# **Chapter 3 Installation**

This chapter describes the procedures for the installation of FUJITSU Enterprise Postgres.



- The installation must be performed by a user with administrator privileges (a user ID that belongs to the Administrators group).
- Stop all applications before starting the installation.
- The Windows Installer service must be running.
- The remote desktop service is installed in application server mode, it is necessary to switch to install mode by executing the command shown below before installation. Also, after the installation is completed, execute the command shown below to switch back to execute mode.

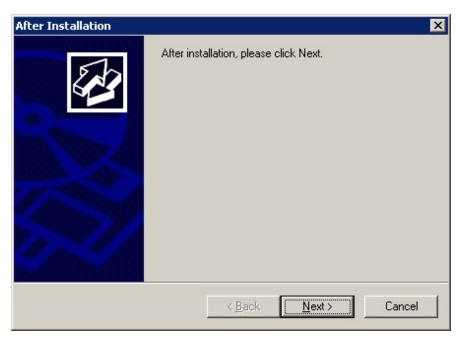
[Before the installation]

CHANGE USER /INSTALL

[After the installation]

CHANGE USER /EXECUTE

- When installing the Fujitsu Enterprise Postgres Client (32 bit), do not specify a destination folder under the environment ProgramFiles variable.
- The following window may be displayed when executing the installation program:



If this window is displayed, perform the following operations:

- 1. Perform the installation steps until the [InstallShield Wizard Complete] window is displayed.
- 2. At the window shown above, click [Next].

3. The window shown below is displayed. Click [Finish].





- If a [User Account Control] dialog box is displayed at the start of the installation, click [Yes] to continue processing:
  - If [No] is clicked, permission to continue is denied and an [Error] dialog box will be displayed.

    To continue the installation, click [Retry] at the [Error] dialog box. To end the installation, click [Cancel].
- If installation is suspended or processing terminates abnormally, the [Program Compatibility Assistant] dialog box may be displayed. Click [This program installed correctly] and continue operation.

### 3.1 Pre-installation Tasks

Check the system environment below before installing FUJITSU Enterprise Postgres.

### Check the disk space

Ensure that there is sufficient disk space to install FUJITSU Enterprise Postgres.

Refer to "2.6 Disk Space Required for Installation" for information on the required disk space.

### Check the installed products and determine the installation method

In Windows, click [All Programs] or [All apps], then [Fujitsu], and then [Uninstall (middleware)]. In the displayed window, check the installed products.

If FUJITSU Enterprise Postgres is already installed, determine which installation method to use:

- Reinstallation
- Multi-version installation

#### Remove applied updates

If you perform reinstallation as the installation method, remove applied updates using the procedure shown below.



If a product is installed without removing applied updates, the following problems will occur:

- Performing reinstallation

If an update with the same update and version number is applied, an error informing you that the update has already been applied is displayed.

#### 1. Display the applied updates

Execute the following command to display the applied updates:

C:\>uam showup

#### 2. Remove the updates

Execute the command below to remove the updates. If an update with the same update number was applied more than once, the updates are removed in order, starting from the highest version number.

C:\>uam remove -i update-number

### Determine the preferred WebAdmin configuration

Starting with FUJITSU Enterprise Postgres 9.5, WebAdmin can be installed in two configurations:

- Single-server
- Multiserver



Refer to "Appendix H Determining the Preferred WebAdmin Configuration" for details.

# 3.2 Installation in Interactive Mode

The installation must be performed by a user with administrator privileges (a user ID that belongs to the Administrators group).



When reinstalling the product, back up the following folder in which the WebAdmin instance management information is stored:

webAdminInstallFolder\data\fepwa

Follow the procedure below to perform the backup.

- 1. Stop the WebAdmin server. Refer to "B.1.3 Stopping the Web Server Feature of WebAdmin" for details.
- 2. Back up the following folder:

webAdminInstallFolder\data\fepwa

Replace the above folder with the backed up folder when the reinstallation is complete.



For installation in interactive mode, default values are set for the installation information. The following settings can be changed for a new installation or a multi-version installation:

- Installation destination
  - It is necessary to specify a local disk as the installation destination of FUJITSU Enterprise Postgres.
  - If using WebAdmin, do not use fullwidth characters or halfwidth katakana characters in [Installation destination folder].
- WebAdmin setup information, if WebAdmin is selected

  To change the port number, confirm that it is an unused port number between 1024 and 49151. Additionally, take note of the Web server port number for the Windows Firewall settings.

The installation procedure is described below.

### 1. Stop applications and programs

When reinstalling the product, all applications and programs that use the product must be stopped.

Before starting the installation, stop the following:

- Applications that use the product
- Instance
- Web server feature of WebAdmin

If you are using WebAdmin, stop WebAdmin.

Refer to "B.1.3 Stopping the Web Server Feature of WebAdmin" for details.

- Mirroring Controller

Execute the mc\_ctl command with the stop mode option specified and stop the Mirroring Controller.

Example

```
> mc_ctl stop -M D:\mcdir\inst1
```

- pgAdmin

### 2. Inserting the DVD

Insert the FUJITSU Enterprise Postgres DVD into the drive.

#### 3. Run the installation

The installation menu will be displayed. Click [Installation].



If the Autorun feature of Windows is disabled, or a remote desktop service (terminal service) is used, the installation program is not automatically started. Execute the following file using [Run] or Windows Explorer.

Z:\autorun.exe

Z: The drive into which the DVD is inserted.

#### 4. Select the products to install

The [Installation product] window will be displayed.

Select the products to install, and then click [Next].

If a selected product can only be reinstalled, refer to "6. Check the settings".



- To develop or execute a 32-bit application in a 64-bit environment, FUJITSU Enterprise Postgres Client (32bit) is required.
- The FUJITSU Enterprise Postgres server component and WebAdmin can be installed on the same machine by selecting the "FUJITSU Enterprise Postgres server component" and the "WebAdmin component".
- If the selected product has been installed, the [Select installation method] window is displayed for each product. To perform a multiversion installation, click [Next].

### 5. Checking the installation content

The [Confirm installation] window will be displayed.

Click [Next] to start the installation.

To modify the settings, select [Modify], and then click [Next]. Follow the on-screen instructions.

If you have not set up WebAdmin during installation, refer to "Appendix B Setting Up and Removing WebAdmin for details.



If using WebAdmin for operation, make a note of the Web server port number displayed in the settings, for use in the Windows firewall settings.

### 6. Check the settings

The [Confirm installation] window is displayed for reinstallation, or if the installation information is modified.

Click [Install] to start the installation.

To change any settings, click [Back].

### 7. Completing installation

The [Installation complete] window is displayed. Click [Finish].

From [All Programs] or [All apps], click [Fujitsu] >> [Uninstall (middleware)]. If the installed product names have been added under [Software Name], installation is complete.

## 3.3 Installation in Silent Mode

Installation in silent mode can be performed only when the installation method is one of the following:

- New installation
- Multi-version installation



See

Refer to the FUJITSU Enterprise Postgres product website for information on installation in silent mode, such as the installation parameters and error messages.

......

The installation procedure is described below.

### 1. Insert the DVD

Insert the "server program" DVD in the DVD drive.

The [Install Menu] window will be displayed automatically. Click [Finish].

### 2. Create an installation parameters CSV file

Consider the features that will be required for system operations, and then create an installation parameters CSV file that uses the following specification format.

```
sectionName, parameterName, value sectionName, parameterName, value :
```



If using WebAdmin for operation, make a note of the Web server port number displayed in the settings (the port number defined in WebPortNumber1), for use in the Windows firewall settings.

......



The template for the installation parameters CSV file is "Z:\sample\sample.csv" (Z is the drive into which the DVD is inserted).

#### 3. Start the command prompt

In Windows, right-click [Command Prompt] and then select [Run as administrator].

#### 4. Run the installation

Execute the command below.

Z:\>silent.bat c:\temp\inspara.csv

Z: The drive into which the DVD is inserted.

Also in the example above, c:\temp\inspara.csv is the installation parameter CSV file name.

If the installer ends in an error, a message is output to the log file and return values are returned.

# Chapter 4 Setup

This chapter describes the setup procedures to be performed after installation completes.

## 4.1 Operating Method Types and Selection

This section describes how to operate FUJITSU Enterprise Postgres.

There are two methods of managing FUJITSU Enterprise Postgres operations - select one that suits your purposes:

The Operation Guide describes the operating method using WebAdmin, and the equivalent operating method using the server commands.

### Simple operation management using a web-based GUI tool (WebAdmin)

Suitable when using frequently used basic settings and operations for operation management.

This method allows you to perform simple daily tasks such as starting the system before beginning business, and stopping the system when business is over, using an intuitive operation.

#### Usage method

Usage is started by using WebAdmin to create the instance.

By using an external scheduler and the pgx\_dmpall command, periodic backups can be performed, which can then be used in recovery using WebAdmin.



Do not use a server command other than pgx\_dmpall and pgx\_keystore or a server application. Operation modes that use server commands and server applications cannot be used in conjunction with WebAdmin. If used, WebAdmin will not be able to manage the instances correctly.

In addition, to perform a backup by copy command from the pgx\_dmpall command, select the operating method using the server commands.

Refer to Reference and the PostgreSQL Documentation for information on server commands and server applications.

### Advanced operation management using server commands

When operating in a system that is automated by operation management middleware (Systemwalker Centric Manager, for example), this method allows you to use more detailed settings and operations and perform higher level operation management.

An overview of the operating method using the GUI, and its relationship with the operating method using the server commands, are shown below.

Refer to the Operation Guide for details.

Operation		Operation with the GUI	Operation with commands
Setup	Creating an instance	WebAdmin is used. The server machine capacity, and the optimum parameter for operations using WebAdmin, are set automatically.	The configuration file is edited directly using the initdb command.
	Creating a standby instance	WebAdmin is used.  WebAdmin performs a base backup of the source instance and creates a standby instance.	A standby instance is created using the pg_basebackup command.
	Changing the configuration files	WebAdmin is used.	The configuration file is edited directly.

Operation		Operation with the GUI	Operation with commands	
Starting and stopping an instance		WebAdmin is used.	The net command or sc command of the operating system is used.	
Creating a database			This is defined using pgAdmin of the GUI tool, or using the psql command or the application after specifying the DDL statement.	
Backing up the database		WebAdmin, or the pgx_dmpall command, is used.	It is recommended that the pgx_dmpall command be used. Recovery to the latest database can be performed.	
Database recovery		WebAdmin is used.	To use the backup that was performed using the pgx_dmpall command, the pgx_rcvall command is used.	
Monitoring	Database errors	The status in the WebAdmin window can be checked. (*1)	The messages that are output to the database server log are monitored (*1)	
	Disk space	The status in the WebAdmin window can be checked. A warning will be displayed if the free space falls below 20%. (*1)	This is monitored using the fsutil command (check free space), and the dir command (check used space), of the operating system, for example.  (*1)	
	Connection status	This can be checked using pgAdmin of the GUI tool, or referencing pg_stat_activity of the standard statistics view from psql or the application.		

<sup>\*1:</sup> This can be used together with system log monitoring using operations management middleware (Systemwalker Centric Manager, for example).



Refer to "Periodic Operations" and "Actions when an Error Occurs" in the Operation Guide for information on monitoring and database recovery.

.....

## 4.2 Preparations for Setup

This section describes the preparation required before setting up FUJITSU Enterprise Postgres.

## 4.2.1 Creating an Instance Administrator

Decide the OS user account that will become the instance administrator. Use either a new user, or a user that already exists.

To create a user in Windows, select [Administrative Tools], [Computer Management], and then create the user in [Local Users and Groups]. Refer to "Help and Support" for details.

The following characters can be used for user names:

- - (hyphen)
- \_ (underscore)
- Space
- A-Z, a-z, 0-9 (alphanumeric)



The following notes apply if using WebAdmin for operations:

- The instance administrator must have a local OS user account.
- After creating the user account of the instance administrator, log in to the operating system. A profile directory is created for the user when logging in to the operating system for the first time. This directory will be used by WebAdmin.
- If changing the password for the user account of the instance administrator, always ensure to stop the instance and log out of WebAdmin before making the change. If you mistakenly change the password while logged in to WebAdmin or while the instance is running, log out from WebAdmin, and then log in again, and stop and start the instance.
- If the password is changed for the user account of the instance administrator, set the changed password using ALTER ROLE WITH ENCRYPTED PASSWORD. This will encrypt the password with the default md5 algorithm. To encrypt the password with the scramsha-256 algorithm, set the password encryption parameter to "scram-sha-256" using SET password\_encryption = 'scram-sha-256' prior to using the ALTER ROLE function.

### 4.2.1.1 Security policy settings

If using commands for operation, security settings that allow logon as a service are required for the operating system user account of the instance administrator in order to start and stop an instance using a Windows service.



If using WebAdmin for operation, these settings are not required as WebAdmin performs the settings automatically for the user ID (operating system user account) that logged in to the database server.

The following explains how to perform the security settings to allow logon as a service:

1. Displaying the Local Security Policy window

In Windows, select [Administrative Tools], and then click [Local Security Policy].

- 2. Setting up security
  - 1. In the [Local Security Policy] window, select [Security Settings], select [Local Policies], and then click [User Rights Assignment].
  - 2. Under [Policy] in the [User Rights Assignment] window, double-click [Log on as a service].
  - 3. In the [Log on as a service Properties] window, set the following:
    - a. Select the [Local Security Setting] tab.
    - b. On the [Local Security Setting] tab, click [Add User or Group].
    - c. In the [Select Users or Groups] window, enter the operating system user account of the instance administrator in [Enter the object names to select].
    - d. Click [OK].
  - 4. In the [Log on as a service Properties] window, click [OK].

## 4.2.2 Preparing Directories for Resource Deployment

Prepare the directories required when creating instances.

#### Considerations when deploying resources

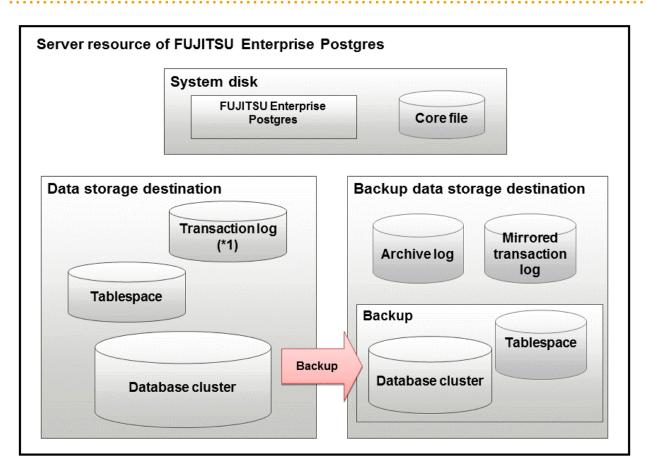
The disk configuration on the resource deployment destination is important, because it affects not only recovery following disk corruption, but normal operation as well. The points for determining the disk configuration are as follows:

1. If the backup data storage destination and the data storage destination are both lost, it will not be possible to recover the data, so deploy them to separate disks.

- 2. To shorten the recovery time following a single disk fault, deploy the system disk and data storage destination to separate disks.
- 3. The backup data storage destination requires at least double the capacity of the data storage destination, so deploy it to the disk with the most space available.
- 4. When large amounts of data are updated, the write-to load for the data storage destination, transaction log storage destination, and backup data storage destination (mirrored transaction log) will also be great. For this reason, deploy them to separate disks, out of consideration for performance.



- When using the volume manager provided by the operating system, be aware of which physical disk the file system has been created on, for example, by deploying the data storage destination and the backup data storage destination to separate disks.
- If using WebAdmin, specify an NTFS volume for the data storage destination and backup data storage destination. A network drive cannot be specified.



\*1: To distribute the I/O load, place the transaction log on a different disk from the data storage destination.

Resource	Role	
Database cluster	The area where the database is stored. It is a collection of databases managed by an instance.	
Tablespace	Stores table files and index files in a separate area from the database cluster.	
	Specify a space other than that under the database cluster.	
Transaction log	Stores log information in preparation for a crash recovery or rollba	
	This is the same as the WAL (Write Ahead Log).	

Resource	Role
Archive log	Stores log information for recovery
Mirrored transaction log (mirrored WAL)	Enables a database cluster to be restored to the state immediately before an error even if both the database cluster and transaction log fail when performing backup/recovery operations using the pgx_dmpall command or WebAdmin.
Corefile	FUJITSU Enterprise Postgres process corefile output when an error occurs with a FUJITSU Enterprise Postgres process.

### Examples of disk deployment

The following are examples of disk deployment:

Number of disks	Disk	Deployment
3 System disk		FUJITSU Enterprise Postgres program
		Corefile
Connected physical disk Data storage destination		Data storage destination, transaction log storage destination
	Connected physical disk	Backup data storage destination
2	System disk	FUJITSU Enterprise Postgres program
		Corefile
		Data storage destination, transaction log storage destination
	Connected physical disk	Backup data storage destination

### Proposal for disk deployment using WebAdmin

To generate an instance using WebAdmin, we recommend an optimum deployment that takes into account the status of all disks at the time of instance generation, and items 1 to 3 in the "Considerations when deploying resources" subheading above, based on the limitations below (note that a different deployment can also be specified).

- The instance administrator has read and write privileges for the volumes.

### **Preparing directories**

The directories to be prepared depend on the way that you create the instances.

The following shows the directories that need to be prepared:

Directory to be prepared	Using WebAdmin	Using the initdb command
Data storage destination	Y (*1)	Y
Backup data storage destination	O (*1) (*4)	0
Transaction log storage destination	O (*1) (*2)	О
Corefile output destination	N (*3)	0

Y: Required

O: Optional

N: Not required

\*1: WebAdmin automatically creates a directory

\*2: The default is to create in a directory in the data storage destination. When it is necessary to distribute the I/O load for the database data and the transaction log, consider putting the transaction log storage destination on a different disk from the data storage destination

\*3: The corefile path is as follows:

userProfileFolder\localSettingsFolder\Fujitsu\fsep\_version\instanceNamePortNumber\core

version: product version\_WA\_architecture

Note: The product version is normally the version of WebAdmin used to create the instance. For example, WebAdmin 12 allows a user to create a FUJITSU Enterprise Postgres 9.6 instance on a database server having WebAdmin 9.6. In this case, because WebAdmin 9.6 is used to create the instance, the product version will be "96".

PortNumber: port number specified when creating the instance

#### Example

C:\Users\naomi\AppData\Local\Fujitsu\fsep\_120\_WA\_64\myinst27599\core

To change the output destination, specify in the core\_directory parameter and core\_contents parameter in postgresql.conf. Refer to "Parameters" in the Operation Guide for information on the settings for these parameters.

\*4: This directory is required when instance backup is enabled.



- The directories must meet the following conditions:
  - The directory owner must be the OS user account that you want to be the instance administrator
  - The directory must have write permission
  - The directory must be empty
- If using WebAdmin, network drives cannot be used.
- If using WebAdmin, the following halfwidth characters can be used for directory names:
  - \(backslash)
  - - (hyphen)
  - \_ (underscore)
  - : (colon)
  - Space
  - A-Z, a-z, 0-9 (alphanumeric)
- If anti-virus software is used, set scan exception settings for folders so that none of the server resources that comprise FUJITSU Enterprise Postgres are scanned for viruses. Alternatively, if the server resources that comprise FUJITSU Enterprise Postgres are to be scanned for viruses, stop the instance and perform the scan when tasks that use FUJITSU Enterprise Postgres are not operating.

### Confirm and configure directory access permissions

If the instance administrator user has "Administrator" permissions (user ID belonging to the Administrators group), it is necessary to configure the settings so that each directory inherits the file and directory access permissions for the instance administrator user.

Therefore, ensure that the setting to inherit permissions has been configured.

The following is an explanation on how to confirm and configure the settings.

How to confirm access permissions

Perform the following operations in Windows Explorer on the directories to be prepared in advance:

- 1. Right-click on the applicable directory, and then click [Properties] from the menu that is displayed.
- 2. In the [applicableDir Properties] window, select [Security] >> [Advanced].
- 3. In the [Advanced Security Settings for *applicableDir*] window, and in the [Permission entries] list under the [Permissions] tab, confirm that [Applies to] for the instance administrator user is "This folder, subfolders and files".

4. Click [OK].

How to configure the access permissions

Perform the following operations in Windows Explorer if there are any directories that have not been configured for the access permissions to be inherited.

- 1. Right-click on the applicable directory, and then click [Properties] from the menu that is displayed.
- 2. In the [applicableDir Properties] window, select [Security] >> [Advanced].
- 3. In the [Advanced Security Settings for applicableDir] window, click [Add].
- 4. In the [Permission Entry for applicableDir] window, click [Select a principal].
- 5. In the [Select User or Group] window, enter the instance administrator user name as the object name to select, and then click [OK].
- 6. In the [Permission Entry for *applicableDir*] window, set [This folder, subfolders and files] for [Apply to:] ([Applies to:] for Windows Server(R) 2016 and Windows Server(R) 2019), and under [Basic permissions], allow read and write permissions, and then click [OK].
- 7. In the [Advanced Security Settings for *applicableDir*] window, confirm that the instance administrator user has been added, with [This folder, subfolders and files] set for [Apply to] ([Applies to] for Windows Server(R) 2016 and Windows Server(R) 2019) in the [Permission entries] list.
- 8. Click [OK].



The access permissions can also be configured using the icacls command provided by the operating system.

The following is an execution example in which the application destination is set to "(OI)(CI)" and the access permissions are set to "(F) (Full access permissions)" when the data storage destination is "D:\database\inst1" and the instance administrator user is "fsepuser":

```
>icacls D:\database\inst1 /grant fsepuser:(OI)(CI)(F)
processed file: D:\database\inst1
Successfully processed 1 files; Failed processing 0 files
```

## 4.2.3 Estimating Resources

Estimate the resources to be used on the FUJITSU Enterprise Postgres.

Refer to "Appendix E Estimating Database Disk Space Requirements" for information on estimating database disk space requirements.

Refer to "Parameters automatically set by WebAdmin according to the amount of memory" when creating multiple instances with WebAdmin.

Refer to "Appendix F Estimating Memory Requirements" when creating instances with the initdb command, to estimate memory usage.

## 4.2.4 Windows Firewall Settings

This section explains the Windows firewall settings required if using WebAdmin for operation.

These settings are not required if using server commands for operation.

If the Windows firewall feature is to be enabled, you should enable a port number on the Web server. The following explains how to enable a port number:

### Windows Server(R) 2019:

- 1. Select [Systems and Security] from [Control Panel] and click [Windows Defender Firewall].
- 2. In the [Windows Defender Firewall] window, click [Advanced settings].
- 3. In the [Windows Defender Firewall with Advanced Security] window, click [Inbound Rules] on the left side of the window.
- 4. Click [New Rule] on the right side of the window.

- 5. In the [New Inbound Rule Wizard] window, select [Port], and then click [Next].
- 6. Select [TCP] and [Specific local ports], then specify the Web server port number specified during the WebAdmin setup, and then click [Next].
- 7. Select [Allow the connection], and then click [Next].
- 8. Select the profiles for which this rule applies, and then click [Next].
- 9. In [Name], specify the desired name, and then click [Finish].
- 10. In the [Windows Defender Firewall with Advanced Security] window, check if the added rule is enabled under [Inbound Rules] in the center of the window.

#### In cases other than the above:

- 1. Select [Systems and Security] from [Control Panel] and click [Windows Firewall].
- 2. In the [Windows Firewall] window, click [Advanced settings].
- 3. In the [Windows Firewall with Advanced Security] window, click [Inbound Rules] on the left side of the window.
- 4. Click [New Rule] on the right side of the window.
- 5. In the [New Inbound Rule Wizard] window, select [Port], and then click [Next].
- 6. Select [TCP] and [Specific local ports], then specify the Web server port number specified during the WebAdmin setup, and then click [Next].
- 7. Select [Allow the connection], and then click [Next].
- 8. Select the profiles for which this rule applies, and then click [Next].
- 9. In [Name], specify the desired name, and then click [Finish].
- 10. In the [Windows Firewall with Advanced Security] window, check if the added rule is enabled under [Inbound Rules] in the center of the window.

### 4.2.5 Preparing for Output to the Event Log

This section provides an explanation on the preparation to be carried out if you are outputting error logs to the event log.

If outputting error logs to the event log, you should register an event source name beforehand.

If you do not register an event source name, the message content output to the event log may be incomplete.

Due to the default event source name "FUJITSU Enterprise Postgres Server" being output to the event log when using the following commands, you should register this default event source name beforehand:

- pg\_ctl command
- pgx\_dmpall command
- pgx\_rcvall command

The following is an example in which the DLL of a 64-bit product is registered under the default event source name:

```
> regsvr32 "C:\Program Files\Fujitsu\fsepv<x>server64\lib\pgevent.dll"
```

Note that this step is not required if using WebAdmin to create an instance.

### If using multiple instances

You can output messages corresponding to the event source name assigned by the user, so that messages output to the event log can be identified by instance.

The following is an example in which the DLL of a 64-bit product is registered under the event source name "FUJITSU Enterprise Postgres inst1":

> regsvr32 /n /i:"FUJITSU Enterprise Postgres instl" "C:\Program Files\Fujitsu\fsepv<x>server64\lib
\pgevent.dll"

You will need to edit the parameters for each instance, therefore, after creating an instance, refer to "4.5.1 Error Log Settings" when performing this setting.

### If installing multiple versions

If FUJITSU Enterprise Postgres is already installed on the same machine, search for the key below in Registry Editor, and make a note of the path of the registered DLL. Afterwards, register a new DLL under the default event source name.

Use the DLL path that you made a note of in the above step when re-registering the default event source name during an uninstall.

FUJITSU Enterprise Postgres Server



Refer to "Registering Event Log on Windows" in "Server Setup and Operation" in the PostgreSQL Documentation for information on how to register event source names.

......

## 4.3 Creating Instances

There are two methods that can be used to create an instance:

- 4.3.1 Using WebAdmin
- 4.3.2 Using the initdb Command

### Creating multiple instances

Multiple instances can be created.

The memory allocated needs to be adjusted when multiple instances are created with WebAdmin (refer to "Parameters automatically set by WebAdmin according to the amount of memory" for details).

### Features that cannot be set up using WebAdmin

The "Storage data protection using transparent data encryption" feature cannot be set up using WebAdmin.

To set up this feature in an instance created by WebAdmin, perform the additional setup tasks detailed in "Storage Data Protection using Transparent Data Encryption" in the Operation Guide.



- Instances created using the initidb command (command line instances) can be managed using WebAdmin, however, they must first be imported into WebAdmin. Refer to "4.3.1.4 Importing Instances" for details.

- Always use WebAdmin to delete instances that were created or imported using WebAdmin. Because WebAdmin management information cannot be deleted, WebAdmin will determine that the instance is abnormal.
- Databases with the names 'template0' and 'template1' are automatically created when an instance is created.

  These databases are used as the templates for databases created later. Furthermore, a default database with the name 'postgres' is automatically created, which will be used with FUJITSU Enterprise Postgres commands. It is important that you do not delete these databases created by default.
- When an instance that uses WebAdmin is created successfully, the following Windows service is registered:

fsep\_version\_WA\_architecture\_userName\_instanceNamePortNumber

The account and password of the instance administrator are registered in the Windows service.

If the password for this account is changed, you must also change the password registered in the service.

Change this at the Properties window registered in the Windows service.

Note: The product version is normally the version of WebAdmin used to create the instance. In addition, WebAdmin 12 also allows a user to create a FUJITSU Enterprise Postgres 9.6 instance on a database server having WebAdmin 9.6. In this case the product version will be "96".

- Refer to "4.5.2 Configuring Automatic Start and Stop of an Instance" for information on how to start and stop the operating system of the database server, and how to start and stop linked instances.

### 4.3.1 Using WebAdmin

This section describes how to create an instance using WebAdmin.

WebAdmin must be set up correctly before it can be used. Refer to "B.1 Setting Up WebAdmin" for details. Additionally, if WebAdmin needs to be configured to use an external repository database, refer to "B.3 Using an External Repository for WebAdmin" for details.

It is recommended to use the following browsers with WebAdmin:

- Internet Explorer 11
- Microsoft Edge (Build41 or later)

WebAdmin will work with other browsers, such as Firefox and Chrome, however, the look and feel may be slightly different.

Configure your browser to allow cookies and pop-up requests from the server on which FUJITSU Enterprise Postgres is installed.

Refer to "Appendix A Recommended WebAdmin Environments" for information on how to change the pop-up request settings and other recommended settings.



- WebAdmin does not run in Windows(R) safe mode.
- If the same instance is operated from multiple WebAdmin windows, it will not work correctly.
- If the same instance is operated from multiple WebAdmin versions, it will not work correctly. Always use the latest version of WebAdmin for instance operations.
- For efficient use of WebAdmin, it is recommended not to use the browser [Back] and [Forward] navigation buttons, the [Refresh] button, and context-sensitive menus, including equivalent keyboard shortcuts.
- Copying and pasting the WebAdmin URLs are not supported. Additionally, bookmarking of WebAdmin URLs is not supported.
- It is recommended to match the language between the instance server locale and WebAdmin.
- WebAdmin supports only two languages: English and Japanese.
- It is recommended to change the WebAdmin language setting from the instance details page only.
- It is recommended to operate WebAdmin using the WebAdmin launcher.
- WebAdmin uses the labels "Data storage path", "Backup storage path" and "Transaction log path" to indicate "data storage destination", "backup data storage destination" and "transaction log storage destination" respectively. In this manual these terms are used interchangeably.
- If the browser was not operated for a fixed period (about 30 minutes), the session will time out and the login page will be displayed again for the next operation.
- Port access permissions
  - If a port is blocked (access permissions have not been granted) by a firewall, enable use of the port by granting access. Refer to the vendor document for information on how to grant port access permissions.

    Consider the security risks carefully when opening ports.
- When creating or importing an instance in WebAdmin, set the log\_directory parameter in postgresql.conf in the following format: log\_directory='userProfileFolder\\localSettingsFolder\\Fujitsu\\fsep\_version\\instanceNamePortNumber\\log'

Example: userProfileFolder\\localSettingsFolder\will be C:\\Users\\userName\\AppData\\Local.

### 4.3.1.1 Logging in to WebAdmin

This section describes how to log in to WebAdmin.

### Startup URL for WebAdmin

In the browser address bar, type the startup URL of the WebAdmin window in the following format:

http://hostNameOrIpAddress:portNumber/

- hostNameOrIpAddress: Host name or IP address of the server where WebAdmin is installed
- portNumber. Port number of WebAdmin. The default port number is 27515.

#### Example

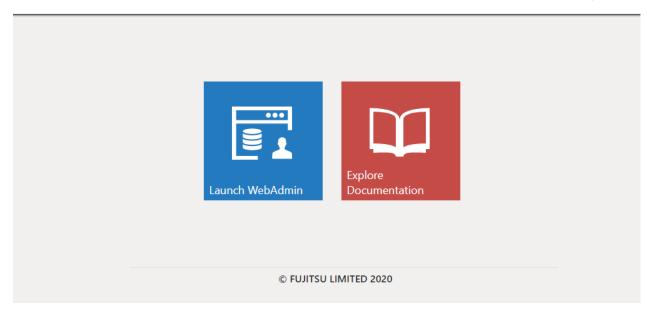
For a server with IP address "192.0.2.0" and port number "27515":

http://192.0.2.0:27515/

The startup URL window shown below is displayed. From this window you can log in to WebAdmin or access the product documentation.

### **FUJITSU Enterprise Postgres**





### Logging in to WebAdmin

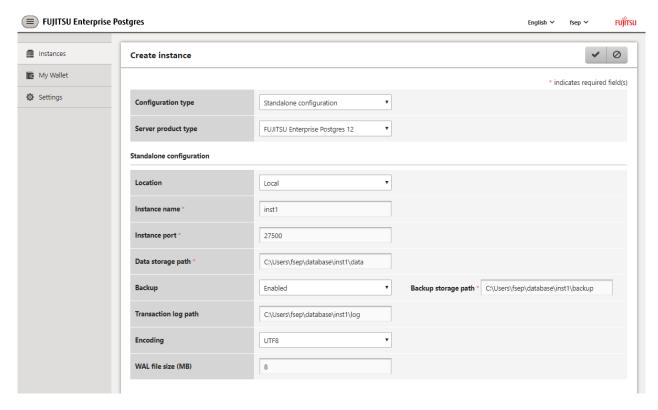
Click [Launch WebAdmin] in the startup URL window to start WebAdmin and display the login window. Enter the instance administrator user name (operating system user account name) and password, and log in to WebAdmin. User credential (instance administrator user ID and password) should not contain hazardous characters. Refer to "Appendix C WebAdmin Disallow User Inputs Containing Hazardous Characters".

### 4.3.1.2 Creating an Instance

This section describes how to create an instance.

- 1. Start WebAdmin, and log in to the database server.
- 2. In the [Instances] tab, click -

3. Enter the information for the instance to be created.



#### Enter the following items:

- [Configuration type]: Whether to create a standalone instance or an instance that is part of a cluster.
- [Server product type]: Sets which of the following instances to create:
  - FUJITSU Enterprise Postgres 9.5 Instances
  - FUJITSU Enterprise Postgres 9.6 Instances
  - FUJITSU Enterprise Postgres 10 Instances
  - FUJITSU Enterprise Postgres 11 Instances
  - FUJITSU Enterprise Postgres 12 Instances

The default is "FUJITSU Enterprise Postgres 12".

WebAdmin can create and manage instances compatible with the following, but new functionality in FUJITSU Enterprise Postgres 12 may not support the instance or it may be disabled.

- FUJITSU Enterprise Postgres 9.5
- FUJITSU Enterprise Postgres 9.6
- FUJITSU Enterprise Postgres 10
- FUJITSU Enterprise Postgres 11
- [Location]: Whether to create the instance in the server that the current user is logged into, or in a remote server. The default is "Local", which will create the instance in the server machine where WebAdmin is currently running.
- [Instance name]: Name of the database instance to manage

The name must meet the conditions below:

- Maximum of 16 characters
- The first character must be an ASCII alphabetic character
- The other characters must be ASCII alphanumeric characters

- [Instance port]: Port number of the database server
- [Data storage path]: Directory where the database data will be stored
- [Backup]: Whether to enable or disable the WebAdmin backup feature. The default is "Enabled". Select "Disabled" to disable all backup and restore functionality for the instance. If "Enabled" is selected, enter the following item:
  - [Backup storage path]: Directory where the database backup will be stored
- [Transaction log path]: Directory where the transaction log will be stored
- [Encoding]: Database encoding system
- [WAL file size]: Allow the WAL file size to be set when creating an instance. The default is 16 MB if the field is blank. The size specified must be a power of 2 between 1 and 1024. This option is not available for standby instances.

If "Remote" is selected for [Location], enter the following additional items:

- [Host name]: Name of the host where the instance is to be created
- [Operating system credential]: Operating system user name and password for the remote machine where the instance is to be created
- [Remote WebAdmin port for standalone]: Port in which WebAdmin is accessible in the remote machine



- Refer to "4.2.2 Preparing Directories for Resource Deployment" "Considerations when deploying resources" for information on points to consider when determining the data storage path, backup storage path, and transaction log path.
- The following items can be modified after the instance has been created. These items cannot be modified on instances that have compatibility with FUJITSU Enterprise Postgres 9.5.
  - Instance name
  - Port number
  - Backup storage path

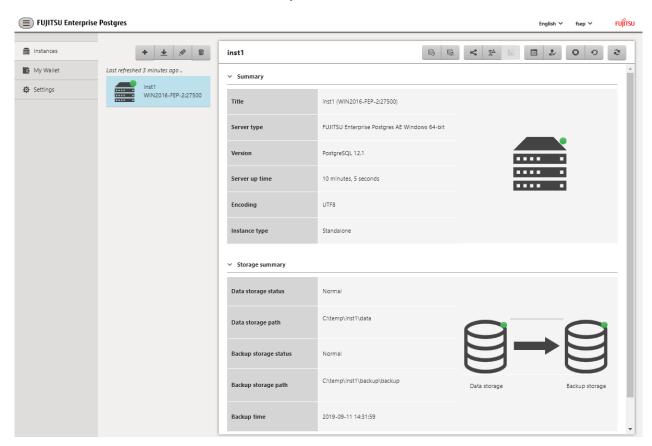
Refer to "Editing instance information" for details.

- Do not specify shortcuts for the data storage path, backup storage path, or transaction log path.
- In the instance that is created using WebAdmin, the locale of the character set to be used in the database, and the locale of the collating sequence, are fixed using C.
- The following characters can be used for the data storage path, backup storage path, and transaction log path:
  - \ (backslash)
  - - (hyphen)
  - \_ (underscore)
  - : (colon)
  - Space
  - A-Z, a-z, 0-9 (alphanumeric)
- Instance administrator read and write permissions are required for the data storage path, backup storage path, and transaction log path.
- For the port number, specify an unused port number in the following range:
  - 1024 to 49151
- Make a note of the port number for use in the Windows firewall settings.
- Refer to "4.5.2 Configuring Automatic Start and Stop of an Instance" for information on configuring the automatic start and stop of instances.

- For enhanced security, WebAdmin encrypts the superuser password using SCRAM-SHA-256 authentication for all FUJITSU Enterprise Postgres 10 or later instances. The client/driver must therefore support SCRAM-SHA-256 authentication if they need to connect to FUJITSU Enterprise Postgres 10 or later instances created by WebAdmin with superuser credentials.
- Host name and Operating system credential (Operating system user name and password) should not contain hazardous characters. Refer to "Appendix C WebAdmin Disallow User Inputs Containing Hazardous Characters".
- 4. Click **v** to create the instance.

If the instance is created successfully, a message indicating the same will be displayed.

5. The instance will be started when it is created successfully.



6. Back up the basic information that was set

Back up the WebAdmin management information periodically to ensure operational continuity when a fault occurs on the system disk. Follow the procedure below to perform the backup.

- a. Stop the WebAdmin server. Refer to "B.1.3 Stopping the Web Server Feature of WebAdmin" for details.
- b. Back up the following directory:

webAdminInstallDir\data\fepwa



The following message is output during startup of an instance when the startup process is operating normally, therefore, the user does not need to be aware of this message.

.....

......

FATAL: the database system is starting up (XXXXX)

### 4.3.1.3 Changing Instance Settings

You can change the information that is set when an instance is created.

Change the following settings to suit the operating and management environment for FUJITSU Enterprise Postgres.

- Instance configuration
  - Character encoding
  - Communication
  - SQL options
  - Memory
  - Streaming replication
- Changing client authentication information
- Editing instance information



These settings are the same as the parameters that can be set in the files shown below. Refer to "Appendix D Configuring Parameters" for information on the equivalence relationship between the item name and the parameter.

•••••••••••

- postgresql.conf
- pg\_hba.conf



The files shown below can also be modified directly, however if a parameter not described in "Appendix D Configuring Parameters" was edited by mistake, WebAdmin may not run correctly.

- postgresql.conf
- pg\_hba.conf

### Instance configuration

- 1. Start WebAdmin and log in to the database server.
- 2. In the [Instances] tab, click =
- 3. Click to change the configuration.
- 4. Click to save your changes.



Select a client-side encoding system that can be converted to/from the database encoding system. Refer to "Automatic Character Set Conversion Between Server and Client" in "Server Administration" in the "PostgreSQL Documentation" for information on the encoding system combinations that can be converted.

.....

### Changing client authentication information

- 1. Start WebAdmin and log in to the database server.

Click • to register new authentication information.

To change authentication information, select the information, and then click



To delete authentication information, select the information, and then click in



When creating the instance, do not delete the entry below, because it is a connection required for WebAdmin to monitor the operational status of the database:

Type= host, Database=all, User=all, and Method=md5

### Editing instance information

Use the [Edit instance] page to modify the following items for an instance:

- Instance name
- Port number
- Backup storage path
- 1. In the [Instances] tab, click . The [Edit instance] page is displayed.
- 2. Modify the relevant items.

If [Backup storage path] is changed, [Backup management] is enabled. Select the required option:

Retain existing backup: Create a backup in [Backup storage path] and retain the existing backup in its original location.

Copy existing backup to new path: Copy the existing backup to [Backup storage path]. A new backup will not be created. The existing backup will be retained in its original location.

Move existing backup to new path: Move the existing backup to [Backup storage path]. A new backup will not be created.

Remove existing backup: Create a backup in [Backup storage path]. The existing backup will be removed.

3. Click to save your changes.



- The [Edit instance] page is also displayed when the user selects 'Navigate to the "Edit instance" page' from the [Anomaly Error] dialog box. Refer to "Anomaly Detection and Resolution" in the Operation Guide for information on what takes place when an anomaly is detected.

- When [Instance name] or [Instance port] is modified, the log\_directory and core\_directory parameters in postgresql.conf are updated. Also, the specified directories are created if they do not exist.

Refer to "4.3.1.4 Importing Instances" for information on the format of these directories.

### 4.3.1.4 Importing Instances

Instances can be created using WebAdmin, or via the command line using the initdb command. Instances created using the initdb command (command line instances) can be managed using WebAdmin, however, they must first be imported into WebAdmin.

This section explains how to import command line instances into WebAdmin.

- 1. In the [Instances] tab, click . The [Import instance] page is displayed.
- 2. Enter the information for the instance being imported. Refer to "4.3.1.2 Creating an Instance" for information on the items that need to be entered.
- 3. Click to import the instance.



- Importing neither starts nor stops the instance.
- A Windows service is automatically registered when an instance is imported into WebAdmin. If a Windows service was registered by the user prior to importing the instance, that service will not be deleted. After importing the instance into WebAdmin, it is recommended to discontinue the use of the user-created service.
- The following restrictions apply to instance import:
  - Any instance already managed by WebAdmin cannot be imported again.
  - The postgresql.conf file must be located in the same directory as [Data storage path].
  - Read/write permissions are required for [Data storage path].
  - The location specified in postgresql.conf for the following files must not have been changed:
    - hba\_file
    - ident\_file
  - If the instance is part of a cluster that is monitored by Mirroring Controller, WebAdmin will be unable to detect the Mirroring Controller settings.
  - Instances making use of Mirroring Controller functionality should not be imported, because subsequent operations on those instances may cause unexpected and undesirable side-effects.
  - It is not possible to import and operate an instance that uses a directory mounted by Network File System (NFS).
  - You must make the following changes to the parameters in postgresql.conf prior to importing the instance in WebAdmin.

Parameter	Requirements
port	The port parameter should be uncommented.

The log\_directory and core\_directory parameters in postgresql.conf are updated during import. Also, the specified directories are created if they do not exist.

The format of these directories is as follows:

log\_directory: 'userProfileFolder\\localSettingsFolder\\Fujitsu\\fsep\_version\\instanceNamePortNumber\\log'

core\_directory: 'userProfileFolder\\localSettingsFolder\\Fujitsu\\fsep\_version\\instanceNamePortNumber\\core'

version: product version\_WA\_architecture

PortNumber: port number specified when creating the instance

Examples:

 $log\_directory: \label{log_directory} Local \label{log_directory: C:\local} log\_directory: \local \$ 

 $core\_directory: \label{local} Local \labell} Local \label{local} Local \label{local} Local \labell{local$ 

- When a standby instance is imported, a valid entry, using the IP address of the standby instance, must exist in the pg\_hba.conf file of the master instance to allow the standby instance to connect to the master instance.
- When a standby instance is imported, the value for "host" in the primary\_conninfo parameter of postgresql.auto.conf should match the host name of the master instance.
- When a standby instance is imported, you cannot specify "passfile" in the primary\_conninfo parameter of postgresql.auto.conf. Be sure to specify "password".
- Instances created by other operating systems cannot be imported.
- If a FUJITSU Enterprise Postgres 10 or later instance is being imported while it is running, WebAdmin will encrypt the superuser password using SCRAM-SHA-256 authentication.

## 4.3.2 Using the initdb Command

This section describes the procedure to create an instance using the initdb command.



If a port is blocked (access permissions have not been granted) by a firewall, enable use of the port by granting access. Refer to the vendor document for information on how to grant port access permissions.

Consider the security risks carefully when opening ports.

### 4.3.2.1 Creating an Instance

Create an instance, with the database cluster storage destination specified in the PGDATA environment variable or in the -D option. Furthermore, the user that executed the initdb command becomes the instance administrator.



- Instances created using the initdb command (command line instances) can be managed using WebAdmin, however, they must first be imported into WebAdmin. Refer to "4.3.1.4 Importing Instances" for details.

- If creating multiple instances, ensure that there is no duplication of port numbers or the directories that store database clusters.



Refer to "initdb" in "Reference" in the PostgreSQL Documentation for information on the initdb command.

The procedure to create an instance is described below.

1. Use the OS user account that you want as the instance administrator.

Connect with the server using the OS user account that you want as the instance administrator.

2. Configure the environment variables

Configure the environment variables in the server with the newly created instance.

Set the following environment variables:

- PATH environment variables

Add installDir\bin and installDir\lib.

### Example

The following is a setting example for environment variables in which "C:\Program Files\Fujitsu\fsepv<*x*>server64" is used as the installation folder:

Note that "<x>" indicates the product version.

> SET PATH=C:\Program Files\Fujitsu\fsepv<x>server64\bin;C:\Program Files\Fujitsu\fsepv<x>server64\lib;%PATH%

#### 3. Create a database cluster

Create the database cluster with the initdb command, specifying the storage destination directory.

Specify the transaction log storage destination and the locale setting option as required.

### Example

> initdb -D D:\database\instl --waldir=E:\transaction\instl --lc-collate="C" --lc-ctype="C" --encoding=UTF8

# Point

In some features, instance names are requested, and those names are required to uniquely identify the instance within the system. These features allow names that conform to WebAdmin naming conventions, so refer to the following points when determining the names:

- Maximum of 16 characters
- The first character must be ASCII alphabetic character
- The other characters must be ASCII alphanumeric characters



- To balance I/O load, consider deploying the transaction log storage destination to a disk device other than the database cluster storage destination and the backup data storage destination.

- Messages may not display correctly if a value other than "C" is specified as the display language for messages.
- Specify "C" for collation and character category. Performance deteriorates if you specify a value other than "C", although the behavior will follow the rules for particular languages, countries and regions. Furthermore, this may need to be revised when running applications on systems with different locales.

For example, specify as follows:

```
initdb --locale="C" --lc-messages="C"
initdb --lc-collate="C" --lc-ctype="C"
```

- Specify an encoding system other than SQL\_ASCII for the database. If SQL\_ASCII is used, there is no guarantee that the encryption system for data in the database will be consistent, depending on the application used to insert the data.



#### See

Refer to "Locale Support" in "Localization" in "Server Administration" in the Postgre SQL Documentation for information on locales.

4. Set port number.

Specify a port number in the port parameter of postgresql.conf. Ensure that the specified port number is not already used for other software. If a port number is not specified, "27500" is selected.

Register the specified port numbers in the C:\Windows\System32\drivers\etc\services file if WebAdmin is used to create other instances. WebAdmin uses the services file to check if port numbers specified as available candidates have been duplicated.

Register any name as the service name.



Make a note of the port number for use in the Windows firewall settings.

5. Set the corefile output destination.

Specify the output destination of the corefile, which can later be used to collect information for investigation, by setting the core\_directory and core\_contents parameters of postgresql.conf.



#### See

Refer to "Parameters" in the Operation Guide for information on the settings for these parameters.

6. Set the backup storage destination.

Specify the backup data storage destination and other backup settings when backup is to be performed as a provision against database errors.



#### See

Refer to "Backup Methods" in the Operation Guide for information on specifying backup settings.

7. Register an instance in the Windows service

Use the register mode of the pg\_ctl command to register an instance in the Windows service.

Specify the service name, user name, password and path to the instance in the pg\_ctl command, and register the instance in the Windows service.

#### Example

The following is a setting example, in which the service name to register is "inst1", the user name is "fepuser", and the storage destination directory of the database cluster is "D:\database\inst1":

> pg\_ctl register -N "instl" -U fepuser -P \*\*\*\*\*\*\* -D D:\database\instl



- This command must be executed by an instance administrator user with administrator privileges. Execute the command from the [Administrator: Command Prompt] window. Right-click [Command Prompt], and then select [Run as administrator] from the menu to display the [Administrator: Command Prompt] window.

- For the following reasons, a user name and password must always be specified:

Because the Windows service is started up by the Network Service account, all user resources are created as resources of that account. This can result in error events such as failing to access database resources and not being able to perform backups/recovery.

Note that if not specifying a user name and password for security reasons, you should specify the account from the Windows services list immediately after registering the instance in Windows services.

- When entering the password that is specified in the pg\_ctl command, for security reasons, you should be careful not to allow other users to access it.

Commands such as sc query can be used to check the registration status.

### 8. Start an instance

Use the following procedure to start the service:

a. Display the [Services] window.

In Windows, select [Administrative Tools], and then click [Services].

b. Start the service

From the services list, select the instance name that you wish to start, and click [Start Service].

If using commands to start the service, specify the service name using either the net start command or sc start command from the command prompt.

# 4.4 Configuring Remote Connections

This section describes the settings required when connecting remotely to FUJITSU Enterprise Postgres from a database application or a client command.

### 4.4.1 When an Instance was Created with WebAdmin

#### Settings related to connection

The default is to accept connections from remote computers to the database.

Change "listen\_addresses" in postgresql.conf to modify the default behavior.

Refer to "Appendix D Configuring Parameters" for information on postgresql.conf.

#### **Client Authentication Information settings**

The following content is set by default when WebAdmin is used to create an instance.

- Authentication of remote connections from local machines is performed.

When changing Client Authentication Information, select [Client Authentication] from [Setting], and then change the settings.

### 4.4.2 When an Instance was Created with the initdb Command

### **Connection settings**

The default setting only permits local connections from the client to the database. Remote connections are not accepted.

Change "listen\_addresses" in postgresql.conf to perform remote connection.

All remote connections will be allowed when changed as shown below.

#### Example

```
listen_addresses = '*'
```

Also, configure the parameters shown below in accordance with the applications and number of client command connections.

Parameter name	Parameter description
superuser_reserved_connections	Number of connections reserved for database maintenance, for example backup or index rebuilding. If you need to simultaneously perform a large number of processes that exceed the default value, change this value accordingly.
max_connections	Set the value as:  numberOfSimultaneousConnectionsToInstance +
	superuser_reserved_connections

### Client authentication information settings

When trying to connect from a client to a database, settings are required to determine whether the instance permits connections from the client - if it does, then it is possible to make settings to determine if authentication is required.



Refer to "The pg\_hba.conf File" in "Server Administration" in the PostgreSQL Documentation for details.

### 4.4.3 Windows Firewall Settings

If the Windows firewall feature is to be enabled, you should enable a port number on the database server. The following explains how to enable a port number:

### Windows Server(R) 2019:

- 1. Select [Systems and Security] from [Control Panel] and click [Windows Defender Firewall].
- 2. In the [Windows Defender Firewall] window, click [Advanced settings].

- 3. In the [Windows Defender Firewall with Advanced Security] window, click [Inbound Rules] on the left side of the window.
- 4. Click [New Rule] on the right side of the window.
- 5. In the [New Inbound Rule Wizard] window, select [Port], and then click [Next].
- 6. Select [TCP] and [Specific local ports], then specify the Web server port number specified during the WebAdmin setup, and then click [Next].
- 7. Select [Allow the connection], and then click [Next].
- 8. Select the profiles for which this rule applies, and then click [Next].
- 9. In [Name], specify the desired name, and then click [Finish].
- 10. In the [Windows Defender Firewall with Advanced Security] window, check if the added rule is enabled under [Inbound Rules] in the center of the window.

#### In cases other than the above:

- 1. Select [Systems and Security] from [Control Panel] and click [Windows Firewall].
- 2. In the [Windows Firewall] window, click [Advanced settings].
- 3. In the [Windows Firewall with Advanced Security] window, click [Inbound Rules] on the left side of the window.
- 4. Click [New Rule] on the right side of the window.
- 5. In the [New Inbound Rule Wizard] window, select [Port], and then click [Next].
- 6. Select [TCP] and [Specific local ports], then specify the Web server port number specified during the WebAdmin setup, and then click [Next].
- 7. Select [Allow the connection], and then click [Next].
- 8. Select the profiles for which this rule applies, and then click [Next].
- 9. In [Name], specify the desired name, and then click [Finish].
- 10. In the [Windows Firewall with Advanced Security] window, check if the added rule is enabled under [Inbound Rules] in the center of the window.

## 4.5 Other Settings

This section describes settings that are useful for operations.

## 4.5.1 Error Log Settings

This section explains the settings necessary to monitor errors in applications and operations, and to make discovering the causes easier.

Make error log settings only when instances are created with the initdb command.

When creating instances with WebAdmin, these settings are already made and hence do not need to be set. Furthermore, some parameters are used by WebAdmin, and if changed, may cause WebAdmin to no longer work properly. Refer to "Appendix D Configuring Parameters" for details.

Edit the following parameters in postgresql.conf:

Parameter name	Parameter description	How to enable the settings
event_source	Specify the event source name to be attached to messages, for identifying messages output to the event log when using multiple instances.	<ul> <li>Restart services from the Windows services window.</li> <li>Use the net command or sc command to stop and start services.</li> </ul>

Parameter name	Parameter description	How to enable the settings
logging_collector	Specify "on" to ensure that messages are output by FUJITSU Enterprise Postgres to the server log file. The server log file is created in the pg_log directory in the database cluster.	<ul> <li>Restart services from the Windows services window.</li> <li>Use the net command or sc command to stop and start services.</li> </ul>
log_destination	Specify "stderr, eventlog" to output messages from FUJITSU Enterprise Postgres to the screen and either the system log or the event log.	reload option of the pg_ctl mode
log_line_prefix	Specify information to be added at the start of messages output by an instance. This information is useful for automatic monitoring of messages.	reload option of the pg_ctl mode
	You can output the SQLSTATE value, output time, executing host, application name, and user ID.	
	Refer to "What To Log" in the PostgreSQL Documentation for details.	
	Example: log_line_prefix = '%e: %t [%p]: [%l-1] user = %u,db = %d,remote = %r app = %a '	

# Point

- If you want fewer application errors being output to the eventlog, refer to "When To Log" and "What To Log" in the PostgreSQL Documentation for information on how to reduce the output messages.
- If you want to separate errors output from other software, refer to "Where To Log" in the PostgreSQL Documentation to change the output destination to the server log file rather than the system log.

## 4.5.2 Configuring Automatic Start and Stop of an Instance

You can automatically start or stop an instance when the operating system on the database server is started or stopped.

Use the following procedure to configure automatic start and stop of an instance.

Note that, if an instance is started in a failover operation, the cluster system will control the start or stop, therefore this feature should not be used. Also, when performing database multiplexing, refer to "Enabling Automatic Start and Stop of Mirroring Controller and Multiplexed Instances" in the Cluster Operation Guide (Database Multiplexing).

### When an instance was created with WebAdmin

When an instance is created with WebAdmin, the instance is registered in the Windows service and automatic start and stop is set for the instance.

To change the automatic start and stop setting for an instance, select the service for the applicable instance in the Windows services window, and in [Startup Type], select [Automatic] or [Manual].

### When an instance was created with the initdb command

When the startup type of the service is set to [Manual], change it to [Automatic]. By setting the startup type to [Automatic], the service will start up automatically when the Windows(R) system is started up, and will stop automatically when the Windows(R) system is shut down.



The settings should be performed by a user with administrator privileges.

Use the following procedure to switch the service:

1. Display the [Services] window.

In Windows, select [Administrative Tools], and then click [Services].

2. Switch the startup type

Select the FUJITSU Enterprise Postgres service name, display the [Properties] dialog box, and then switch the startup type from [Manual] to [Automatic].

The above setting can also be changed using the sc config command.

### 4.5.3 Settings when Using the Features Compatible with Oracle Databases

To use the features compatible with Oracle databases, create a new instance and execute the following command for the "postgres" and "template1" databases:

```
CREATE EXTENSION oracle_compatible;
```

Features compatible with Oracle databases are defined as user-defined functions in the "public" schema created by default when database clusters are created, so they can be available for all users without the need for special settings.

For this reason, ensure that "public" (without the double quotation marks) is included in the list of schema search paths specified in the search\_path parameter.

There are also considerations for use the features compatible with Oracle databases. Refer to "Precautions when Using the Features Compatible with Oracle Databases" in the Application Development Guide for details.

## 4.6 Integration with Message-Monitoring Software

To monitor messages output by FUJITSU Enterprise Postgres using software, configure the product to monitor SQLSTATE, instead of the message text - this is because the latter may change when FUJITSU Enterprise Postgres is upgraded.

Configure FUJITSU Enterprise Postgres to output messages in a format that can be read by the message-monitoring software by specifying "%e" in the log\_line\_prefix parameter of postgresql.conf to output the SQLSTATE value.

A setting example is shown below - it outputs the output time, executing host, application name, and user ID, in addition to the SQLSTATE value.

### Example

```
log\_line\_prefix = '%e: %t [%p]: [%l-1] user = %u,db = %d,remote = %r app = %a '
```



Refer to "What To Log" in the PostgreSQL Documentation for information on how to configure the settings.

## 4.7 Setting Up and Removing OSS

This section explains how to set up OSS supported by FUJITSU Enterprise Postgres.

If you want to use OSS supported by FUJITSU Enterprise Postgres, follow the setup procedure.

If you decide not to use the OSS supported by FUJITSU Enterprise Postgres, follow the removing procedure.



- In this section, the applicable database that enables the features of each OSS is described as "postgres".
- Execute CREATE EXTENSION for the "template1" database also, so that each OSS can be used by default when creating a new database.

Refer to "OSS Supported by FUJITSU Enterprise Postgres" in the General Description for information on OSS other than those described below.

### 4.7.1 oracle\_fdw

### 4.7.1.1 Setting Up oracle\_fdw

- Add the path of the OCI library to the environment variable. The available version of the OCI library is 11.2 or later.
   Add the installation path of the OCI library to the PATH environment variable.
- 2. Open a command prompt with administrator privileges and run the following command:

- 3. Restart FUJITSU Enterprise Postgres.
- 4. Execute CREATE EXTENSION for the database that will use this feature. Use the psql command to connect to the "postgres" database.

```
postgres=# CREATE EXTENSION oracle_fdw;
CREATE EXTENSION
```



- If the OCI library is not installed on the server, install it using the Oracle client or Oracle Instant Client. Refer to the relevant Oracle manual for information on the installation procedure.
- If the version of the OCI library is updated, change the path of the OCI library in the PATH environment variable to the updated path.



This feature cannot be used on instances created in WebAdmin. It can only be used via server commands.

### 4.7.1.2 Removing oracle\_fdw

1. Execute DROP EXTENSION for the database that will use this feature. Use the psql command to connect to the "postgres" database.

```
postgres=# DROP EXTENSION oracle_fdw CASCADE;
DROP EXTENSION
```

2. Open a command prompt with administrator privileges and run the following command:

> del "c:\Program Files\Fujitsu\fsepv<x>server64\filesCopiedDuringSetup"



The files copied during setup can be checked below.

> dir /b /s "c:\Program Files\Fujitsu\fsepv<x>server64\OSS\oracle\_fdw"

## 4.7.2 pg\_hint\_plan

## 4.7.2.1 Setting Up pg\_hint\_plan

Set the postgresql.conf file parameters.
 Add "pg\_hint\_plan" to the "shared\_preload\_libraries" parameter.

2. Open a command prompt with administrator privileges and run the following command:

- 3. Restart FUJITSU Enterprise Postgres.
- 4. Run CREATE EXTENSION for the database that uses this feature.

The target database is described as "postgres" here.

Use the psql command to connect to the "postgres" database.

```
postgres=# CREATE EXTENSION pg_hint_plan;
CREATE EXTENSION
```



#### See

Refer to "Optimizer Hints" in the Application Development Guide for details.

## 4.7.2.2 Removing pg\_hint\_plan

1. Execute DROP EXTENSION for the database that will use this feature. Use the psql command to connect to the "postgres" database.

```
postgres=# DROP EXTENSION pg_hint_plan CASCADE;
DROP EXTENSION
```

2. Open a command prompt with administrator privileges and run the following command:

```
> del "c:\Program Files\Fujitsu\fsepv<x>server64\filesCopiedDuringSetup"
```



#### Information

The files copied during setup can be checked below.

> dir /b /s "c:\Program Files\Fujitsu\fsepv<x>server64\OSS\pg\_hint\_plan"

- 3. Set the postgresql.conf file parameters.

  Delete "pg\_hint\_plan" to the shared\_preload\_libraries parameter.
- 4. Restart FUJITSU Enterprise Postgres.

## 4.7.3 pg\_dbms\_stats

#### 4.7.3.1 Setting Up pg\_dbms\_stats

- Set the postgresql.conf file parameter.
   Add "pg\_dbms\_stats" to the "shared\_preload\_libraries" parameter.
- 2. Open a command prompt with administrator privileges and run the following command:

- 3. Restart FUJITSU Enterprise Postgres.
- 4. Run CREATE EXTENSION for the database that will use this feature.

The target database is described as "postgres" here.

Use the psql command to connect to the "postgres" database.

postgres=# CREATE EXTENSION pg\_dbms\_stats; CREATE EXTENSION



Refer to "Optimizer Hints" in the Application Development Guide for details.

### 4.7.3.2 Removing pg\_dbms\_stats

1. Execute DROP EXTENSION for the database that will use this feature. Use the psql command to connect to the "postgres" database.

postgres=# DROP EXTENSION pg\_dbms\_stats CASCADE; DROP EXTENSION

2. Open a command prompt with administrator privileges and run the following command:

> del "c:\Program Files\Fujitsu\fsepv<x>server64\filesCopiedDuringSetup"



The files copied during setup can be checked below.

> dir /b /s "c:\Program Files\Fujitsu\fsepv<x>server64\OSS\pq\_dbms\_stats"

- 3. Set the postgresql.conf file parameters. Delete "pg\_dbms\_stats" to the shared\_preload\_libraries parameter.
- 4. Restart FUJITSU Enterprise Postgres.

# 4.8 Deleting Instances

This section explains how to delete an instance.

- 4.8.1 Using WebAdmin
- 4.8.2 Using Server Commands



- Always use WebAdmin to delete instances that were created or imported using WebAdmin. Because WebAdmin management information cannot be deleted, WebAdmin will determine that the instance is abnormal.

## 4.8.1 Using WebAdmin

This section explains how to delete an instance using WebAdmin.

Use the following procedure to delete an instance.

1. Stop the instance

In the [Instances] tab, select the instance to stop and click .



#### 2. Back up files.

Before deleting the instance, back up any required files under the data storage destination, the backup data storage destination, and the transaction log storage destination.

#### 3. Delete the instance

In the [Instances] tab, select the instance to delete and click iii.



Deleting an instance deletes only the following lowest-level directories. If they are not required, delete them manually.

- Data storage destination
- Backup data storage destination
- Transaction log storage destination (if different from the data storage destination)

## 4.8.2 Using Server Commands

This section explains how to delete an instance using server commands.

Use the following procedure to delete an instance.

#### 1. Stop the instance

Stop the instance in Windows services, or use the pg\_ctl command stop mode.

Use the following procedure to stop a service in Windows services:

a. Display the [Services] window.

In Windows, select [Administrative Tools], and then click [Services].

b. Stop the service

Select the instance name that you wish to stop from the services list, and click [Stop Service]. If you stop a service while applications and commands are running, FUJITSU Enterprise Postgres will force those applications and commands to close and will stop normally.

You can also stop a service by specifying the service name in the net stop command or sc stop command.

#### 2. Back up files

Before deleting the instance, you should back up all necessary files contained in the data storage destination, backup data storage destination, and transaction log storage destination.

#### 3. Delete the instance

Use a command such as rmdir to delete the following directories:

- Data storage destination directory
- Backup data storage destination directory
- Transaction log storage destination directory (if a different directory to the data storage destination directory was specified)

# Chapter 5 Uninstallation

This chapter explains the uninstallation of FUJITSU Enterprise Postgres.

Note that "x SPz" indicates the version and level of the installed product and "<x>" in paths indicates the product version.



- If a [User Account Control] dialog box is displayed at the start of the uninstallation, click [Yes] to continue processing.

If [No] is clicked, permission to continue is denied and an [Error] dialog box will be displayed. To continue the uninstallation, click [Retry] in the [Error] dialog box. To end the operation, click [Cancel].

- If uninstallation is suspended or processing terminates abnormally, the [Program Compatibility Assistant] dialog box may be displayed. Click [This program uninstalled correctly] and continue operation.

## 5.1 Uninstallation in Interactive Mode

The uninstallation procedure is described below.



- Uninstalling removes all files and directories under the installation directory. If you have placed user files in the installation directory, you may need to save them before uninstalling.
- If performing operation with WebAdmin, back up the following folder before uninstallation.
   Instances will not be recognized by WebAdmin even if FUJITSU Enterprise Postgres is reinstalled after uninstallation.
   If performing operation with WebAdmin after reinstalling FUJITSU Enterprise Postgres, replace the following backed up file after installation.

Follow the procedure below to perform the backup.

- 1. Stop the WebAdmin server. Refer to "B.1.3 Stopping the Web Server Feature of WebAdmin" for details.
- 2. Back up the following folder:

webAdminInstallFolder\data\fepwa

By replacing the above folder in the installation folder after installation, the instance will be recognized by WebAdmin, and the recognized instance will be set to automatically start and stop.

To disable the automatic start and stop setting for an instance, select the service for the applicable instance in the Windows services window, and in [Startup Type], select [Manual].

- If using database multiplexing mode, refer to "Uninstalling in Database Multiplexing Mode" in the FUJITSU Enterprise Postgres Cluster Operation Guide (Database Multiplexing) before performing the uninstallation.

......

.....



#### See

Refer to the Installation and Setup Guide for Client when uninstalling the FUJITSU Enterprise Postgres client feature.



If an error occurs while the product is being uninstalled, refer to "Uninstall (middleware) Messages" in the FUJITSU Enterprise Postgres product website, and take the required action.

#### 1. Stop applications and programs

Before starting the uninstallation, stop the following:

- Applications that use the product
- pgAdmin

#### 2. Stop instances

Stop all instances that are using the product to be uninstalled.

Stopping of instances should be performed by the appropriate instance administrator.

When an instance was created with WebAdmin

In the [Instances] tab, select the instance to stop and click .



When an instance was created with the initdb command

Use the following procedure to stop a service:

a. Display the [Services] window

In Windows, select [Administrative Tools], and then click [Services].

b. Stop the service

Select the instance name that you wish to stop from the services list, and click [Stop Service]. If you stop a service while applications and commands are running, FUJITSU Enterprise Postgres will force those applications and commands to close and will stop normally.

You can also stop a service by specifying the service name in the net stop command or sc stop command.

#### 3. Stop WebAdmin

If you are using WebAdmin, stop WebAdmin.

Refer to "B.1.3 Stopping the Web Server Feature of WebAdmin" for details.

#### 4. Unregister Windows services

Perform this step if the instance was created with the initdb command.

Unregister the instance registered in Windows services.

Use the unregister mode of the pg\_ctl command to specify the registered service name and unregister the instance from Windows services.

#### Example

The following is an example showing execution of this command on the registered service name "inst1".

pg\_ctl unregister -N "inst1"



You should unregister services before uninstalling FUJITSU Enterprise Postgres. If you uninstall FUJITSU Enterprise Postgres while services are running, several files will remain after the uninstallation.

If you have carried out the uninstallation without unregistering services beforehand, use the server command sc delete to unregister the services.

This command must be executed by an instance administrator user with administrator privileges. Execute the command from the [Administrator: Command Prompt] window. Right-click [Command Prompt], and then select [Run as administrator] from the menu to display the [Administrator: Command Prompt] window.

#### 5. Delete registrations related to the event log

If you are outputting to the event log, a DLL registration mentioned in "4.2.5 Preparing for Output to the Event Log" has been performed.

To prevent unnecessary issues from occurring, you should delete this registration. Refer to "Server Setup and Operation", "Registering Event Log on Windows" in the PostgreSQL Documentation for details.

The following is an example showing deletion of the DLL registration for a 64-bit product under the default event source name.

```
> regsvr32 /u "c:\Program Files\Fujitsu\fsepv<x>server64\lib\pgevent.dll"
```

#### If using multiple instances

DLL registration is performed so that you can output messages corresponding to the event source name assigned by the user, allowing you to identify messages output to the event log by instance.

Since it is necessary to delete the DLL registration for each instance, delete the DLL registration by event source name.

The following is an example showing deletion of the DLL of a 64-bit product registered under the event source name "Enterprise Postgres inst1".

```
> regsvr32 /u /i: "Enterprise Postgres inst1" "C:\Program Files\Fujitsu\fsepv<x>server64\lib \pgevent.dll"
```

Note that this step is not required if the instance was created with WebAdmin.

#### If installing multiple versions

If the instances you created using this package have been set to output error logs to the event log, use the DLL path name that you took note of previously as explained in "4.2.5 Preparing for Output to the Event Log" to reregister the default event source name.



Ensure to delete DLLs before the uninstallation. If you perform the uninstallation without doing so, you may not be able to delete the DLLs at a later time.

#### 6. Start Uninstall (middleware)

In Windows, click [All Programs] or [All apps], then [Fujitsu], and then [Uninstall (middleware)].

#### 7. Select the software

Select the product to be uninstalled from [Software Name], and then click [Remove].

#### 8. Start the uninstallation

Click [Uninstall].

#### 9. Finish the uninstallation

The uninstallation completion window will be displayed. Click [Finish].

The installation folder may remain after uninstallation. If it is not required, delete it.

#### 10. Stop Uninstall (middleware)

In Uninstall (middleware), click [Close].

## 5.2 Uninstallation in Silent Mode

The uninstallation procedure is described below.



- Uninstalling removes all files and directories under the installation directory. If you have placed user files in the installation directory, you may need to save them before uninstalling.

- If performing operation with WebAdmin, back up the following folder before uninstallation. Instances will not be recognized by WebAdmin even if FUJITSU Enterprise Postgres is reinstalled after uninstallation. If performing operation with WebAdmin after reinstalling FUJITSU Enterprise Postgres, replace the following backed up file after installation.

Follow the procedure below to perform the backup.

- 1. Stop the WebAdmin server. Refer to "B.1.3 Stopping the Web Server Feature of WebAdmin" for details.
- 2. Back up the following folder:

webAdminInstallFolder\data\fepwa

By replacing the above folder in the installation folder after installation, the instance will be recognized by WebAdmin, and the recognized instance will be set to automatically start and stop.

To disable the automatic start and stop setting for an instance, select the service for the applicable instance in the Windows services window, and in [Startup Type], select [Manual].

- If using database multiplexing mode, refer to "Uninstalling in Database Multiplexing Mode" in the FUJITSU Enterprise Postgres Cluster Operation Guide (Database Multiplexing) before performing the uninstallation.



#### See

- Refer to the Installation and Setup Guide for Client when uninstalling the FUJITSU Enterprise Postgres client feature.
- Refer to the FUJITSU Enterprise Postgres product website for information on uninstallation in silent mode, such as the error messages.

.....

#### 1. Stop applications and programs

Before starting the uninstallation, stop the following:

- Applications that use the product
- pgAdmin

#### 2. Stop all instances

Stop all instances that are using the product to be uninstalled.

When an instance was created with WebAdmin

In the [Instances] tab, select the instance to stop and click .



When an instance was created with the initdb command

Use the following procedure to stop a service:

- a. Display the [Services] window.
  - In Windows, select [Administrative Tools], and then click [Services].
- b. Stop the service

Select the instance name that you wish to stop from the services list, and click [Stop Service]. If you stop a service while applications and commands are running, FUJITSU Enterprise Postgres will force those applications and commands to close and will stop normally.

You can also stop a service by specifying the service name in the net stop command or sc stop command.

#### 3. Stop WebAdmin

If you are using WebAdmin, stop WebAdmin.

Refer to "B.1.3 Stopping the Web Server Feature of WebAdmin" for details.

#### 4. Unregister Windows services

Perform this step if the instance was created with the initdb command.

Unregister the instance registered in Windows services.

Use the unregister mode of the pg\_ctl command to specify the registered service name and unregister the instance from Windows services.

#### Example

The following is an example showing execution of this command for the registered service name "inst1".

> pg\_ctl unregister -N "instl"



You should unregister services before uninstalling FUJITSU Enterprise Postgres. If you uninstall FUJITSU Enterprise Postgres while services are running, several files will remain after the uninstallation.

If you have carried out the uninstallation without unregistering services beforehand, use the server command sc delete to unregister the services.

This command must be executed by an instance administrator user with administrator privileges. Execute the command from the [Administrator: Command Prompt] window. Right-click [Command Prompt], and then select [Run as administrator] from the menu to display the [Administrator: Command Prompt] window.

#### 5. Delete registrations related to the event log

If you are outputting to the event log, a DLL registration mentioned in "4.2.5 Preparing for Output to the Event Log" has been performed.

To prevent unnecessary issues from occurring, you should delete this registration. Refer to "Server Setup and Operation", "Registering Event Log on Windows" in the PostgreSQL Documentation for details.

The following is an example showing deletion of the DLL registration for a 64-bit product under the default event source name.

> regsvr32 /u "c:\Program Files\Fujitsu\fsepv<x>server64\lib\pgevent.dll"

#### If using multiple instances

DLL registration is performed so that you can output messages corresponding to the event source name assigned by the user, allowing you to identify messages output to the event log by instance.

Since it is necessary to delete the DLL registration for each instance, delete the DLL registration by event source name.

The following is an example showing deletion of the DLL of a 64-bit product registered under the event source name "Enterprise Postgres inst1".

Note that this step is not required if the instance was created with WebAdmin.

#### If installing multiple versions

If the instances you created using this package have been set to output error logs to the event log, use the DLL path name that you took note of previously as explained in "4.2.5 Preparing for Output to the Event Log" to reregister the default event source name.



Ensure to delete DLLs before the uninstallation. If you perform the uninstallation without doing so, you may not be able to delete the DLLs at a later time.

#### 6. Start the command prompt

In Windows, right-click [Command Prompt] and then select [Run as administrator].

#### 7. Run the uninstaller

Execute the command below.

The installation folder may remain after uninstallation. If it is not required, delete it.

#### Example

X:> installFolder\suninst.bat

X: Drive on which the product is installed

# Appendix A Recommended WebAdmin Environments

This appendix describes the recommended WebAdmin environment. The following explanation is based on the assumption that Internet Explorer 11 or later is used unless otherwise stated.

## A.1 Recommended Browser Settings

- Use a display resolution of 1280 x 768 or higher, and 256 colors or more.
- Select [View] >> [Text size] >> [Medium].
- Select [View] >> [Zoom] >> [100%].
- Click [Tools] >> [Internet options] >> [General] >> [Fonts], and then:
  - Set [Webpage font] to [Times New Roman].
  - Set [Plain text font] to [Courier New].

## A.2 How to Set Up the Pop-up Blocker

If the Pop-up Blocker is enabled, use the procedure below to configure settings to allow pop-ups from the server where FUJITSU Enterprise Postgres is installed.

- 1. Click [Tools] >> [Internet options], and then select the [Privacy] tab.

  If [Turn on Pop-up Blocker] is not selected, the Pop-up Blocker feature will not operate, and therefore steps below are not required.
- 2. Click [Settings].
- 3. In the [Pop-up Blocker Settings] window, enter in the [Address of website to allow] the URL of the server where FUJITSU Enterprise Postgres is installed, and then click [Add].
- 4. Click [Close].
- 5. In the [Internet Options] window, click [OK].

# Appendix B Setting Up and Removing WebAdmin

This appendix describes how to set up and remove WebAdmin.

Note that "<x>" in paths indicates the product version.

## **B.1 Setting Up WebAdmin**

This section explains how to set up WebAdmin.

## **B.1.1 Setting Up WebAdmin**

Follow the procedure below to set up WebAdmin.

1. Log in

Log in as a user that belongs to the Administrators group.

2. Display the setup window

In Windows, click [All Programs] or [All apps], then [Product name], and then [WebAdmin Setup].



If the same [User Account Control] dialog box as that shown below is displayed, click [Yes] to continue processing.

#### 3. Specify the port number

Specify the following port numbers to be used in WebAdmin.

Refer to the services file. Only change to a different port number if the same port number is being used by another service.

Make a note of the Web server port number, because it will be required for starting the WebAdmin window.

Item	Value (recommended value)
Web server port number	27515
WebAdmin internal port number	27516
WebAdmin automatic start	Selected

#### Web server port number

Specify an unused port number in the following range for the port number used for communication between the Web browser and Web server:

- 1024 to 49151

The Web server port number is registered as a port number of the following service name in the services file. "<*x*>" indicates the product version.

fep\_xSPz\_WA\_64\_WebAdmin\_Port1

#### WebAdmin internal port number

Specify an unused port number in the following range for the port number used for communication between the Web server and WebAdmin runtime environment:

- 1024 to 49151

The WebAdmin internal port number is registered as a port number of the following service name in the services file. "< x >" indicates the product version.

fep\_xSPz\_WA\_64\_WebAdmin\_Port2

#### WebAdmin automatic start

Select whether or not to start WebAdmin when the machine is started.



- Make a note of the Web server port number for use in the Windows firewall settings.
- Unused port numbers

In the operating system and other products, regardless of the information in the service file, unused port numbers may be automatically numbered and then used, or port numbers specified in environment files within products may also be used. Check the port numbers used by the OS and other products, and ensure that these are not duplicated.

Access restrictions

Prevent unauthorized access and maintain security by using a firewall product, or the packet filtering feature of a router device, to restrict access to the server IP address and the various specified port numbers.

- Port access permissions

If a port is blocked (access permissions have not been granted) by a firewall, enable use of the port by granting access. Refer to the vendor document for information on how to grant port access permissions.

Consider the security risks carefully when opening ports.

- Changing port numbers

When using WebAdmin in multiserver mode, it is recommended not to change WebAdmin ports after creating instances. Otherwise, the created instances may not be accessible through WebAdmin after the port is changed.

#### 4. Prepare for setup

Click [OK] in the setup window, and after completing the WebAdmin setup, refer to "4.2 Preparations for Setup" and perform the required preparations for setting up FUJITSU Enterprise Postgres if using WebAdmin for operation.

## **B.1.2 Starting the Web Server Feature of WebAdmin**

Follow the procedure below to start the Web server feature of WebAdmin.

1. Display the Services window

In Windows, select [Administrative Tools], and then click [Services].

2. Start the service

Select the displayed name "FUJITSU Enterprise Postgres WebAdmin version", and then click [Start Service].

You can also start the service by specifying the service name of the Web server feature of WebAdmin in the net start command or sc start command.

## **B.1.3 Stopping the Web Server Feature of WebAdmin**

Follow the procedure below to stop the Web server feature of WebAdmin.

1. Display the Services window

In Windows, select [Administrative Tools], and then click [Services].

2. Stop the service

Select the displayed name "FUJITSU Enterprise Postgres WebAdmin version", and then click [Stop Service].

You can also stop the service by specifying the service name of the Web server feature of WebAdmin in the net stop command or sc stop command.



- For efficient operation of WebAdmin, it is recommended to stop the Web server feature only during a scheduled maintenance period.

- When WebAdmin is used to create and manage instances in a multiserver configuration, the Web server feature must be started and running on all servers at the same time.

# B.2 Removing WebAdmin

This section explains how to remove WebAdmin.

This removal procedure stops WebAdmin and ensures that it no longer starts automatically when the machine is restarted.

To remove the setup, execute the command shown below.

#### Example

When WebAdmin is installed in "C:\Program Files\Fujitsu\fsepv<*x*>webadmin":

> C

> cd C:\Program Files\Fujitsu\fsepv<x>webadmin\sbin

> WebAdminSetup --delete



- The removal of the WebAdmin setup must be performed by a user with administrator privileges (a user ID that belongs to the Administrators group).
- Commands that require administrator privileges must be executed from the [Administrator: Command Prompt] window. Right-click [Command Prompt], and then select [Run as administrator] from the menu to display the [Administrator: Command Prompt] window.

## **B.3 Using an External Repository for WebAdmin**

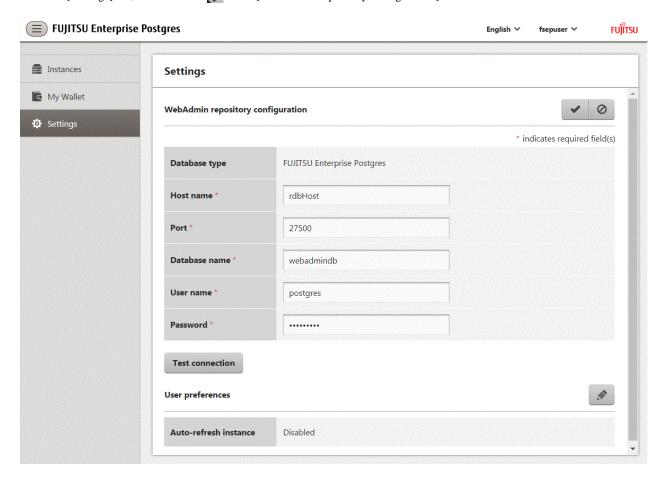
WebAdmin can be configured to use an external database, where it can store the various metadata information it uses. WebAdmin will use this database as a repository to store the information it uses to manage all the created instances. This can be a FUJITSU Enterprise Postgres database or an Open Source PostgreSQL V9.2 or later database.

Using an external database as a WebAdmin repository provides you with more flexibility in managing WebAdmin. This repository can be managed, backed up and restored as needed using pgAdmin or command line tools, allowing users to have greater flexibility and control.

Follow the procedure below to set up the repository.

1. Start WebAdmin, and log in to the database server.

2. Click the [Settings] tab, and then click in the [WebAdmin repository configuration] section.



#### Enter the following items:

- [Host name]: Host name of the database server
- [Port]: Port number of the database server
- [Database name]: Name of the database
- [User name]: User name to access the database
- [Password]: Password of the database user



- Database type
  - It is recommended to use a FUJITSU Enterprise Postgres database as a repository. A compatible PostgreSQL database can also be used as an alternative.
- It is recommended to click [Test connection] to ensure that the details entered are valid and WebAdmin is able to connect to the target database.
- Host name, Database name, User name, Password should not contain hazardous characters. Refer to "Appendix C WebAdmin Disallow User Inputs Containing Hazardous Characters".
- 3. Click to register the repository details.



- Once the repository is set up, it can be changed any number of times by the user logged into WebAdmin. When a repository is changed:
  - It is recommended to preload the backup into this database.
  - If the data is not preloaded, WebAdmin will create a new repository.
- The database repository can be set up even after WebAdmin was already used to create instances. In that scenario, the instances already created are retained and can continue to be operated on.
- If the instance used as a repository is stopped, WebAdmin will be unusable. For this reason, it is recommended to be familiar with starting an instance from the command line. If the instance is stopped for any reason, start it from the command line and WebAdmin will be usable again.

# **B.4** Using the WebAdmin Auto-Refresh Feature

The WebAdmin auto-refresh feature automatically refreshes the operating status of all instances in the Instance list at the specified interval. It also refreshes the details of the selected instance.

Follow the procedure below to configure the auto-refresh options.

- 1. Click the [Settings] tab, and then click in the [User preferences] section.
- 2. Enter the following items:
  - [Auto-refresh instance]: To use the auto-refresh feature, select "Enabled". The default is "Disabled".
  - [Refresh interval (seconds)]: Number of seconds between each refresh. This is a countdown timer, which is reset every time the instance status is refreshed by any operation. Specify a value from 30 to 3600 (seconds). The default is 30.

3. Click to save the auto-refresh settings.



- Auto-refresh will run only if the [Instances] page is displayed and no user-initiated operation is in progress.
- A text indicator, which is independent of auto-refresh, is displayed at the top of the Instance list. It is dynamically updated to display when the page was last refreshed.

# Appendix C WebAdmin Disallow User Inputs Containing Hazardous Characters

WebAdmin considers the following as hazardous characters, which are not allowed in user inputs

| (pipe sign)
& (ampersand sign)
; (semicolon sign)
\$ (dollar sign)
% (percent sign)
@ (at sign)
' (single apostrophe)
" (quotation mark)
\' (backslash-escaped apostrophe)
\" (backslash-escaped quotation mark)
<> (triangular parenthesis)
() (parenthesis)
+ (plus sign) [14] CR (Carriage return, ASCII 0x0d)
LF (Line feed, ASCII 0x0a)
, (comma sign)

\ (backslash)

# Appendix D Configuring Parameters

WebAdmin operates and manages databases according to the contents of the following configuration files:

- postgresql.conf

Contains various items of information that define the operating environment of FUJITSU Enterprise Postgres.

- pg\_hba.conf

Contains various items of information related to client authentication.

These configuration files are deployed to a data storage destination. Data is written to them when the instance is created by WebAdmin and when settings are changed, and data is read from them when the instance is started and when information from the [Setting] menu is displayed.

Direct editing of each configuration file is possible with a text editor.



#### See

Refer to "Server Configuration" and "Client Authentication" in "Server Administration" in the PostgreSQL Documentation for information on the parameters.

......



WebAdmin checks for port number and backup storage path anomalies when various operations are performed. An anomaly occurs when the value of [Port number] and/or [Backup storage path] in WebAdmin is different from the value of the corresponding parameter in postgresql.conf. Refer to "Anomaly Detection and Resolution" in the Operation Guide for details.

#### postgresql.conf

Parameters that can be changed in WebAdmin

The postgresql.conf parameters that can be changed in WebAdmin are shown below:

Section	WebAdmin item	postgresql.conf file parameter
Instance Configuration		
Character encoding	Character set	client_encoding
	Message locale	lc_messages
Communication	Max connections	max_connections
SQL options	Transform NULL format	transform_null_equals
	Date output format	DateStyle (*1)
	Interval output format	IntervalStyle
	Number of digits for floating values	extra_float_digits
	Transaction isolation levels	default_transaction_isolation
	Currency format	lc_monetary
	Date and time format	lc_time
	Numerical value format	lc_numeric
Memory	Sort memory (KB)	work_mem
	Shared buffers (KB)	shared_buffers
Streaming replication	WAL level	wal_level

Section	WebAdmin item	postgresql.conf file parameter
	Maximum WAL senders	max_wal_senders
	WAL keep segments	wal_keep_segments
	Hot standby	hot_standby
	Synchronous standby names	synchronous_standby_names
	WAL receiver timeout (ms)	wal_receiver_timeout
Edit instance		
	Instance name	n/a
	Instance port	port
	Backup storage path	backup_destination

<sup>\*1:</sup> If you specify "Postgres" as the output format, dates will be output in the "12-17-1997" format, not the "Wed Dec 17 1997" format used in the PostgreSQL Documentation.



- Calculate the maximum number of connections using the formula below:

```
maximumNumberOfConnections = maximumNumberOfConnectionsFromApplications + 3 (*1)
```

\*1: 3 is the default number of connections required by the system.

Calculate the maximum number of connections using the following formula when changing superuser\_reserved\_connections (connections reserved for use by the superuser) in postgresql.conf.

 ${\it maximumNumberOfConnections} = {\it maximumNumberOfConnectionsFromApplications} + {\it superuser\_reserved\_connections}$ 

- Also check if the memory used exceeds the memory installed (refer to "Parameters automatically set by WebAdmin according to the amount of memory").

#### Parameters set by WebAdmin

The following postgresql.conf parameters are set by WebAdmin during instance startup (they will be ignored even if specified in postgresql.conf):

Parameter	Value
event_source (*1)	'fsep_version_userName_instanceNamePortNumber'
listen_addresses	*
log_destination	'stderr,eventlog'
logging_collector	on
log_line_prefix	'%e: %t [%p]: [%l-1] user = %u,db = %d,remote = %r app = %a '
log_filename (*2) (*3)	'logfile-%a.log'
log_truncate_on_rotation	on
log_rotation_age	1d

<sup>\*1:</sup> PortNumber is the port number of the database server specified when creating the instance.

<sup>\*2:</sup> The server logs are split into files based on the day of the week, and are rotated after each week.

<sup>\*3:</sup> If the date changes while the instance is stopped, old logs are not deleted and continue to exist. Manually delete old logs that are no longer required to release disk space.

Parameters automatically set by WebAdmin according to the amount of memory

The postgresql.conf parameters automatically set according to the amount of installed memory, during the creation of instances by WebAdmin, are shown below:

Parameter	Value
shared_buffers	30% of the machine's installed memory
work_mem	30% of the machine's installed memory / max_connections / 2
effective_cache_size	75% of the machine's installed memory
maintenance_work_mem	10% of the machine's installed memory / (1 + autovacuum_max_workers) (*1)

<sup>\*1:</sup> The value will be capped at 2097151 KB.

When determining the values to be configured in the above parameters, you must take into account any anticipated increases in access volume or effects on performance during business operations, such as the number of applications and commands that will access the instance, and the content of processes. Also, note that in addition to FUJITSU Enterprise Postgres, other software may be running on the actual database server. You will need to determine the degree of priority for the database and other software, as well as the memory allocation size.

WebAdmin automatically configures complex parameter settings such as those mentioned above, based on the size of the internal memory of the machine. This enables maximum leverage of the machine memory to facilitate resistance against fluctuations during business operations.

Accordingly, the effects of the above-mentioned factors must be estimated and taken into account when determining and configuring parameter values, so that memory resources can be effectively allocated among other software or instances, and so that adverse effects can be mutually avoided. Refer to "Memory" in "Resource Consumption", and "Planner Cost Constants" in "Query Planning", under "Server Administration" in the PostgreSQL Documentation for information on parameter values and required considerations.

Parameter values can be modified using the WebAdmin [Setting] menu, or edited directly using a text editor.

If adding an instance, determine the parameter values, including for existing instances, and make changes accordingly.



- Do not directly edit the following postgresql.conf parameters with a text editor, otherwise WebAdmin may not work properly if you make a mistake):
  - archive\_mode
  - archive\_command
  - wal level
  - core\_directory
  - wal\_sync\_method
  - log\_line\_prefix
  - log\_destination
  - logging\_collector
  - log\_directory
  - log\_file\_mode
  - log\_filename
  - log\_truncate\_on\_rotation
  - log\_rotation\_age

- event\_source
- You must take care with the following parameter:
  - superuser\_reserved\_connections

Set it to a number that includes the 3 connections required in WebAdmin (the default is 3).

#### pg\_hba.conf

Refer to "Client Authentication" in "Server Administration" in the PostgreSQL Documentation for information on content that can be configured in pg\_hba.conf.



- Use the following client authentication settings to allow the instance administrator to connect to the database using WebAdmin:

......

- The connection type: "host"
- The IP address is a loopback address ("127.0.0.1/32")
- If you specify an item or value that cannot be set by WebAdmin when editing the pg\_hba.conf file with a text editor, it will not be possible to reference that line from WebAdmin.

# Appendix E Estimating Database Disk Space Requirements

This appendix describes how to estimate database disk space requirements.

# E.1 Estimating Table Size Requirements

The following tables provide the formulas for estimating table size requirements.

Table E.1 Estimation formula when the record length is 2032 bytes or less

Item	Estimation formula (bytes)
(1) Record length	27(*1) + NULL map + OID + column data
	NULL map: Number of columns / 8 (*2) OID: 4
	Column data: Sum of column lengths
	*1: Record header section
	*2: Round the result up to the next integer.
	<ul> <li>Because the column data is placed in boundaries of 8 bytes, you need to make an adjustment so that the sum of the record header section, NULL map and OID is a multiple of 8.</li> <li>For example, if the calculated length is 27 + 1 / 8 (rounded up) + 0 = 28 bytes, add</li> </ul>
	4 to make the length 32 bytes.
	<ul> <li>Because the data of each column is placed in boundaries of the defined data type, take the boundary of each data type into account for the length of the column data.</li> <li>For example, the length of the column data in the table below will not be the sum of the data types, which is 37 bytes, but will instead be 64 bytes following boundary adjustment.</li> </ul>
	Definition: create table tb1(c1 char(1), c2 long, c3 int, c4 box) Estimation: CHAR type 1 byte + boundary adjustment of 7 bytes for LONG type 8 bytes + LONG type 8 bytes + INT type 4 bytes + boundary adjustment of 12 bytes for BOX type 32 bytes + BOX type 32 bytes = 64 bytes
	- Because each record is placed in boundaries of 8 bytes, you need to make an adjustment so that the length of the column data is a multiple of 8.
	<ul> <li>If the calculated record length exceeds 2,032 bytes, the variable length data in the record might be compressed automatically. If so, use the estimation formulas in "Table E.2 Estimation formula when the record length exceeds 2032 bytes" to estimate the table size.</li> </ul>
(2) Page size requirement	8192 (*1) × fillfactor (*2) - 24 (*3)
	*1: Page length (8192)
	*2: Value of the fillfactor specified in the table definitions (if omitted, 100%)
	*3: Page header (24)
	- The calculated (2) page size requirement will be rounded down to the nearest integer.
(3) Number of records per page	(2) Page size requirement / ((1) record length + 4 (*1))
	*1: Pointer length (4)
	- The result will be rounded down to the nearest integer.

Item	Estimation formula (bytes)
(4) Number of pages required	Total number of records / (3) number of records per page
for storing records	- The result will be rounded up to the next integer.
(5) Amount of space	(4) Number of pages required for storing records x page length x safety factor (*1)
	*1: Specify 2.0 or higher.
	<ul> <li>This is the safety factor assumed if vacuuming is performed for garbage collection in tables and indexes.</li> </ul>

Table E.2 Estimation formula when the record length exceeds 2032 bytes

Item	Estimation formula (bytes)
(5) Amount of space	Total number of records x (1) record length x safety factor (*1)
	*1: Specify 2.0 or higher.
	<ul> <li>This is the safety factor assumed if vacuuming is performed for garbage collection in tables and indexes.</li> </ul>

# **E.2 Estimating Index Size Requirements**

This section provides the formulas for estimating index size requirements.

FUJITSU Enterprise Postgres provides six index types: B-tree, Hash, GiST, GIN, SP-GiST, and VCI. If you do not specify the index type in the CREATE INDEX statement, a B-tree index is generated.

The following describes how to estimate a B-tree index. Refer to "E.7 Estimating VCI Disk Space Requirements" for information on how to estimate VCI.

A B-tree index is saved as a fixed-size page of 8 KB. The page types are meta, root, leaf, internal, deleted, and empty. Since leaf pages usually account for the highest proportion of space required, you need to calculate the requirements for these only.

Table E.3 Estimation formula when the key data length is 512 bytes or less

Item	Estimation formula (bytes)
(1) Entry length	8 (*1) + key data length (*2)
	*1: Entry header
	*2: The key data length depends on its data type (refer to "E.3 Sizes of Data Types" for details).
	Because each entry is placed in boundaries of 8 bytes, you need to make an adjustment so that the length of the key data is a multiple of 8.
	For example, if the calculated length is 28 bytes, add 4 to make the length 32 bytes.
	- If the key data length exceeds 512 bytes, key data may be automatically compressed.
	In this case, use the estimation formula given in "Table E.4 Estimation formula when the key data length exceeds 512 bytes" to estimate the key data length.
(2) Page size requirement	8192 (*1) × fillfactor (*2) - 24 (*3) - 16 (*4)
	*1: Page length (8192)
	*2: Value of the fillfactor specified in the index definitions (if omitted, 90%)
	In the case of indexes of primary key constraints and unique constraints,
	the value of the fill factor specified for each constraint in the table definitions (if omitted, 90%)
	*3: Page header (24)
	*4: Special data (16)

Item	Estimation formula (bytes)
	- The calculated (2) page size requirement will be rounded down to the nearest integer.
(3) Number of entries per page	(2) Page size requirement / ((1) entry length + 4 (*1))
	*1: Pointer length
	- Result of (3) number of entries per page will be rounded down to the nearest integer.
(4) Number of pages required	Total number of records / (3) number of entries per page
for storing indexes	- Result of (4) number of pages required for storing indexes will be rounded up to the nearest integer.
(5) Space requirement	(4) Number of pages required for storing indexes x 8192 (*1) / usage rate (*2)
	*1: Page length
	*2: Specify 0.7 or lower.

Table E.4 Estimation formula when the key data length exceeds 512 bytes

Item	Estimation formula (bytes)
(5) Space requirement	Total number of records x key data length x compression ratio (*1) / usage rate (*2)
	*1: The compression ratio depends on the data value, so specify 1.
	*2: Specify 0.7 or lower as the usage rate.

# E.3 Sizes of Data Types

This section lists the sizes of the data types.

# **E.3.1 Sizes of Fixed-Length Data Types**

The following table lists the sizes of fixed-length data types.

Data type	Size (bytes)
SMALLINT (INT2)	2
INTEGER (INT4)	4
BIGINT (INT8)	8
REAL	4
DOUBLE PRECISION	8
SERIAL (SERIAL4)	4
BIGSERIAL (SERIAL8)	8
MONEY	8
FLOAT	8
FLOAT (1-24)	4
FLOAT (25-53)	8
TIMESTAMP WITHOUT TIME ZONE	8
TIMESTAMP WITH TIME ZONE	8
DATE	4
TIME WITHOUT TIME ZONE	8
TIME WITH TIME ZONE	12

Data type	Size (bytes)
INTERVAL	12
BOOLEAN	1
CIDR	IPv4: 7 IPv6: 19
INET	IPv4: 7 IPv6: 19
MACADDR	6
MACADDR8	8
POINT	16
LINE	32
LSEG	32
BOX	32
CIRCLE	24

# **E.3.2 Sizes of Variable-Length Data Types**

The following table lists the sizes of variable-length data types.

Data type	Data type Size (bytes) Remarks	
path	Length of size portion + 12 + 16 x number of vertices	1) When carrying out division, round to the next integer.
polygon	Length of size portion + 36 + 16 x number of vertices	2) If the real data length is less than 127, then the length of the size portion is 1 byte, otherwise it is 4 bytes.
decimal	Length of size portion + 2 + (integer	3) The number of bytes per character depends on the
numeric	precision / 4 + decimal precision / 4) x 2	character set (refer to "E.3.4 Number of Bytes per Character" for details).
bytea	Length of size portion + real data length	Character for details).
character varying(n), varchar(n)	Length of size portion + number of characters x number of bytes per character	
character(n), char(n)	Length of size portion + n x number of bytes per character	
text	Length of size portion + number of characters x number of bytes per character	

# **E.3.3** Sizes of Array Data Types

The following table lists the sizes of array data types.

Data type	Size (bytes)	Remarks
Array	Length of size portion + 12 + 8 x number of dimensions + data size of each item	If the real data length is less than 127, then the length of the size portion is 1 byte, otherwise it is 4 bytes.  - Example of estimation when array data is "ARRAY[[1,2,3], [1,2,3]]"
		Number of dimensions: 2
		INTEGER data size: 4  Total size = $1+12+8x2+6x4 = 53$

## E.3.4 Number of Bytes per Character

The following table lists the number of bytes per character.

The given values relate to the common character sets EUC-JP and UTF8.

Character type	Character set	Number of bytes per character
ASCII	EUC_JP	1
Halfwidth katakana	EUC_JP	2
JIS X 0208 kanji characters	EUC_JP	2
JIS X 0212 kanji characters	EUC_JP	3
ASCII	UTF8	1
Halfwidth katakana	UTF8	3
JIS X 0208 kanji characters	UTF8	3
JIS X 0212 kanji characters	UTF8	3

## **E.4 Estimating Transaction Log Space Requirements**

This section provides the formula for estimating transaction log space requirements.

```
Transaction log space requirements = max_wal_size
```

However, if the update volume is extremely high (for example, due to a large data load and batch processing), disk writing at a checkpoint may not be able to keep up with the load, and a higher number of transaction logs than indicated here may temporarily be accumulated.

## E.5 Estimating Archive Log Space Requirements

This section explains how to estimate archive log space requirements.

The archive log is an archive of the transaction logs from the time of a previous backup to the present, so it fluctuates depending on the backup period and the content of update transactions.

The longer the backup period and the more update transactions, the greater the space required for the archive log.

Therefore, measure the actual archive log space by using a test environment to simulate backup scheduling and database update in a real operating environment.

## E.6 Estimating Backup Disk Space Requirements

This section provides the formula for estimating backup disk space requirements.

Backup disk space requirements = size of the database cluster x 2 + transaction log space requirements + archive log space requirements



If the pgx\_dmpall command performs a backup using a user exit, the backup disk size differs according to the database resources targeted for backup and the copy method.

# E.7 Estimating VCI Disk Space Requirements

This section provides the formula for estimating VCI disk space requirements.

Disk space = (number of rows in tables) x (number of bytes per row) x (compression ratio) + (WOS size)

#### Number of bytes per row

```
Number of bytes per row = (19 + (number of columns specified in CREATE INDEX) / 8 + (number of bytes per single column value)) x 1.1
```

Note: Round up the result to the nearest integer.

#### Compression ratio

Specify a value between 0 and 1. Since compression ratio depends on the data being compressed, use actual data or test data that simulates it, then compare the value with the estimation result. As a guide, the compression ratio measured with the Fujitsu sample data is shown below:

- Data with high degree of randomness (difficult to compress): Up to approximately 0.9 times.
- Data with high degree of similarity (easy to compress): Up to approximately 0.5 times.

#### WOS size

```
WOS size = (number of WOS rows) / 185 x 8096
```

One row is added to the number of WOS rows for each INSERT and DELETE, and two rows are added for UPDATE. On the other hand, the number decreases to 520,000 rows or less during conversion to ROS performed by the ROS control daemon.



VCI does not support retrieval of disk space usage using the database object size function pg\_indexes\_size. To find out the actual total VCI disk space, check the disk space of the storage directory using an OS command or other method.

# Appendix F Estimating Memory Requirements

This appendix explains how to estimate the memory.

## F.1 FUJITSU Enterprise Postgres Memory Requirements

This section describes the formulas for estimating FUJITSU Enterprise Postgres memory requirements.

Use the following formula to obtain a rough estimate of memory required for FUJITSU Enterprise Postgres:

fujitsuEnterprisePostgresRequiredMemory = sharedMemoryAmount + localMemoryAmount

#### **Shared memory amount**

Refer to "Shared Memory and Semaphores" under "Server Administration" in the PostgreSQL Documentation for information on shared memory. If you enable the Global Meta Cache feature, you must also add the value of pgx\_global\_metacache. Refer to "Parameters" in the Operation Guide for the setting values.

However, note that if instances have been created using WebAdmin, the parameters below will be configured automatically when the instances are created. Take this into account when calculating the shared memory size.

Parameter name	Set value	
shared_buffers	30 percent of the internal memory of the machine.	
max_connections	100	
max_prepared_transactions	100	

#### Local memory amount

localMemoryAmount = processStackArea

- + memoryUsedInDbSessionsThatUseTempTables
- $+ \ \textit{memoryUsedInDbSessionsThatPerformSortAndHashTableOperations}$
- + memoryUsedInMaintenanceOperations
- + baseMemoryUsedInEachProcess
- $+ \ {\it memoryUsedPreparingForDataAccess}$

#### Process stack area

processStackArea

= max\_stack\_depth x (max\_connections + autovacuum\_max\_workers + 9)

This formula evaluates to the maximum value.

Actually it is used according to the growth of the stack.

In the formula above, 9 is the number of processes that perform roles specific to servers.

Memory used in database sessions that use temporary tables

memoryUsedInDbSessionsThatUseTempTables
= temp\_buffers x max\_connections

This formula evaluates to the maximum value.

Memory is gradually used as temporary buffers are used, and is released when the session ends.

Memory used in database sessions that perform sort and hash table operations

This formula evaluates to the maximum value.

Memory is gradually used as operations such as sort are performed, and is released when the query ends.

#### Memory used in maintenance operations

```
memoryUsedInMaintenanceOperations
= maintenance_work_mem x (numOfSessionsPerformingMaintenance + autovacuum_max_workers)
```

Note that 'maintenance operations' are operations such as VACUUM, CREATE INDEX, and ALTER TABLE ADD FOREIGN KEY.

#### Base memory used in each process

Use the result of the following formula for memory consumed per process. This formula evaluates to the memory used when server processes are running.

In the formula above, 9 is the number of processes that perform roles specific to servers.

The amount of memory consumed per process is determined by the number of tables, indexes, and all columns of all tables that the process accesses. If your system has about 100 tables, you can estimate it to be 3 MB, but otherwise use the following estimate:

There are variable length information. This value takes that into account.

#### Memory used preparing for data access

This formula evaluates to the memory required to access the database cache in the shared memory.

In the formula above, among the processes that perform roles specific to servers, 4 is the number of processes that access the database.

## F.2 Database Multiplexing Memory Requirements

This section describes the formula for estimating database multiplexing memory requirements for the database server.

Use the following formula to obtain a rough estimate of memory required for database multiplexing:

```
Memory usage of the database multiplexing feature for the database server

= Peak memory usage of the Mirroring Controller processes
+ Peak memory usage of the Mirroring Controller commands

Peak memory usage of the Mirroring Controller processes=150 MB

Peak memory usage of the Mirroring Controller commands=50 MB x Number of commands executed simultaneously
```

## F.3 VCI Memory Requirements

This section describes the formula for estimating VCI memory requirements.

Use the following formula to obtain a rough estimate of memory requirements:

```
memUsedByVci = memForData + memForEachProcess
```

<sup>\*1)</sup> Safety Factor (1.5)

#### Memory required to store data in memory

Secure the space estimated using the formula below on the stable buffer (part of shared\_buffers).

```
memForData = (numOfRowsInTables) x (numOfBytesPerRow) + (wosSize)
```

#### Number of bytes per row

```
numOfBytesPerRow =
  (19 + (numOfColsInCreateIndexStatement) / 8 + (numOfBytesPerSingleColValue)) x 1.1
```

Note: Round up the result to the nearest integer.

#### WOS size

```
wosSize = (numOfWosRows) / 185 x 8096
```

One row is added to the number of WOS rows for each INSERT and DELETE, and two rows are added for UPDATE. On the other hand, the number decreases to 520,000 rows or less during conversion to ROS performed by the ROS control daemon.

#### Memory required for each process

```
memForEachProcess =
  memUsedPerScanning +
  memUsedForVciMaintenace +
  memUsedByCreateIndexStatement
```

#### Memory used per scanning

- Parallel scan

```
memUsedPerScanning =
  vci.shared_work_mem + (numOfParallelWorkers + 1) x vci.maintenance_work_mem
```

Note: The number of parallel workers used by VCI simultaneously in the entire instance is equal to or less than vci.max\_parallel\_degree.

- Non-parallel scan

```
memUsedPerScanning = vci.max_local_ros + vci.maintenance_work_mem
```



- vci.shared\_work\_mem, and vci.max\_local\_ros are used to create local ROS. If local ROS exceeds these sizes, execute a query without using VCI according to the conventional plan.
- vci.maintenance\_work\_mem specifies the memory size to be secured dynamically. If it exceeds the specified value, a disk temporary file is used for operation.

#### Memory used for VCI maintenance

```
memUsedForVciMaintenace = vci.maintenance_work_mem x vci.control_max_workers
```

#### Memory used by CREATE INDEX

```
memUsedByCreateIndexStatement = vci.maintenance_work_mem
```



vci.maintenance\_work\_mem specifies the memory to be secured dynamically. If it exceeds the specified value, a disk temporary file is used for operation.

## F.4 High-Speed Data Load Memory Requirements

This section describes the formula for estimating memory requirements for the high-speed data load feature.

Use the following formula to obtain a rough estimate of memory requirements:

```
Memory usage of high speed data load

= (Peak memory usage of pgx_loader processes + Peak memory usage of the pgx_loader commands)

x Number of commands executed simultaneously

Peak memory usage of pgx_loader processes

= Peak memory usage of the backend process (6 MB)

+ Peak memory usage of parallel workers (6 MB x number of parallel workers)

+ Peak memory usage of dynamic shared memory (80 MB x number of parallel workers)

Peak memory usage of the pgx_loader commands=9 MB
```



In addition to the size calculated using the formula above, the database cache on the shared memory estimated using the shared\_buffers parameter is consumed according to the size of the data (table and index keys) loaded using this feature. Refer to "E.1 Estimating Table Size Requirements" and "E.2 Estimating Index Size Requirements" for information on estimating an appropriate shared buffers value.

# F.5 Global Meta Cache Memory Requirements

This section describes the formula for estimating Global Meta Cache memory requirements.

The memory calculated by "Size of the GMC area" is allocated to the shared memory. The memory calculated by the per-process meta cache management information is allocated to the local memory. Refer to the graphic in "Architecture of Global Meta Cache Feature" in the "Memory usage reduction by Global Meta Cache" in the General Description for more information.

Use the following formula to obtain a rough estimate of memory requirements:

```
Amount of memory used by the Global Meta Cache feature

= Size of GMC area + Per-process meta cache management information

Size of GMC area = (All user tables x 0.4 KB

+ All user indexes x 0.3 KB

+ All user columns x 0.8 KB) x 1.5 (*1)

Per-process meta cache management information

= (All user tables + All user indexes + All user columns) x 1.0KB x max_connections x 1.5 (*1)
```

This value takes into account the case where both GMC before and after the change temporarily exist at the same time in shared memory when the table definition is changed or the row of the system catalog is changed.

<sup>\*1)</sup> Safety Factor (1.5)

# Appendix G Quantitative Limits

This appendix lists the quantitative limits of FUJITSU Enterprise Postgres.

Table G.1 Length of identifier

Item	Limit
Database name	Up to 63 bytes (*1) (*2)
Schema name	Up to 63 bytes (*1) (*2)
Table name	Up to 63 bytes (*1) (*2)
View name	Up to 63 bytes (*1) (*2)
Index name	Up to 63 bytes (*1) (*2)
Tablespace name	Up to 63 bytes (*1) (*2)
Cursor name	Up to 63 bytes (*1) (*2)
Function name	Up to 63 bytes (*1) (*2)
Aggregate function name	Up to 63 bytes (*1) (*2)
Trigger name	Up to 63 bytes (*1) (*2)
Constraint name	Up to 63 bytes (*1) (*2)
Conversion name	Up to 63 bytes (*1) (*2)
Role name	Up to 63 bytes (*1) (*2)
Cast name	Up to 63 bytes (*1) (*2)
Collation sequence name	Up to 63 bytes (*1) (*2)
Encoding method conversion name	Up to 63 bytes (*1) (*2)
Domain name	Up to 63 bytes (*1) (*2)
Extension name	Up to 63 bytes (*1) (*2)
Operator name	Up to 63 bytes (*1) (*2)
Operator class name	Up to 63 bytes (*1) (*2)
Operator family name	Up to 63 bytes (*1) (*2)
Rewrite rule name	Up to 63 bytes (*1) (*2)
Sequence name	Up to 63 bytes (*1) (*2)
Text search settings name	Up to 63 bytes (*1) (*2)
Text search dictionary name	Up to 63 bytes (*1) (*2)
Text search parser name	Up to 63 bytes (*1) (*2)
Text search template name	Up to 63 bytes (*1) (*2)
Data type name	Up to 63 bytes (*1) (*2)
Enumerator type label	Up to 63 bytes (*1) (*2)

<sup>\*1:</sup> This is the character string byte length when converted by the server character set character code.

Table G.2 Database object

Item	Limit
Number of databases	Less than 4,294,967,296 (*1)

<sup>\*2:</sup> If an identifier that exceeds 63 bytes in length is specified, the excess characters are truncated and it is processed.

Item	Limit
Number of schemas	Less than 4,294,967,296 (*1)
Number of tables	Less than 4,294,967,296 (*1)
Number of views	Less than 4,294,967,296 (*1)
Number of indexes	Less than 4,294,967,296 (*1)
Number of tablespaces	Less than 4,294,967,296 (*1)
Number of functions	Less than 4,294,967,296 (*1)
Number of aggregate functions	Less than 4,294,967,296 (*1)
Number of triggers	Less than 4,294,967,296 (*1)
Number of constraints	Less than 4,294,967,296 (*1)
Number of conversion	Less than 4,294,967,296 (*1)
Number of roles	Less than 4,294,967,296 (*1)
Number of casts	Less than 4,294,967,296 (*1)
Number of collation sequences	Less than 4,294,967,296 (*1)
Number of encoding method conversions	Less than 4,294,967,296 (*1)
Number of domains	Less than 4,294,967,296 (*1)
Number of extensions	Less than 4,294,967,296 (*1)
Number of operators	Less than 4,294,967,296 (*1)
Number of operator classes	Less than 4,294,967,296 (*1)
Number of operator families	Less than 4,294,967,296 (*1)
Number of rewrite rules	Less than 4,294,967,296 (*1)
Number of sequences	Less than 4,294,967,296 (*1)
Number of text search settings	Less than 4,294,967,296 (*1)
Number of text search dictionaries	Less than 4,294,967,296 (*1)
Number of text search parsers	Less than 4,294,967,296 (*1)
Number of text search templates	Less than 4,294,967,296 (*1)
Number of data types	Less than 4,294,967,296 (*1)
Number of enumerator type labels	Less than 4,294,967,296 (*1)
Number of default access privileges defined in the ALTER DEFAULT PRIVILEGES statement	Less than 4,294,967,296 (*1)
Number of large objects	Less than 4,294,967,296 (*1)
Number of index access methods	Less than 4,294,967,296 (*1)

<sup>\*1:</sup> The total number of all database objects must be less than 4,294,967,296.

## Table G.3 Schema element

Item	Limit
Number of columns that can be defined in one table	From 250 to 1600 (according to the data type)
Table row length	Up to 400 gigabytes
Number of columns comprising a unique constraint	Up to 32 columns
Data length comprising a unique constraint	Less than 2,000 bytes (*1) (*2)

Item	Limit
Table size	Up to 32 terabyte
Search condition character string length in a trigger definition statement	Up to 800 megabytes (*1) (*2)
Item size	Up to 1 gigabyte

<sup>\*1:</sup> Operation might proceed correctly even if operations are performed with a quantity outside the limits.

#### Table G.4 Index

Item	Limit
Number of columns comprising a key (including VCI)	Up to 32 columns
Key length (other than VCI)	Less than 2,000 bytes (*1)

<sup>\*1:</sup> This is the character string byte length when converted by the server character set character code.

Table G.5 Data types and attributes that can be handled

Item		Limit	
Character	Data length		Data types and attributes that can be handled (*1)
	Specification length (n)		Up to 10,485,760 characters (*1)
Numeric	External decimal expression		Up to 131,072 digits before the decimal point, and up to 16,383 digits after the decimal point
	Internal binary	2 bytes	From -32,768 to 32,767
	expression	4 bytes	From -2,147,483,648 to 2,147,483,647
		8 bytes	From -9,223,372,036,854,775,808 to 9,223,372,036,854,775,807
	Internal decimal expression		Up to 13,1072 digits before the decimal point, and up to 16,383 digits after the decimal point
	Floating point expression	4 bytes	From -3.4E+38 to -7.1E-46, 0, or from 7.1E-46 to 3.4E+38
		8 bytes	From -1.7E+308 to -2.5E-324, 0, or from 2.5E-324 to 1.7E+308
bytea			Up to one gigabyte minus 53 bytes
Large object			Up to 4 terabyte

<sup>\*1:</sup> This is the character string byte length when converted by the server character set character code.

#### Table G.6 Function definition

Item	Limit
Number of arguments that can be specified	Up to 100
Number of variable names that can be specified in the declarations section	No limit
Number of SQL statements or control statements that can be specified in a function processing implementation	No limit

<sup>\*2:</sup> This is the character string byte length when converted by the server character set character code.

Table G.7 Data operation statement

Maximum number of connections for one process in an application (remote access)  Number of expressions that can be specified in a FROM clause  Number of tubles that can be specified in a FROM clause  Number of unique expressions that can be specified in a selection list DISTINCT clause/ORDER BY clause/GROUP BY clause within one SELECT statement  Number of expressions that can be specified in an ORDER BY clause  Number of expressions that can be specified in an ORDER BY clause  Number of expressions that can be specified in an UNION clauses/INTERSECT clause/EXCEPT clause  Number of sexpressions that can be specified in a UNION clauses/INTERSECT clause/EXCEPT clause  Number of nestings in joined tables that can be specified in one view  Number of functions or operator expressions that can be specified in one row constructor  Number of expressions that can be specified in one row constructor  Number of expressions that can be specified in one row constructor  Number of expressions that can be specified in one row of a VALUES  list  Number of expressions that can be specified in one row of a VALUES  list  Number of expressions that can be specified in a RETURNING clause  Up to 1,664  Number of expressions that can be specified in the argument list of one for expressions that can be specified in a RETURNING clause  Up to 1,664  Number of expressions that can be specified in the argument list of one for processed simultaneously by one session  Number of input parameter specifications that can be specified in one for processed simultaneously by one session  Number of input parameter specifications that can be specified in one SQL statement  Number of tokens that can be specified as a list in a WHERE clause IN so limit  Number of 20 IN shat can be specified in a COALESCE  No limit  Number of 20 IN shat can be specified in a COALESCE  No limit  Number of bythe shat can be specified in COALESCE  No limit	Item	Limit
Number of tables that can be specified in a FROM clause  Number of unique expressions that can be specified in a selection list/ DISTINCT clause/ORDER BY clause/GROUP BY clause within one SELECT statement  Number of expressions that can be specified in a GROUP BY clause  Number of expressions that can be specified in an ORDER BY clause  Number of expressions that can be specified in a UNION clause/INTERSECT clause/EXCEPT clause  Number of Inactions or operator expressions that can be specified in one view  Number of Inactions or operator expressions that can be specified in one expression that can be specified in an UPDATE statement  SET clause  Number of expressions that can be specified in one row constructor  Number of expressions that can be specified in one row of a VALUES list  Number of expressions that can be specified in one row of a VALUES list  Number of expressions that can be specified in a RETURNING clause  Number of expressions that can be specified in the argument list of one function specification  Number of cursors that can be specified in the argument list of one function specification  Number of input parameter specifications that can be specified in one specified in one dynamic SQL statement  Number of input parameter specifications that can be specified in one dynamic SQL statement  Number of values that can be specified as a list in a WHERE clause IN syntax  Number of expressions that can be specified in a USING clause  Number of expressions that can be specified in a USING clause  Number of expressions that can be specified in a OOALESCE  Number of wHEND clauses that can be specified in COALESCE  Number of WHEND clauses that can be specified in COALESCE  Number of WHEND clauses that can be specified or inserted by one SQL  statement  Vip to one gigabyte minus 53 bytes	1	4,000 connections
Number of unique expressions that can be specified in a selection list/ DISTINCT clause/ORDER BY clause/GROUP BY clause within one SELECT statement  Number of expressions that can be specified in a GROUP BY clause within one SELECT statement by the specified in a GROUP BY clause on the specified in a UNION clause/INTERSECT clause/EXCEPT clause  Number of sexpressions that can be specified in a UNION clause/INTERSECT clause/EXCEPT clause  Number of functions or operator expressions that can be specified in one view one expression shat can be specified in one row constructor  Number of functions or operator expressions that can be specified in one row constructor  Number of expressions that can be specified in one row constructor  Number of expressions that can be specified in one row of a VALUES list  Number of expressions that can be specified in a RETURNING clause  Number of expressions that can be specified in the argument list of one function specification  Number of cursors that can be specified in the argument list of one function specification  Number of cursors that can be processed simultaneously by one session  Character string length of one SQL statement  Number of input parameter specifications that can be specified in one SQL statement  Number of tokens that can be specified as a list in a WHERE clause IN syntax  Number of expressions that can be specified in a USING clause  No limit  Number of Appressions that can be specified in a USING clause  No limit  Number of expressions that can be specified in a USING clause  No limit  Number of expressions that can be specified in a USING clause  No limit  Number of expressions that can be specified in a USING clause  No limit  Number of expressions that can be specified in a USING clause  No limit  Number of expressions that can be specified in a USING clause  No limit  Number of expressions that can be specified in a USING clause  No limit  Number of or pressions that can be specified in a USING clause  No limit  Number of or pressions that can be specif	Number of expressions that can be specified in a selection list	Up to 1,664
DISTINCT clause/ORDER BY clause/GROUP BY clause within one SELECT statement  Number of expressions that can be specified in a GROUP BY clause No limit  Number of expressions that can be specified in an ORDER BY clause No limit  Number of SELECT statements that can be specified in a UNION clause/INTERSECT clause/EXCEPT clause  Number of nestings in joined tables that can be specified in one view Number of functions or operator expressions that can be specified in one row constructor Number of expressions that can be specified in one row constructor Number of expressions that can be specified in one row constructor Number of expressions that can be specified in one row of a VALUES list Number of expressions that can be specified in one row of a VALUES list Number of expressions that can be specified in one row of a VALUES list Number of expressions that can be specified in the argument list of one function specification Number of cursors that can be processed simultaneously by one session Character string length of one SQL statement Number of input parameter specifications that can be specified in one SQL statement Number of tokens that can be specified in one SQL statement Number of values that can be specified in a USING clause No limit Number of expressions that can be specified in a USING clause No limit Number of values that can be specified in a USING clause No limit Number of values that can be specified in a USING clause No limit Number of values that can be specified in a USING clause No limit Number of values that can be specified in a USING clause No limit Number of values that can be specified in a USING clause No limit Number of values that can be specified in a USING clause No limit Number of values that can be specified in a USING clause No limit Number of values that can be specified in cOALESCE No limit Number of wHEN clauses that can be specified in COALESCE No limit  Data size per record that can be updated or inserted by one SQL statement  Data size per record that can be updated or inserted b	Number of tables that can be specified in a FROM clause	No limit
Number of expressions that can be specified in an ORDER BY clause  Number of SELECT statements that can be specified in a UNION clause/INTERSECT clause/EXCEPT clause  Number of nestings in joined tables that can be specified in one view  Number of functions or operator expressions that can be specified in one view  Number of functions or operator expressions that can be specified in one row constructor  Number of expressions that can be specified in one row constructor  Number of expressions that can be specified in an UPDATE statement  SET clause  Number of expressions that can be specified in one row of a VALUES  list  Number of expressions that can be specified in a RETURNING clause  Number of expressions that can be specified in the argument list of one function specification  Number of cursors that can be specified in the argument list of one function specification  Number of cursors that can be processed simultaneously by one session  Character string length of one SQL statement  Number of input parameter specifications that can be specified in one dynamic SQL statement  Number of tokens that can be specified in one SQL statement  Number of values that can be specified in one SQL statement  Number of values that can be specified in a USING clause  No limit  Number of expressions that can be specified in a Joined table  Number of expressions that can be specified in OOALESCE  No limit  Number of wHEN clauses that can be specified for CASE in a simple format or a searched format  Data size per record that can be updated or inserted by one SQL statement  Up to one gigabyte minus 53 bytes	DISTINCT clause/ORDER BY clause/GROUP BY clause within one	Up to 1,664
Number of SELECT statements that can be specified in a UNION clause/INTERSECT clause/EXCEPT clause  Number of nestings in joined tables that can be specified in one view  Number of functions or operator expressions that can be specified in one view  Number of functions or operator expressions that can be specified in one row constructor  Number of expressions that can be specified in one row constructor  Number of expressions that can be specified in an UPDATE statement SET clause  Number of expressions that can be specified in one row of a VALUES list  Number of expressions that can be specified in a RETURNING clause  Number of expressions that can be specified in the argument list of one function specification  Number of cursors that can be processed simultaneously by one session  Number of cursors that can be processed simultaneously by one session  Character string length of one SQL statement  Number of input parameter specifications that can be specified in one dynamic SQL statement  Number of tokens that can be specified in one SQL statement  Number of values that can be specified in one SQL statement  Number of values that can be specified as a list in a WHERE clause IN syntax  Number of expressions that can be specified in a USING clause  No limit  Number of expressions that can be specified in a Joined table  Number of expressions that can be specified in a joined table  Number of WHEN clauses that can be specified for CASE in a simple format or a searched format  Data size per record that can be updated or inserted by one SQL statement  Up to one gigabyte minus 53 bytes	Number of expressions that can be specified in a GROUP BY clause	No limit
clause/INTERSECT clause/EXCEPT clause  Number of nestings in joined tables that can be specified in one view one expression  Number of functions or operator expressions that can be specified in one expression  Number of expressions that can be specified in one row constructor  Number of expressions that can be specified in an UPDATE statement SET clause  Number of expressions that can be specified in one row of a VALUES list  Number of expressions that can be specified in a RETURNING clause  Number of expressions that can be specified in the argument list of one function specification  Number of cursors that can be processed simultaneously by one session  Character string length of one SQL statement  Number of input parameter specifications that can be specified in one SQL statement  Number of tokens that can be specified as a list in a WHERE clause IN syntax  Number of values that can be specified in a USING clause  No limit  Number of JOINs that can be specified in COALESCE  No limit  No limit  No limit  No limit  No limit  No limit  Number of wHEN clauses that can be specified in COALESCE  No limit  No limit  No limit  No limit  No limit  No limit	Number of expressions that can be specified in an ORDER BY clause	No limit
Number of functions or operator expressions that can be specified in one expression  Number of expressions that can be specified in one row constructor  Number of expressions that can be specified in an UPDATE statement SET clause  Number of expressions that can be specified in one row of a VALUES list  Number of expressions that can be specified in one row of a VALUES list  Number of expressions that can be specified in a RETURNING clause  Total expression length that can be specified in the argument list of one function specification  Number of cursors that can be processed simultaneously by one session  Character string length of one SQL statement  Number of input parameter specifications that can be specified in one dynamic SQL statement  Number of tokens that can be specified in one SQL statement  Number of values that can be specified as a list in a WHERE clause IN syntax  Number of expressions that can be specified in a USING clause  No limit  Number of syntax is a list in a WHERE clause IN syntax  No limit  Number of expressions that can be specified in COALESCE  No limit  Number of WHEN clauses that can be specified for CASE in a simple format or a searched format  Data size per record that can be updated or inserted by one SQL statement  Up to one gigabyte minus 53 bytes	-	Up to 4,000 (*1)
one expression       Up to 1,664         Number of expressions that can be specified in one row constructor       Up to 1,664         Number of expressions that can be specified in an UPDATE statement SET clause       Up to 1,664         Number of expressions that can be specified in one row of a VALUES list       Up to 1,664         Number of expressions that can be specified in a RETURNING clause       Up to 1,664         Total expression length that can be specified in the argument list of one function specification       Up to 800 megabytes (*2)         Number of cursors that can be processed simultaneously by one session       No limit         Character string length of one SQL statement       Up to 800 megabytes (*1) (*3)         Number of input parameter specifications that can be specified in one dynamic SQL statement       No limit         Number of tokens that can be specified in one SQL statement       Up to 10,000         Number of values that can be specified as a list in a WHERE clause IN syntax       No limit         Number of expressions that can be specified in a USING clause       No limit         Number of expressions that can be specified in a joined table       Up to 4,000 (*1)         Number of WHEN clauses that can be specified for CASE in a simple format or a searched format       No limit         Data size per record that can be updated or inserted by one SQL statement       Up to one gigabyte minus 53 bytes	Number of nestings in joined tables that can be specified in one view	Up to 4,000 (*1)
Number of expressions that can be specified in an UPDATE statement SET clause  Number of expressions that can be specified in one row of a VALUES list  Number of expressions that can be specified in a RETURNING clause Total expression length that can be specified in the argument list of one function specification  Number of cursors that can be processed simultaneously by one session  Character string length of one SQL statement  Number of input parameter specifications that can be specified in one dynamic SQL statement  Number of tokens that can be specified in one SQL statement  Up to 10,000  No limit  Vup to 10,000  No limit  No limit  No limit  Number of values that can be specified in a USING clause  No limit  Number of expressions that can be specified in a Joined table  Number of expressions that can be specified in COALESCE  No limit  Number of WHEN clauses that can be specified for CASE in a simple format or a searched format  Data size per record that can be updated or inserted by one SQL statement  Up to one gigabyte minus 53 bytes		Up to 4,000 (*1)
Number of expressions that can be specified in one row of a VALUES list  Number of expressions that can be specified in a RETURNING clause  Total expression length that can be specified in the argument list of one function specification  Number of cursors that can be processed simultaneously by one session  Character string length of one SQL statement  Up to 800 megabytes (*2)  No limit  Number of input parameter specifications that can be specified in one dynamic SQL statement  Number of tokens that can be specified in one SQL statement  Up to 10,000  Number of values that can be specified in a WHERE clause IN syntax  Number of expressions that can be specified in a USING clause  No limit  Number of JOINs that can be specified in COALESCE  No limit  Number of WHEN clauses that can be specified for CASE in a simple format or a searched format  Data size per record that can be updated or inserted by one SQL statements Up to one gigabyte minus 53 bytes	Number of expressions that can be specified in one row constructor	Up to 1,664
Number of expressions that can be specified in the argument list of one function specification  Number of cursors that can be processed simultaneously by one session  Character string length of one SQL statement  Number of input parameter specifications that can be specified in one dynamic SQL statement  Number of tokens that can be specified in one SQL statement  Number of values that can be specified as a list in a WHERE clause IN syntax  Number of JOINs that can be specified in a USING clause  Number of syressions that can be specified in a joined table  Number of wHEN clauses that can be specified in COALESCE  Number of WHEN clauses that can be specified for CASE in a simple format or a searched format  Data size per record that can be updated or inserted by one SQL statement  Up to 1,664  Up to 800 megabytes (*2)  No limit  Vip to 4,000 (*1)  No limit  Up to 4,000 (*1)  No limit  Up to 4,000 (*1)  No limit  Up to one gigabyte minus 53 bytes		Up to 1,664
Total expression length that can be specified in the argument list of one function specification  Number of cursors that can be processed simultaneously by one session  Character string length of one SQL statement  Up to 800 megabytes (*2)  Number of input parameter specifications that can be specified in one dynamic SQL statement  Number of tokens that can be specified in one SQL statement  Up to 10,000  Number of values that can be specified as a list in a WHERE clause IN syntax  Number of expressions that can be specified in a USING clause  Number of JOINs that can be specified in a joined table  Up to 4,000 (*1)  Number of wHEN clauses that can be specified for CASE in a simple format or a searched format  Data size per record that can be updated or inserted by one SQL statement  Up to one gigabyte minus 53 bytes		Up to 1,664
Number of cursors that can be processed simultaneously by one session  Character string length of one SQL statement  Up to 800 megabytes (*1) (*3)  Number of input parameter specifications that can be specified in one dynamic SQL statement  Number of tokens that can be specified in one SQL statement  Up to 10,000  Number of values that can be specified as a list in a WHERE clause IN syntax  Number of expressions that can be specified in a USING clause  Number of JOINs that can be specified in a joined table  Up to 4,000 (*1)  Number of WHEN clauses that can be specified for CASE in a simple format or a searched format  Data size per record that can be updated or inserted by one SQL statement  Up to one gigabyte minus 53 bytes	Number of expressions that can be specified in a RETURNING clause	Up to 1,664
Character string length of one SQL statement  Up to 800 megabytes (*1) (*3)  Number of input parameter specifications that can be specified in one dynamic SQL statement  Number of tokens that can be specified in one SQL statement  Up to 10,000  Number of values that can be specified as a list in a WHERE clause IN syntax  Number of expressions that can be specified in a USING clause  No limit  Number of JOINs that can be specified in a joined table  Up to 4,000 (*1)  Number of expressions that can be specified in COALESCE  No limit  Number of WHEN clauses that can be specified for CASE in a simple format or a searched format  Data size per record that can be updated or inserted by one SQL statement  Up to one gigabyte minus 53 bytes		Up to 800 megabytes (*2)
Number of input parameter specifications that can be specified in one dynamic SQL statement  Number of tokens that can be specified in one SQL statement  Up to 10,000  Number of values that can be specified as a list in a WHERE clause IN syntax  Number of expressions that can be specified in a USING clause  Number of JOINs that can be specified in a joined table  Up to 4,000 (*1)  Number of expressions that can be specified in COALESCE  Number of WHEN clauses that can be specified for CASE in a simple format or a searched format  Data size per record that can be updated or inserted by one SQL statement  Up to one gigabyte minus 53 bytes		No limit
Number of tokens that can be specified in one SQL statement  Number of values that can be specified as a list in a WHERE clause IN syntax  Number of expressions that can be specified in a USING clause  Number of JOINs that can be specified in a joined table  Number of expressions that can be specified in COALESCE  Number of WHEN clauses that can be specified for CASE in a simple format or a searched format  Data size per record that can be updated or inserted by one SQL statement  Up to 10,000  No limit  No limit  No limit  Up to one gigabyte minus 53 bytes	Character string length of one SQL statement	Up to 800 megabytes (*1) (*3)
Number of values that can be specified as a list in a WHERE clause IN syntax  Number of expressions that can be specified in a USING clause  No limit  Number of JOINs that can be specified in a joined table  Number of expressions that can be specified in COALESCE  No limit  Number of WHEN clauses that can be specified for CASE in a simple format or a searched format  Data size per record that can be updated or inserted by one SQL statement  Up to one gigabyte minus 53 bytes		No limit
Number of expressions that can be specified in a USING clause  No limit  Number of JOINs that can be specified in a joined table  Up to 4,000 (*1)  Number of expressions that can be specified in COALESCE  No limit  Number of WHEN clauses that can be specified for CASE in a simple format or a searched format  Data size per record that can be updated or inserted by one SQL statement  Up to one gigabyte minus 53 bytes	Number of tokens that can be specified in one SQL statement	Up to 10,000
Number of JOINs that can be specified in a joined table  Up to 4,000 (*1)  Number of expressions that can be specified in COALESCE  No limit  Number of WHEN clauses that can be specified for CASE in a simple format or a searched format  Data size per record that can be updated or inserted by one SQL statement  Up to 4,000 (*1)  No limit  Up to one gigabyte minus 53 bytes		No limit
Number of expressions that can be specified in COALESCE  No limit  Number of WHEN clauses that can be specified for CASE in a simple format or a searched format  Data size per record that can be updated or inserted by one SQL statement  Up to one gigabyte minus 53 bytes	Number of expressions that can be specified in a USING clause	No limit
Number of WHEN clauses that can be specified for CASE in a simple format or a searched format  Data size per record that can be updated or inserted by one SQL statement  Up to one gigabyte minus 53 bytes	Number of JOINs that can be specified in a joined table	Up to 4,000 (*1)
format or a searched format  Data size per record that can be updated or inserted by one SQL statement  Up to one gigabyte minus 53 bytes	Number of expressions that can be specified in COALESCE	No limit
statement		No limit
Number of objects that can share a lock simultaneously  Up to 256,000 (*1)		Up to one gigabyte minus 53 bytes
	Number of objects that can share a lock simultaneously	Up to 256,000 (*1)

<sup>\*1:</sup> Operation might proceed correctly even if operations are performed with a quantity outside the limits.

<sup>\*2:</sup> The total number of all database objects must be less than 4,294,967,296.

<sup>\*3:</sup> This is the character string byte length when converted by the server character set character code.

## Table G.8 Data size

Item	Limit
Data size per record for input data files (COPY statement, psql command \copy meta command)	Up to 800 megabytes (*1)
Data size per record for output data files (COPY statement, psql command \copy meta command)	Up to 800 megabytes (*1)

<sup>\*1:</sup> Operation might proceed correctly even if operations are performed with a quantity outside the limits.

# Appendix H Determining the Preferred WebAdmin Configuration

This appendix describes the two different configurations in which WebAdmin can be used and how to select the most suitable configuration.

## H.1 WebAdmin Configurations

WebAdmin can be installed in two configurations:

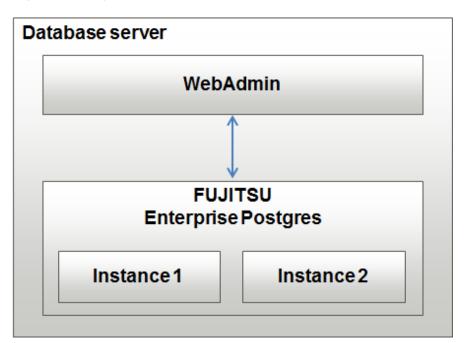
- Single-server
- Multiserver

WebAdmin does not support encrypted communication between browser and server or between servers. Therefore, when using WebAdmin in either configuration, build the communication path with the browser or each server on a network that cannot be accessed externally.

## **H.1.1 Single-Server Configuration**

A single-server configuration enables you to create and operate instances on a single server. In this configuration, WebAdmin must be installed on the same database server as the FUJITSU Enterprise Postgres Server component.

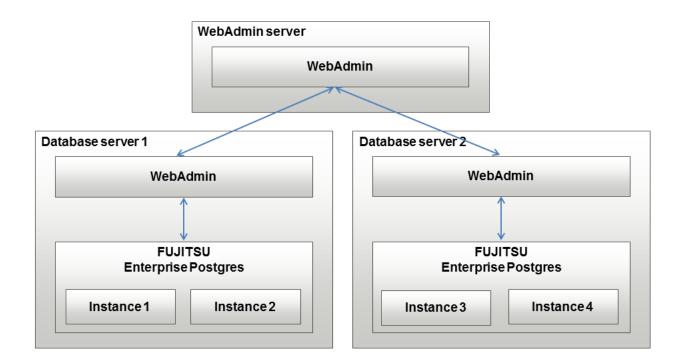
Single-server configuration



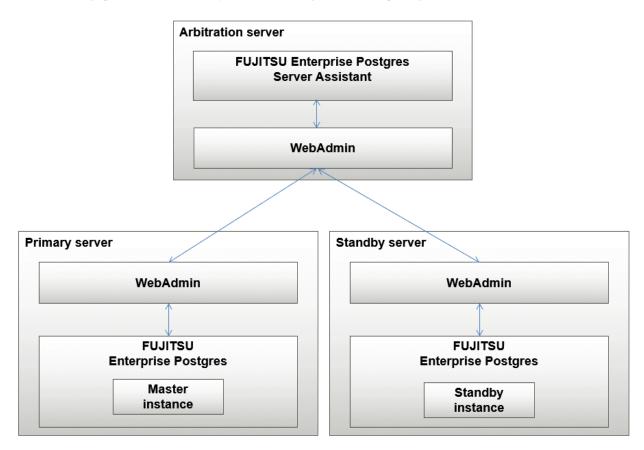
## H.1.2 Multiserver Configuration

A multiserver configuration enables you to create and operate instances stored on multiple database servers. As shown in the figure below, WebAdmin can be installed on a dedicated WebAdmin server and used to collectively manage the instances stored on the database servers.

Multiserver configuration



Also, when setting up the arbitration server by WebAdmin during database multiplexing mode, install WebAdmin on the arbitration server.



# H.2 Installing WebAdmin in a Single-Server Configuration

To install WebAdmin in a single-server configuration, the FUJITSU Enterprise Postgres Server component and WebAdmin must be installed on the same machine.

Select the following items when installing FUJITSU Enterprise Postgres in a single-server configuration:

- FUJITSU Enterprise Postgres Advanced Edition or FUJITSU Enterprise Postgres Standard Edition
- WebAdmin

# H.3 Installing WebAdmin in a Multiserver Configuration

In a multiserver configuration, install WebAdmin on one server, and both WebAdmin and the FUJITSU Enterprise Postgres Server component on any number of database servers.

Select the following items when installing FUJITSU Enterprise Postgres in a multiserver configuration:

- WebAdmin server:
  - WebAdmin
- Database server:
  - FUJITSU Enterprise Postgres Advanced Edition or FUJITSU Enterprise Postgres Standard Edition
  - WebAdmin

Also, when setting up the arbitration server by WebAdmin during database multiplexing mode, select the following when installing FUJITSU Enterprise Postgres.

.....

- Arbitration server
  - FUJITSU Enterprise Postgres Server Assistant
  - WebAdmin



Refer to the Installation and Setup Guide for Server Assistant for details on how to install the Server Assistant.

# Appendix I Supported contrib Modules and Extensions Provided by External Projects

FUJITSU Enterprise Postgres supports PostgreSQL contrib modules, and extensions provided by external projects.

Refer to the following for details on the supported contrib modules:

- "Additional Supplied Modules" in the PostgreSQL Documentation
- "Additional Supplied Programs" in the PostgreSQL Documentation



You can also check the list of available extensions using the pg\_available\_extensions view.

Refer to "OSS Supported by FUJITSU Enterprise Postgres" in the General Description for information on supported extensions provided by external projects.

# <u>Index</u>

[C]
Changing client authentication information
Changing instance settings
Check the disk space8
Check the installed products and determine the installation
method8
Client Authentication Information settings
Creating an Instance
Creating an Instance Administrator
Creating Instances21
[D]
• •
Disk Space Required for Installation5
[E]
Editing instance information
Excluded Software4
[F]
Firewall
17,55
[H]
Hardware Environment5
How to Set Up the Pop-up Blocker47
[1]
Importing Instances
Installation
Installation in Interactive Mode
Installation in Silent Mode
Installation Procedure1
Installation Types1
Instance configuration
[L]
Logging in to WebAdmin23
[M]
Multi-Version Installation
[N]
New Installation1
101
[0]
Operating Environment
Operating Method Types and Selection
rD1
[P]
Port number to use when Tomcat is stopped48
postgresql.conf54
Pre-installation Tasks
Preparations for Setup
[R]
Recommended Browser Settings47
Reinstallation1
Related Software
Remove applied updates8

Removing WebAdmin	50
Required Operating System	3
Required Patches	5
[S]	
Settings related to connection	33
Setting Up and Removing WebAdmin	48
Setting Up WebAdmin	48
Starting the Web Server Feature of WebAdmin	49
Stopping the Web Server Feature of WebAdmin	49
Supported contrib Modules and Extensions Provided by Extensions	ernal
Projects	76
[T]	
TCP/IP Protocol	5
n n	
[U]	
Uninstallation	
Uninstallation in Interactive Mode	
Uninstallation in Silent Mode	
uninstaller	
Using the initdb Command	
Using WebAdmin	22
ПАЛ	
[W]	40
WebAdmin automatic start	
Web server port number	
When an Instance was Created with the initdb Command	
When an Instance was Created with WebAdmin	33



# FUJITSU Enterprise Postgres 12



# Installation and Setup Guide for Client

# Preface

# Purpose of this document

This document describes how to install, uninstall and set up the "FUJITSU Enterprise Postgres client feature".

#### Intended readers

This document is intended for those who install and operate FUJITSU Enterprise Postgres.

Readers of this document are assumed to have general knowledge of:



- PostgreSQL
- SQL
- Windows



- PostgreSQL
- SQL
- Linux

#### Structure of this document

This document is structured as follows:

# Chapter 1 Overview of Installation

Describes the features that can be installed, and provides an overview of installation methods

## Chapter 2 Installation and Uninstallation of the Windows Client

Describes how to install the FUJITSU Enterprise Postgres client feature (Windows client)

# Chapter 3 Installation and Uninstallation of the Linux Client

Describes how to install the FUJITSU Enterprise Postgres client feature (Linux client)

# Chapter 4 Setup

Describes the setup procedures to be performed after installation completes

# **Export restrictions**

Exportation/release of this document may require necessary procedures in accordance with the regulations of your resident country and/or US export control laws.

# Issue date and version

Edition 1.0: July 2020

# Copyright

Copyright 2015-2020 FUJITSU LIMITED

# Contents

Chapter 1 Overview of Installation	
1.1 Features that can be Installed.	
1.2 Installation Types	
1.2.1 New Installation.	,
1.2.2 Reinstallation	
1.2.3 Multi-Version Installation.	
1.3 Installation Procedure.	
1.3.1 Installation in Interactive Mode	
1.3.2 Installation in Silent Mode	
1.4 Uninstallation	
Chapter 2 Installation and Uninstallation of the Windows Client	:
2.1 Operating Environment.	
2.1.1 Required Operating System.	
2.1.2 Related Software	
2.1.3 Excluded Software	
2.1.4 Hardware Environment.	
2.1.5 Disk Space Required for Installation.	
2.1.6 Supported System Environment	
2.1.7 Versions of Open-Source Software Used as the Base for FUJITSU Enterprise Postgres Drivers	
2.2 Installation.	
2.2.1 Pre-installation Tasks	
2.2.2 Pre-installation Considerations.	
2.2.3 Installation in Interactive Mode	
2.2.4 Installation in Silent Mode	
2.3 Uninstallation.	10
2.3.1 Uninstallation in Interactive Mode	1
2.3.2 Uninstallation in Silent Mode	1
2.4 Registering .NET Data Provider	12
2.4.1 Registering Entries to the machine.config File	13
Chapter 3 Installation and Uninstallation of the Linux Client	11
3.1 Operating Environment	
3.1.1 Required Operating System	
3.1.2 Related Software.	
3.1.3 Excluded Software	
3.1.4 Required Patches.	
3.1.5 Hardware Environment.	
3.1.6 Disk Space Required for Installation.	
3.1.7 Supported System Environment	
3.1.8 Versions of Open-Source Software Used as the Base for FUJITSU Enterprise Postgres Drivers	
3.2 Installation.	
3.2.1 Pre-installation Tasks	
3.2.2 Installation in Interactive Mode	
3.2.3 Installation in Silent Mode	2
3.3 Uninstallation	23
3.3.1 Uninstallation in Interactive Mode	23
3.3.2 Uninstallation in Silent Mode	24
Chapter 4 Setup	0.1
Chapter 4 Setup.	
4.1 Configuring Environment Variables	
4.1.1 Wildows	
T.1.2 Dilux.	
Index	26

# Chapter 1 Overview of Installation

This chapter provides an overview of FUJITSU Enterprise Postgres installation.

# 1.1 Features that can be Installed

FUJITSU Enterprise Postgres provides features to enable access to the database from a variety of platforms and languages, as the connection environment for the client and the database server.

The FUJITSU Enterprise Postgres client package must be installed on the client system to use these features.

The following table shows the relationship between the platforms and the features provided by client packages.

Platform	JDBC	ODBC	.NET Data Provide r	C languag e (libpq)	ed SQL	ed SQL (ECOBP G) in	Connect ion Manage r	High- speed data load	pgAdmi n	Pgpool- II
Windows	Y	Y	Y	Y	Y	Y		Y	Y	
Linux	Y	Y		Y	Y	Y	Y	Y		Y

Y: Supported

# 1.2 Installation Types

The following installation types are available for FUJITSU Enterprise Postgres:

- New installation
- Reinstallation
- Multi-version installation

# 1.2.1 New Installation

In initial installation, the FUJITSU Enterprise Postgres client feature is installed for the first time.

# 1.2.2 Reinstallation

Perform reinstallation to repair installed program files that have become unusable for any reason.

# 1.2.3 Multi-Version Installation

Perform multi-version installation to install different versions to the installed program files separately.

# 1.3 Installation Procedure

The following installation procedures are available for FUJITSU Enterprise Postgres:

- Installation in interactive mode
- Installation in silent mode

Select the installation procedure that corresponds to your environment.

# 1.3.1 Installation in Interactive Mode

Interactive mode enables installation to be performed while the required information is entered interactively.

In interactive mode installation, the FUJITSU Enterprise Postgres installer automatically determines the installation state of FUJITSU Enterprise Postgres. Install FUJITSU Enterprise Postgres using one of the following installation types in accordance with the installation state:

- New installation
- Reinstallation
- Multi-version installation

# 1.3.2 Installation in Silent Mode

Silent mode enables installation to be performed without the need to enter any information interactively.

New installations and multi-version installations can be performed in silent mode.

# 1.4 Uninstallation

Uninstallation removes the system files of the installed FUJITSU Enterprise Postgres client feature.

# Chapter 2 Installation and Uninstallation of the Windows Client

This chapter explains how to install and uninstall the Windows client.

# 2.1 Operating Environment

This section describes the operating environment required to use the Windows client.

# 2.1.1 Required Operating System

One of the following operating systems is required to use the Windows client:

- Windows(R) 8.1
- Windows(R) 8.1 Pro
- Windows(R) 8.1 Enterprise
- Windows(R) 10 Home
- Windows(R) 10 Education
- Windows(R) 10 Pro
- Windows(R) 10 Enterprise
- Microsoft(R) Windows Server(R) 2012 Datacenter
- Microsoft(R) Windows Server(R) 2012 Standard
- Microsoft(R) Windows Server(R) 2012 Essentials
- Microsoft(R) Windows Server(R) 2012 Foundation
- Microsoft(R) Windows Server(R) 2012 R2 Datacenter
- Microsoft(R) Windows Server(R) 2012 R2 Standard
- Microsoft(R) Windows Server(R) 2012 R2 Essentials
- Microsoft(R) Windows Server(R) 2012 R2 Foundation
- Microsoft Windows Server 2016 Datacenter
- Microsoft Windows Server 2016 Standard
- Microsoft Windows Server 2016 Essentials
- Microsoft Windows Server 2019 Datacenter
- Microsoft Windows Server 2019 Standard
- Microsoft Windows Server 2019 Essentials

Note: If Windows is 32 bit, only the Windows client (32 bit) can be installed.



The following components of Windows Server(R) 2016 and Windows Server(R) 2019 are not supported:

- Server Core
- Nano Server
- Windows Server Container

# 2.1.2 Related Software

The following table lists the software compatible (that can operate) with the Windows client.

Before using any of these, confirm that the OS supports the software.

Table 2.1 Related software

No.	Software name	Version
1	Visual Studio	2015 2017 2019
2	.NET Framework	4.5.1 or later 4.6/4.6.x 4.7/4.7.x 4.8
3	C compiler (*1)	-
4	JDK or JRE	Java SE 6 or later
5	NetCOBOL	(*2)
	NetCOBOL for .NET	(*3)

<sup>\*1:</sup> Only operations using the C compiler provided with the operating system are guaranteed.

\*2: NetCOBOL is available in the following editions:

- NetCOBOL Base Edition V10.5.0 or later
- NetCOBOL Standard Edition V10.5.0 or later
- NetCOBOL Professional Edition V10.5.0 or later
- NetCOBOL Enterprise Edition V10.5.0 or later

\*3: NetCOBOL for .NET is available in the following editions:

- NetCOBOL Standard Edition for .NET V5.0.0 or later
- NetCOBOL Professional Edition for .NET V5.0.0 or later
- NetCOBOL Enterprise Edition for .NET V5.0.0 or later



Either the JDK or JRE included with Interstage Application Server, or the Oracle JDK or JRE can be used. It is recommended that the JDK or JRE included with Interstage Application Server is used.

The following table lists servers that can be connected to the Windows client.

Table 2.2 Connectable servers

os	Software name		
Windows Linux	<ul> <li>FUJITSU Software Enterprise Postgres Advanced Edition 9.5 or later</li> <li>FUJITSU Software Enterprise Postgres Standard</li> </ul>		
	Edition 9.4 or later		
Solaris	- FUJITSU Software Enterprise Postgres Standard Edition 9.6 or later		

# 2.1.3 Excluded Software

**FUJITSU Enterprise Postgres** 

The Windows client cannot coexist with the FUJITSU Enterprise Postgres Community Edition (hereafter, "CE") client.

Other products

There are no exclusive products.

# 2.1.4 Hardware Environment

The following hardware is required to use the Windows client.

Memory

At least 256 MB of memory is required.

Mandatory hardware

None.

# 2.1.5 Disk Space Required for Installation

The disk space shown below is required for new installation of the Windows client.

At the Windows (32-bit) client installation destination:

414 MB

At the Windows (64-bit) client installation destination:

432 MB

System folder:

121 MB

# 2.1.6 Supported System Environment

This section describes the supported system environment.

# TCP/IP protocol

FUJITSU Enterprise Postgres supports version 4 and 6 (IPv4 and IPv6) of TCP/IP protocols.



Do not use link-local addresses if TCP/IP protocol version 6 addresses are used.

# File system

You can install FUJITSU Enterprise Postgres only if the system folder is an NTFS volume.

# 2.1.7 Versions of Open-Source Software Used as the Base for FUJITSU Enterprise Postgres Drivers

The following table lists the versions of open-source software used as the base for the various FUJITSU Enterprise Postgres drivers.

Driver	Open-source software version
.NET Data Provider	Npgsql 4.0.10
JDBC	pgjdbc 42.2.8

Driver	Open-source software version
ODBC	psqlodbc 12.01.0000
libpq	PostgreSQL 12.1

# 2.2 Installation

This section explains how to install the Windows client.



When installing the Fujitsu Enterprise Postgres Client (32 bit) in a 64 bit environment, do not specify a destination folder under the environment ProgramFiles variable.

# 2.2.1 Pre-installation Tasks

This section describes the operations to be performed before installing the Windows client.

# hosts file settings

Using an editor, add the IP address and server name of the connected server to the \System32\drivers\etc\hosts file in the folder where Windows is installed.



For Windows(R) 8.1, or Windows(R) 10, assign write privileges in advance to the user responsible for editing the hosts file.

## Check the installed product and determine the installation method

In Windows, click [All Programs] or [All apps], then [Fujitsu], and then [Uninstall (middleware)]. In the displayed window, check the installed products.

If the Windows client is already installed, determine the installation method to use:

- Reinstallation
- Multi-version installation

# Remove applied updates

If you perform reinstallation as the installation method, remove applied updates using the procedure shown below.



If a product is installed without removing applied updates, the following problems will occur:

- Performing reinstallation

If an update with the same update and version number is applied, an error informing you that the update has already been applied is displayed.

1. Display the applied updates

Execute the following command to display the applied updates:

C:\Program Files (x86)\Fujitsu\UpdateAdvisor\UpdateAdvisorMW>uam showup

# 2. Remove the updates

Execute the command below to remove the updates. If an update with the same update number was applied more than once, the updates are removed in order, starting from the highest version number.

C:\Program Files (x86)\Fujitsu\UpdateAdvisor\UpdateAdvisorMW>uam remove -i update-number

#### **Confirm service**

The Windows Installer service must be running.

# Confirm Path system environment variable

Ensure that the Path system environment variable includes "%SystemRoot%\system32".

# 2.2.2 Pre-installation Considerations

This section describes points that the user must take into account prior to installation.

- The remote desktop service is installed in application server mode, it is necessary to switch to install mode by executing the command shown below before installation. Also, after the installation is completed, execute the command shown below to switch back to execute mode.

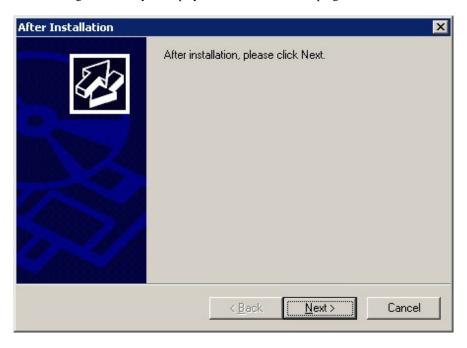
Before installation:

CHANGE USER /INSTALL

After installation:

CHANGE USER /EXECUTE

- The following window may be displayed when the installation program is executed.



If the above window is displayed, follow the procedure below:

- 1. Perform the steps in the installation procedure until the [InstallShield Wizard Complete] window is displayed.
- 2. Click [Next] in the [InstallShield Wizard Complete] window.

3. The following window is displayed. Click [Finish].



4. Restart the system.



- If the [User Account Control] dialog box is displayed when installation or uninstallation starts, click [Yes] to continue the operation.

If [No] is clicked, the [Error] dialog box will be displayed because continuing is not permitted. Click [Retry] in the displayed [Error] dialog box to proceed with installation or uninstallation. To end operations, click [Cancel].

- If installation or uninstallation is suspended or processing terminates abnormally, the [Program Compatibility Assistant] dialog box may be displayed.

If this happens, click [This program installed correctly] or [This program uninstalled correctly] and continue operation.

# 2.2.3 Installation in Interactive Mode

The installation procedure is described below.



For installation in interactive mode, default values are set for the installation information. The following settings can be changed for a new installation or a multi-version installation:

- Installation folder

# 1. Stop applications and programs

If the installation method is the following, all applications and programs that use the product must be stopped:

- Reinstallation

Before starting the installation, stop the following:

- Applications that use the product
- pgAdmin

#### 2. Insert the DVD

Insert the client program DVD into the DVD drive.

#### 3. Run the installation

The installation menu is displayed. Click [Installation].

# 4. Select the product to install

Select the product to install, and then click [Next].

If a selected product can only be reinstalled, refer to "6. Check the settings".



If multi-version installation is possible, the [Select Installation Method] window is displayed for each selected product. Select "Multi-version installation" and click [Next].

#### 5. Confirm the contents for installation

The [Confirm installation] window will be displayed.

Click [Next] to start the installation.

To modify the settings, select [Modify], and then click [Next]. Follow the on-screen instructions.

### 6. Check the settings

The [Confirm installation] window will be displayed for reinstallation, or if the installation information is modified.

Click [Install] to start the installation.

To modify the settings again, click [Back].

#### 7. Completion of installation

The completion window is displayed. Click [Finish].



In the cases below, refer to "2.4 Registering .NET Data Provider" for information on how to set up .NET Data Provider:

- .NET Framework is installed subsequently
- An issue occurs during installation
- A multi-version installation is performed

# 2.2.4 Installation in Silent Mode

Installation in silent mode can be performed only when the installation method is one of the following:

- New installation
- Multi-version installation



See

Refer to the FUJITSU Enterprise Postgres product website for information on installation in silent mode, such as the installation parameters and error messages.

The installation procedure is described below.

#### 1. Insert the DVD

Insert the client program DVD into the DVD drive.

The [Install Menu] window is displayed. Click [Finish].

# 2. Create an installation parameters CSV file

Consider the server type or features that will be required for system operations, and then create an installation parameters CSV file that uses the following specification format.

```
sectionName, parameterName, value
sectionName, parameterName, value
:
```



The template for the installation parameters CSV file is "Z:\sample\sample\_windows.csv (Z is the drive into which the DVD is inserted.).

## 3. Start the command prompt

In Windows, right-click [Command Prompt] and then select [Run as administrator].

#### 4. Run the installation

Execute the following command:

```
Z:\>silent.bat c:\temp\inspara.csv
```

*Z:* The drive into which the DVD is inserted.

c:\temp\inspara.csv: The installation parameter CSV file name.

If the silent installer ends in an error, a message is output to the log file and return values are returned.



In the cases below, refer to "2.4 Registering .NET Data Provider" for information on how to set up .NET Data Provider:

- .NET Framework is installed subsequently
- An issue occurs during installation
- A multi-version installation is performed

# 2.3 Uninstallation

This section describes the procedure for uninstalling the Windows client.



- Before uninstalling the product, close the product program and all applications that are using it.
- Log in using an account that has administrator privileges and then execute the command, or switch to an account that has administrator privileges and then uninstall the product.

# 2.3.1 Uninstallation in Interactive Mode

The uninstallation procedure is described below.



If an error occurs while the product is being uninstalled, refer to "Uninstall (middleware) Messages" in the FUJITSU Enterprise Postgres product website, and take the required action.

# 1. Stop applications and programs

Before starting the uninstallation, stop the following:

- Applications that use the product
- pgAdmin

# 2. Start the Uninstall (middleware) tool

In Windows, click [All Programs] or [All apps], then [Fujitsu], and then [Uninstall (middleware)].

# 3. Select the software

Select the product to be uninstalled from [Software Name], and then click [Remove].

#### 4. Start the uninstallation

Click [Uninstall].

#### 5. Finish the uninstallation

The uninstallation completion window is displayed. Click [Finish].

The installation folder may remain after uninstallation. If it is not required, delete it.

#### 6. Stop the Uninstall (middleware) tool

The [Uninstall (middleware)] window is displayed. Click [Close].

# 2.3.2 Uninstallation in Silent Mode

The uninstallation procedure is described below.



See

Refer to the FUJITSU Enterprise Postgres product website for information on uninstallation in silent mode, such as the error messages.

# 1. Stop applications and programs

Before starting the uninstallation, stop the following:

- Applications that use the product
- pgAdmin

# 2. Start the command prompt

In Windows, right-click [Command Prompt] and then select [Run as administrator].

#### 3. Start the uninstaller

Execute the command below.

The installation folder may remain after uninstallation. If it is not required, delete it.

# Example

X:> installFolder\suninst.bat

X: Drive on which the product is installed

# 2.4 Registering .NET Data Provider

This section explains how to register the .NET Data Provider feature in cases such as below, after installing the FUJITSU Enterprise Postgres client feature.

- .NET Framework is installed subsequently
- An issue occurs during installation of the FUJITSU Enterprise Postgres client feature
- Multiple versions of the FUJITSU Enterprise Postgres client feature are installed

In [Control Panel], select [Programs and Features], right-click [FUJITSU Npgsql .NET Data Provider GAC Register For .NET 4.5], and then select [Change] (if there is more than one listing, check the version column and select the version you wish to repair).

Operation	Description
Change	Install Npgsql Performance Counters or perform the Npgsql GAC Installation again.
	Click [Change], click the relevant item, select "Entire feature will be installed on local hard drive", and then click [Next] >> [Change] >> [Finish].
Repair	Repair a damaged installation.
	Click [Repair] >> [Repair] >> [Finish].

When you select [Change] and then select [Npgsql Installation], note that this consists of two parts:

- "Npgsql GAC only Installation"
- "Npgsql Publisher Policy and machine.config Installation"

If the FUJITSU Enterprise Postgres client feature is installed in an environment where no other versions exist

Perform "Npgsql GAC only Installation" and "Npgsql Publisher Policy and machine.config Installation".

If multiple versions of the FUJITSU Enterprise Postgres client feature are installed

The procedure differs depending on whether "FUJITSU.Npgsql" or another value is used for the invariant value.

Using "FUJITSU.Npgsql"

Perform "Npgsql GAC only Installation" and "Npgsql Publisher Policy and machine.config Installation".



- When "Npgsql Publisher Policy and machine.config Installation" is performed, an entry with the invariant value "FUJITSU.Npgsql" will be registered in the DbProviderFactories section of the machine.config file. All .NET applications that specify "FUJITSU.Npgsql" as an argument of GetFactory method of the DbProviderFactories class will now use the latest FUJITSU Enterprise Postgres Npgsql. Therefore, in a multi-version installation environment, carefully consider whether to use this option.
- If creating an application using Visual Studio tools when using the invariant value "FUJITSU.Npgsql" in a multiversion environment with FUJITSU Enterprise Postgres 9.6 or earlier, use the Visual Studio Integration Add-On

bundled with this version.

Refer to "Setting Up the Visual Studio Integration Add-On" in the Application Development Guide for details.

# Using values other than "FUJITSU.Npgsql"

When using a value other than "FUJITSU.Npgsql" for the invariant value, perform "Npgsql GAC only Installation" and "2.4.1 Registering Entries to the machine.config File".



- If you want to make a modification to use a value other than "FUJITSU.Npgsql" after performing the procedure to use "FUJITSU.Npgsql" for the invariant value, follow the procedure below:
  - In [Control Panel], select [Programs and Features], right-click [FUJITSU Npgsql .NET Data Provider GAC Register For .NET 4.5], and then select [Change]. If there is more than one listing, check the version column and select the version you wish to repair.
  - 2. Open the [Npgsql Installation] feature and deselect [Npgsql Publisher Policy and machine.config Installation].
  - 3. Select [Next] >> [Change] >> [Finish].
- If you want to use different versions of Npgsql for each application in a multi-version environment, use the NpgsqlRegister.exe tool provided. Refer to "2.4.1 Registering Entries to the machine.config File" for details.

# 2.4.1 Registering Entries to the machine.config File

If you want to use different versions of Npgsql for each application during multi-version installation, this can be achieved by registering multiple entries with different invariant values to DbProviderFactories in the machine.config file and using the relevant value for each application.

NpgsqlRegister.exe is a tool provided by FUJITSU.

The NpgsqlRegister.exe tool allows a user to:

- Change the invariant value registered to a different version of FUJITSU Enterprise Postgres to this version (cannot be edited with the exception of the default invariant value "FUJITSU.Npgsql").
- Register an invariant value.



- When changing existing invariant values, a user must edit existing DbProvider entries very carefully. Ensure that no application is using the DbProvider entries being edited.
- If it is likely that a user will install multiple versions of Npgsql, then it is recommended to add and use new invariant values.

The instructions below describe where to find and how to use this tool to register machine.config entries for the version being installed.

# Location of NpgsqlRegister.exe

The NpgsqlRegister.exe tool is stored in the following location.

fujitsuEnterprisePostgresClientInstallDir\DOTNET\NpgsqlRegister.exe

The NpgsqlRegister.exe tool requires .NET Framework 4.5.1 or later to be installed. If not installed, install .NET Framework 4.5.1 or later, and then execute the command again.

# Using NpgsqlRegister.exe

Execute the command below as a user with administrator privileges.

```
> NpgsqlRegister.exe
```

The NpgsqlRegister.exe tool checks if a FUJITSU Enterprise Postgres invariant value exists. If a FUJITSU Enterprise Postgres invariant value does not exist in machine.config, an invariant value with the current version (FUJITSU.Npgsql) is registered to machine.config, and this will be notified to the user. If a FUJITSU Enterprise Postgres invariant value exists, the user will be prompted with the following options.

```
1 - Edit an existing DbProvider entry to point to the current version
2 - Add a new Invariant alias DbProvider entry to point to the current version
3 - Exit
```

#### Editing an existing DbProvider entry

When this option is selected, applications that use an existing DbProvider entry will reference the latest installed Npgsql version from the GAC.

A list of all FUJITSU Enterprise Postgres Npgsql invariant name entries that exist in machine.config will be displayed as follows.

#### Example

```
Please select the Invariant attribute you would like to point to the current version.

1 - Invariant attribute: FUJITSU.Npgsql.11, Npgsql Version: 11.0.0.0

2 - Invariant attribute: FUJITSU.Npgsql.12, Npgsql Version: 12.0.0.0

3 - Invariant attribute: Npgsql, Npgsql Version: 12.0.0.0

4 - Exit
```

Selecting one of the invariant values listed will result in the machine.config entries being updated to point to this version, and any existing application referencing this invariant value will start referencing this version of Npgsql.

# Adding a new entry to DbProvider

When this option is selected, a new invariant value can be registered. When referencing the new registered invariant value from the application, this FUJITSU Enterprise Postgres server version of Npgsql is referenced.

Upon selecting the option to add a new invariant value, the following will be displayed.

```
You have chosen to add a new DbProvider entry with an Invariant Alias to the machine.config

The Invariant attribute can either be in the form of 2 user defined formats (define xyz):

1 - 'FUJITSU.Npgsql.xyz'

2 - 'Npgsql.xyz'

or add a DbProvider entry with Invariant attribute set to 'Npgsql'

3 - 'Npgsql'

4 - Exit

Please enter one of the options (1 .. 4):
```

By specifying a value to "xyz" in the first two options, any invariant value can be created. However, the user must specify at least one alphanumeric character of up to 10 characters.

The option three registers a new invariant value "Npgsql".

If the specified value already exists, a warning will be output as follows.

A DbProvider entry with Invariant attribute 'Npgsql' already exists in machine.config If you require this Invariant attribute, please restart the application and choose 'Edit an existing DbProvider entry'

# Chapter 3 Installation and Uninstallation of the Linux Client

This chapter explains how to install and uninstall the Linux client.

# 3.1 Operating Environment

This section describes the operating environment required to use the Linux client.

# 3.1.1 Required Operating System

One of the following operating systems is required to use the Linux client:

- RHEL7.2 or later
- RHEL8.1 or later
- SLES 12 SP2 or later



- The following packages are required for operations on RHEL7.

Package name	Remarks
glibc	-
glibc.i686	-
libgcc	-
libmemcached	Required when using Pgpool-II.
libstdc++	-
libtool-ltdl	-
ncurses-libs	-
nss-softokn-freebl	-
redhat-lsb	-
unixODBC	Required when using ODBC drivers
unzip	-
xz-libs	-
zlib	-

- The following packages are required for operations on RHEL8.1.

Package name	Remarks
glibc	-
glibc.i686	-
libnsl2	-
libgcc	-
libmemcached	Required when using Pgpool-II.
libstdc++	-
libtool-ltdl	-

Package name	Remarks
ncurses-libs	-
nss-softokn-freebl	-
redhat-lsb	-
unixODBC	Required when using ODBC drivers
unzip	-
xz-libs	-
zlib	-

- The following packages are required for operations on SLES 12.

Package name	Remarks
glibc	-
glibc-32bit	-
libgcc	-
libmemcached	Required when using Pgpool-II.
libstdc++	-
libtool-ltdl	-
ncurses-libs	-
nss-softokn-freebl	-
Oracle JRE 8	Required by installer Use update 31 or later of the 64-bit version.
unixODBC	Required when using ODBC drivers
unzip	-
xz-libs	-
zlib	-

# 3.1.2 Related Software

The following table lists the software required to use the Linux client.

Table 3.1 Related software

No.	Software name	Version
1	C compiler (*1)	-
2	JDK or JRE	Java SE 6 or later
3	NetCOBOL	(*2)

<sup>\*1:</sup> Only operations using the C compiler provided with the operating system are guaranteed.

# (RHEL8.1)

- NetCOBOL Base Edition V12.2.0 or later
- NetCOBOL Standard Edition V12.2.0 or later
- NetCOBOL Enterprise Edition V12.2.0 or later

(RHEL7)

<sup>\*2:</sup> NetCOBOL is available in the following editions:

- NetCOBOL Base Edition V11.1.0 or later
- NetCOBOL Standard Edition V11.1.0 or later
- NetCOBOL Enterprise Edition V11.1.0 or later



Either the JDK or JRE included with Interstage Application Server, or the Oracle JDK or JRE can be used. It is recommended that the JDK or JRE included with Interstage Application Server is used.

The following table lists servers that can be connected to the Linux client.

Table 3.2 Connectable servers

os	Software name
Windows Linux	- FUJITSU Software Enterprise Postgres Advanced Edition 9.5 or later
	- FUJITSU Software Enterprise Postgres Standard Edition 9.4 or later
Solaris	- FUJITSU Software Enterprise Postgres Standard Edition 9.6 or later

# 3.1.3 Excluded Software

**FUJITSU Enterprise Postgres** 

The Linux client cannot coexist with the FUJITSU Enterprise Postgres Community Edition (hereafter, "CE") client.

Other products

There are no exclusive products.

# 3.1.4 Required Patches

There are no required patches.

# 3.1.5 Hardware Environment

The following hardware is required to use the Linux client.

Memory

At least 160 MB of memory is required.

Mandatory hardware

None.

# 3.1.6 Disk Space Required for Installation

The following table lists the disk space requirements of the corresponding directories for new installation of the Linux client. If necessary, increase the size of the file system.

Disk space required for installation

**RHEL** 

Directory	Required disk space Unit: MB
/etc	1 + 1 (*1)

Directory	Required disk space Unit: MB
/var	1 + 2 (*1)
/opt	106 (*1)
Installation destination of the client (64-bit)	126
Installation destination of Pgpool-II	13

<sup>\*1:</sup> Uninstall (middleware) must be installed.

#### **SLES**

Directory	Required disk space Unit: MB
/etc	1 + 1 (*1)
/var	1 + 2 (*1)
/opt	2 (*1)
Installation destination of the client (64-bit)	121
Installation destination of Pgpool-II	10

<sup>\*1:</sup> Uninstall (middleware) must be installed.

# 3.1.7 Supported System Environment

This section describes the supported system environment.

# TCP/IP protocol

FUJITSU Enterprise Postgres supports version 4 and 6 (IPv4 and IPv6) of TCP/IP protocols.



Do not use link-local addresses if TCP/IP protocol version 6 addresses are used.

# 3.1.8 Versions of Open-Source Software Used as the Base for FUJITSU Enterprise Postgres Drivers

The following table lists the versions of open-source software used as the base for the various FUJITSU Enterprise Postgres drivers.

Driver	Open-source software version
JDBC	pgjdbc 42.2.8
ODBC	psqlodbc 12.01.0000
libpq	PostgreSQL 12.1

# 3.2 Installation

This section explains how to install the Linux client.

# 3.2.1 Pre-installation Tasks

Check the system environment for the following before the Linux client is installed.

# Check the disk capacity

Check if sufficient free disk space is available for installing the Linux client.

Refer to "Disk space required for installation" for information on disk space requirements.

If sufficient free disk space is unavailable, reconfigure disk partitions.

# Set JAVA\_HOME (SLES 12 only)

Ensure that Oracle JRE 8 is installed, and export the JAVA\_HOME environment variable.

```
#export JAVA_HOME="oracleJre8InstallDir"
```

# Check the installed product and determine the installation method

Using the operation shown below, start Uninstall (middleware), and check the installed products.

#### Example

```
# /opt/FJSVcir/cimanager.sh -c
Loading Uninstaller...

Currently installed products
1. productName
2. productName
Type [number] to select the software you want to uninstall.
[number,q]
=>q

Exiting Uninstaller.
```

If the Linux client is already installed, determine the installation method to use:

- Reinstallation
- Multi-version installation

# Remove applied updates

If you perform reinstallation as the installation method, remove applied updates using the procedure shown below.



If a product is installed without removing applied updates, the following will occur:

- Performing reinstallation

If an update with the same update and version number is applied, an error informing you that the update has already been applied is displayed.

Perform the reinstallation after removing the update.

1. Display the applied updates

Execute the following command to display the applied updates:

```
# /opt/FJSVfupde/bin/uam showup
```

# 2. Remove the updates

Execute the command below to remove the updates. If an update with the same update number was applied more than once, the updates are removed in order, starting from the highest version number.

# 3.2.2 Installation in Interactive Mode

The installation procedure is described below.



The following characters can be used as input values:

Alphanumeric characters, hyphens, commas and forward slashes



For installation in interactive mode, default values are set for the installation information. The following settings can be changed for a new installation or a multi-version installation:

- Installation directory

# 1. Stop applications and programs

If the installation method is the following, all applications and programs that use the product must be stopped:

- Reinstallation

Before starting the installation, stop the following:

- Applications that use the product
- pgBadger
- Pgpool-II

# 2. Change to the superuser

Run the following command to switch to the superuser on the system.

```
$ su -
Password:*****
```

# 3. Mount the DVD drive

Insert the client program DVD into the DVD drive, and then execute the following command:

Example

```
# mount -t iso9660 -r -o loop /dev/dvd /media/dvd
```

Here /dev/dvd is the device name for the DVD drive (which may vary depending on your environment), and /media/dvd is the mount point (which may need to be created before calling the command).



If the DVD was mounted automatically using the automatic mount daemon (autofs), "noexec" is set as the mount option, so the installer may fail to start. In this case, use the mount command to remount the DVD correctly, and then run the installation. Note that the mount options of a mounted DVD can be checked by executing the mount command without any arguments.

#### 4. Run the installation

Execute the following command:

## Example

# cd /media/dvd
# ./install.sh

In the example above, /media/dvd is the DVD mount point.

## 5. Select the product to install

The list of installation target products is displayed.

Type the number for the product to be installed, or "all", and press Enter.



If the selected product has already been installed, a window for selecting reinstallation or multi-version installation is displayed for each product. Follow the on-screen instructions to select the installation method.

# 6. Check the settings

The window for checking the installation information is displayed as shown below.

Type "y" and press Enter to start the installation.

To change the settings, type "c" and press Enter, and follow the on-screen instructions. This option is not displayed if there is no information that can be modified.

#### 7. Check the changed settings

If the installation information has been modified, the new installation information is displayed.

Type "y" and press Enter to start the installation.

To change the settings again, type "c" and press Enter.

#### 8. Finish the installation

When the installation is complete, a message is displayed showing the status.

If installation was successful, the following message is displayed:

Installed successfully.



If an error occurs during the installation, read the error message and remove the cause of the error, and then reexecute the install.sh command.

# 3.2.3 Installation in Silent Mode

Installation in silent mode can be performed only when the installation method is one of the following:

- New installation

- Multi-version installation



Refer to the FUJITSU Enterprise Postgres product website for information on installation in silent mode, such as the installation parameters and error messages.

......

The installation procedure is described below.

# 1. Change to the superuser

Run the following command to switch to the superuser on the system.

```
$ su -
Password:*****
```

#### 2. Mount the DVD drive

Insert the client program DVD into the DVD drive, and then execute the following command:

# Example

```
# mount -t iso9660 -r -o loop /dev/dvd /media/dvd
```

Here /dev/dvd is the device name for the DVD drive (which may vary depending on your environment), and /media/dvd is the mount point (which may need to be created before calling the command).



If the DVD was mounted automatically using the automatic mount daemon (autofs), "noexec" is set as the mount option, so the installer may fail to start. In this case, use the mount command to remount the DVD correctly, and then run the installation. Note that the mount options of a mounted DVD can be checked by executing the mount command without any arguments.

#### 3. Create an installation parameters CSV file

Consider the features that will be required for system operations, and then create an installation parameters CSV file that uses the following specification format.

```
sectionName, parameterName, value
sectionName, parameterName, value
:
```



The templates for the installation parameters CSV file is "mountPoint/sample/sample\_linux.csv"

#### 4. Run the installation

Execute the following command:

# Example

```
# cd /media/dvd
# ./silent.sh /home/work/inspara.csv
```

In the example above, /media/dvd is the DVD mount point, and /home/work/inspara.csv is the installation parameter CSV.

If the silent installer ends in an error, a message is output to the log file and return values are returned.

# 3.3 Uninstallation

This section describes the procedure for uninstalling the Linux client.



- Before uninstalling the product, close the product program and all applications that are using it.
- On SLES 12, before uninstallation, ensure that Oracle JRE 8 is installed, and export the JAVA\_HOME environment variable.

#export JAVA\_HOME="oracleJre8InstallDir"

# 3.3.1 Uninstallation in Interactive Mode

The uninstallation procedure is described below.



If an error occurs while the product is being uninstalled, refer to "Uninstall (middleware) Messages" in the FUJITSU Enterprise Postgres product website, and take the required action.

# 1. Stop applications and programs

Before starting the uninstallation, stop the following:

- Applications that use the product
- pgBadger
- Pgpool-II

# 2. Change to the superuser

Run the following command to switch to the superuser on the system.

\$ su Password:\*\*\*\*\*

# 3. Start the Uninstall (middleware)

Execute the following command:

# /opt/FJSVcir/cimanager.sh -c

# 4. Select the product

Enter the number for the product to be uninstalled, and press Enter.

# 5. Start the uninstallation

To start the uninstallation, type "y" and press Enter.

To display the list of products again, type "b" and press Enter.

# 6. Finish the uninstallation

If uninstallation is successful, the message below is displayed.

The installation directory may remain after uninstallation. If it is not required, delete it.

# 3.3.2 Uninstallation in Silent Mode

The uninstallation procedure is described below.



See

Refer to the FUJITSU Enterprise Postgres product website for information on uninstallation in silent mode, such as the error messages.

# 1. Stop applications and programs

Before starting the uninstallation, stop the following:

- Applications that use the product
- pgBadger
- Pgpool-II

# 2. Change to the superuser

Run the following command to switch to the superuser on the system.

```
$ su -
Password:****
```

#### 3. Run the uninstallation

Execute the following command:

The installation directory may remain after uninstallation. If it is not required, delete it.

#### Example

```
# /opt/fsepv<x>client64/setup/suninst.sh
```

In the example above, /opt/fsepv<x>client64 is the installation directory, and "<x>" indicates the product version.

```
# /opt/fsepv<x>pgpool-II/setup/suninst.sh
```

In the example above, /opt/fsepv<x>pgpool-II/setup is the name of the installation directory in which the Pgpool-II is installed.

# Chapter 4 Setup

This chapter describes the setup procedures to be performed after installation completes.

# 4.1 Configuring Environment Variables

Configure the following environment variables when using client commands.

# W 4.1.1 Windows

PATH environment variable

Add "installationFolder\bin".

PGLOCALEDIR environment variable

Add "installationFolder\share\locale".

Examples of environment variable configurations are shown below.

#### Example

This example is specific to 32-bit Windows.

Note that "<x>" indicates the product version.

- > SET PATH=%ProgramFiles%\Fujitsu\fsepv<x>client32\bin;%PATH%
- > SET PGLOCALEDIR=%ProgramFiles%\Fujitsu\fsepv<x>client32\share\locale

# 4.1.2 Linux

PATH environment variable

Add "installationDirectory/bin".

MANPATH environment variable

Add "installationDirectory/share/man".

PGLOCALEDIR environment variable

Add "installationDirectory/share/locale".

LD\_LIBRARY\_PATH environment variable

Add "installationDirectory/lib".

Examples of environment variable configurations are shown below.

# Example

This example is specific to 64-bit Linux.

Note that "<*x*>" indicates the product version.

- \$ PATH=/opt/fsepv<x>client64/bin:\$PATH ; export PATH
- \$ MANPATH=/opt/fsepv<x>client64/share/man:\$MANPATH ; export MANPATH
- \$ PGLOCALEDIR=/opt/fsepv<x>client64/share/locale ; export PGLOCALEDIR
- \$ LD\_LIBRARY\_PATH=/opt/fsepv<x>client64/lib:\$LD\_LIBRARY\_PATH ; export LD\_LIBRARY\_PATH

# Index

[C]	
Check the disk capacity	
Configuring Environment Variables	
Confirm service	
Connectable servers4	ŀ
[D]	
Disk Space Required for Installation5,17	7
713K Space Required for installation	
[E]	
Excluded Software5,17	7
[F]	
Features that can be Installed	
ile system5	5
r⊔ı	
[H]	7
Hardware Environment	
osis me seungs	)
[I]	
nstallation6	5
nstallation and Uninstallation of the Linux Client15	
nstallation and Uninstallation of the Windows Client 3	
nstallation in Interactive Mode2,8,11	
nstallation in Silent Mode2,21	
nstallation Procedure1	ĺ
nstallation Types1	
[L]	
[L] .D_LIBRARY_PATH environment variable25	
.D_LIBRARY_PATH environment variable25	
D_LIBRARY_PATH environment variable25	5
D_LIBRARY_PATH environment variable	5
D_LIBRARY_PATH environment variable25	5
D_LIBRARY_PATH environment variable	5
ID_LIBRARY_PATH environment variable	5
[M]  MANPATH environment variable	5
[M]  MANPATH environment variable	5
[M]  MANPATH environment variable	5
[M] MANPATH environment variable	5
[M] MANPATH environment variable	5
Manual	5 5 1
Manual   M	55
Manual   M	55
Manual   M	55
Manual   M	55
Manual   M	55
Manual   M	55
M	55
Manual   M	55 11 15 55 55 55 55 55 55 55 55 55 55 5
Manual   M	55 11 15 55 55 55 55 55 55 55 55 55 55 5
Manual   M	55 55 11 11 55 55 77
Manual   M	55 55 57 77 55

[T]	£ 10
TCP/IP protocol	5,18
[U]	
Uninstallation	2,10,23
Uninstallation in Interactive Mode	23
Uninstallation in Silent Mode	11,24



# FUJITSU Enterprise Postgres 12



# Installation and Setup Guide for Server Assistant

# **Preface**

# Purpose of this document

This document describes how to install and uninstall the FUJITSU Enterprise Postgres Server Assistant.

#### Intended readers

This document is intended for those who install and operate FUJITSU Enterprise Postgres.

Readers of this document are assumed to have general knowledge of:

- PostgreSQL
- SQL
- Linux
- w
- PostgreSQL
- SQL
- Windows

#### Structure of this document

This document is structured as follows:

# Chapter 1 Overview of Installation

Describes the features that can be installed, and provides an overview of installation methods

## Chapter 2 Installation and Uninstallation of the Linux Server Assistant

Describes how to install and uninstall the Linux Server Assistant

# Chapter 3 Installation and Uninstallation of the Windows Server Assistant

Describes how to install and uninstall the Windows Server Assistant

# Chapter 4 Setup of the Server Assistant

Describes the setup to be performed after installation

# Appendix A Estimating Memory Requirements

Describes the formulas for estimating memory requirements

Appendix B Procedure when Modifying the Oracle JRE Installation

Describes the procedure to follow when modifying the Oracle JRE installation.

# **Export restrictions**

Exportation/release of this document may require necessary procedures in accordance with the regulations of your resident country and/or US export control laws.

# Issue date and version

Edition 1.0: July 2020

# Copyright

Copyright 2018-2020 FUJITSU LIMITED

# Contents

Index	25
Appendix B Procedure when Modifying the Oracle JRE Installation	24
A.1 Server Assistant Memory Requirements	
Appendix A Estimating Memory Requirements	
Chapter 4 Setup of the Server Assistant	
3.3.2 Uninstallation in Silent Mode	
3.3.1 Uninstallation in Interactive Mode	
3.3 Uninstallation.	
3.2.3 Installation in Silent Mode	
3.2.2 Installation in Interactive Mode	
3.2.1 Pre-installation Tasks.	
3.2 Installation.	
3.1.7 Supported System Environment	
3.1.6 Disk Space Required for Installation.	
3.1.5 Hardware Environment.	
3.1.4 Required Patches	
3.1.2 Related Software	
3.1.1 Required Operating System.	
3.1 Operating Environment	
Chapter 3 Installation and Uninstallation of the Windows Server Assistant	
2.3.2 Uninstallation in Silent Mode	
2.3 Uninstallation in Interactive Mode	
2.3 Uninstallation.	
2.2.2 Installation in Interactive Mode	
2.2.1 Pre-installation Tasks	
2.2 Installation.	
2.1.7 Supported System Environment	
2.1.6 Disk Space Required for Installation.	
2.1.5 Hardware Environment.	
2.1.4 Required Patches.	
2.1.3 Excluded Software	
2.1.2 Related Software	
2.1.1 Required Operating System	
2.1 Operating Environment	
Chapter 2 Installation and Uninstallation of the Linux Server Assistant	
1.7 Omnounauon	
1.4 Uninstallation.	
1.3.2 Installation in Silent Mode	
1.3 Installation Procedure	
1.2.3 Multi-Version Installation.	
1.2.2 Reinstallation.	
1.2.1 New Installation.	
1.2 Installation Types.	
1.1 Features that can be Installed	
Chapter 1 Overview of Installation	

## Chapter 1 Overview of Installation

This chapter provides an overview of FUJITSU Enterprise Postgres Server Assistant installation and uninstallation.

#### 1.1 Features that can be Installed

The Server Assistant is provided as a Server Assistant package, which is installed on a different server (referred to as the arbitration server) to that of the database server.

The following table shows the relationship between the platforms and the features provided by the Server Assistant package.

Platform	Server Assistant
Linux	Y
Windows	Y

Y: Can be installed

## 1.2 Installation Types

The following installation types are available for FUJITSU Enterprise Postgres:

- New installation
- Reinstallation
- Multi-version installation

#### 1.2.1 New Installation

In initial installation, the FUJITSU Enterprise Postgres Server Assistant is installed for the first time.

#### 1.2.2 Reinstallation

Perform reinstallation to repair installed program files that have become unusable for any reason.

#### 1.2.3 Multi-Version Installation

FUJITSU Enterprise Postgres products can be installed on the same server if the product version (indicated by "x" in "x SPz") is different from that of any version of the product that is already installed.

#### 1.3 Installation Procedure

The following installation procedures are available for FUJITSU Enterprise Postgres:

- Installation in interactive mode
- Installation in silent mode

Select the installation procedure that corresponds to your environment.

#### 1.3.1 Installation in Interactive Mode

Interactive mode enables installation to be performed while the required information is entered interactively.

In the interactive mode installation, the FUJITSU Enterprise Postgres installer automatically determines the installation state of FUJITSU Enterprise Postgres. Install FUJITSU Enterprise Postgres using one of the following installation types in accordance with the installation state:

- New installation

- Reinstallation
- Multi-version installation

## 1.3.2 Installation in Silent Mode

Silent mode enables installation to be performed without the need to enter any information interactively.

New installations and multi-version installations can be performed in silent mode.

## 1.4 Uninstallation

Uninstallation removes the system files of the installed FUJITSU Enterprise Postgres Server Assistant.

## Chapter 2 Installation and Uninstallation of the Linux Server Assistant

This chapter explains how to install and uninstall the Linux Server Assistant.

## 2.1 Operating Environment

This section describes the operating environment required in order to use the Linux Server Assistant.

## 2.1.1 Required Operating System

One of the following operating systems is required in order to use the Linux Server Assistant:

- RHEL7.2 or later
- RHEL8.1 or later
- SLES 12 SP2 or later



- The following packages are required for operations on RHEL7.

Package name	Remarks
glibc	-
glibc.i686	-
libgcc	-
libstdc++	-
ncurses-libs	-
nss-softokn-freebl	-
redhat-lsb	-
unzip	-
xz-libs	-
zlib	-

- The following packages are required for operations on RHEL8.1.

Package name	Remarks
glibc	-
glibc.i686	-
libgcc	-
libstdc++	-
ncurses-libs	-
nss-softokn-freebl	-
redhat-lsb	-
unzip	-
xz-libs	-
zlib	-

- The following packages are required for operations on SLES 12.

Package name	Remarks
glibc	-
glibc-32bit	-
libgcc	-
libstdc++	-
ncurses-libs	-
nss-softokn-freebl	-
Oracle JRE 8	Required when using the installer or Server Assistant. Use update 31 or later of the 64-bit version.
unzip	-
xz-libs	-
zlib	-

#### 2.1.2 Related Software

No other software is required in order to use FUJITSU Enterprise Postgres.

The following table lists servers that can be connected to the Linux Server Assistant.

Table 2.1 Connectable servers

os	Software name
Linux	- FUJITSU Software Enterprise Postgres Standard Edition 11 or later
Windows	- FUJITSU Software Enterprise Postgres Advanced Edition 11 or later

#### 2.1.3 Excluded Software

There is no excluded software.

## 2.1.4 Required Patches

There are no required patches.

#### 2.1.5 Hardware Environment

The following hardware is required in order to use the Linux Server Assistant:

Memory

At least 150 MB of memory is required.

Mandatory hardware

None.

## 2.1.6 Disk Space Required for Installation

The following table lists the disk space requirements of the corresponding directories for new installation of the Linux Server Assistant. If necessary, increase the size of the file system.

Disk space required for installation

RHEL

Directory	Required disk space Unit: MB	
/etc	1 + 1 (*1)	
/var	1 + 2 (*1) +1 (*2)	
/opt	106 (*1) +1 (*2)	
serverAssistantInstallDir	139	

<sup>\*1:</sup> Uninstall (middleware) must be installed.

#### **SLES**

Directory	Required disk space Unit: MB
/etc	1 + 1 (*1)
/var	1 + 2 (*1) +1 (*2)
/opt	2 (*1) +1 (*2)
serverAssistantInstallDir	2

<sup>\*1:</sup> Uninstall (middleware) must be installed.

### 2.1.7 Supported System Environment

This section describes the supported system environment.

#### **TCP/IP Protocol**

FUJITSU Enterprise Postgres supports version 4 and 6 (IPv4 and IPv6) of TCP/IP protocols.



Do not use link-local addresses if TCP/IP protocol version 6 addresses are used.

### 2.2 Installation

This section describes how to install the Linux Server Assistant.

#### 2.2.1 Pre-installation Tasks

Check the following system environment before installing the Linux Server Assistant.

#### Check the disk space

Ensure that there is sufficient disk space to install the Linux Server Assistant.

Refer to "2.1.6 Disk Space Required for Installation" for information on disk space requirements.

If sufficient free disk space is unavailable, reconfigure disk partitions.

#### Set JAVA\_HOME (SLES 12 only)

Ensure that Oracle JRE 8 is installed, and export the JAVA\_HOME environment variable.

#export JAVA\_HOME="oracleJre8InstallDir"

<sup>\*2:</sup> Required for the installation of FJQSS.

<sup>\*2:</sup> Required for the installation of FJQSS.

Refer to "Appendix B Procedure when Modifying the Oracle JRE Installation" for information on modifying Oracle JRE after installation.

#### Check the installed product and determine the installation method

Using the operation shown below, start Uninstall (middleware), and check the installed products.

#### Example

```
# /opt/FJSVcir/cimanager.sh -c
Loading Uninstaller...

Currently installed products
1. productName

Type [number] to select the software you want to uninstall.
[number,q]
=>q

Exiting Uninstaller.
```

If the Linux Server Assistant is already installed, determine the installation method to use:

- Reinstallation
- Multi-version installation

#### Remove applied updates

If you perform reinstallation, remove applied updates using the procedure shown below.



If a product is installed without removing applied updates, the following will occur:

- Performing reinstallation

If an update with the same update and version number is applied, an error informing you that the update has already been applied is displayed.

Perform reinstallation after removing the update.

#### 1. Display the applied updates

Execute the following command to display the applied updates:

```
# /opt/FJSVfupde/bin/uam showup
```

#### 2. Remove the updates

Execute the command below to remove the updates. If an update with the same update number was applied more than once, the updates are removed in order, starting from the highest version number.

```
# /opt/FJSVfupde/bin/uam remove -i update-number
```

#### 2.2.2 Installation in Interactive Mode

The installation procedure is described below.



The following characters can be used as input values:

Alphanumeric characters, hyphens and forward slashes



In interactive mode installation, the default values of the installation information are set. The following setting can be changed for a new installation or a multi-version installation.

- Installation directory

#### 1. Stop the program

If the installation method is the following, the program must be stopped:

- Reinstallation

Before starting the installation, stop the following:

- Mirroring Controller arbitration process

Execute the mc\_arb command in stop mode to stop the Mirroring Controller arbitration process.

Example

```
$ mc_arb stop -M /mcarb_dir/arbiter1
```

#### 2. Change to the superuser

Run the following command to switch to the superuser on the system.

```
$ su -
Password:****
```

#### 3. Mount the DVD drive

Insert the Server Assistant program DVD into the DVD drive, and then execute the following command:

Example

```
# mount -t iso9660 -r -o loop /dev/dvd /media/dvd
```

Here /dev/dvd is the device name for the DVD drive (which may vary depending on your environment), and /media/dvd is the mount point (which may need to be created before calling the command).



If the DVD was mounted automatically using the automatic mount daemon (autofs), "noexec" is set as the mount option, so the installer may fail to start. In this case, use the mount command to remount the DVD correctly, and then run the installation. Note that the mount options of a mounted DVD can be checked by executing the mount command without any arguments.

#### 4. Run the installation

Execute the following command:

Example

```
# cd /media/dvd
# ./install.sh
```

In the example above, /media/dvd is the DVD mount point.

#### 5. Select the product to install

The list of installation target products is displayed.

Type the number for the product to be installed, or "all", and press Enter.



If the selected product has been installed, a window for selecting reinstallation or multi-version installation is displayed for each product. Follow the on-screen instructions to select the installation method.

#### 6. Check the settings

The window for checking the installation information is displayed.

Type "y" and press Enter to start the installation.

To change the settings, type "c", press Enter, and follow the on-screen instructions. This option is not displayed if there is no information that can be modified.

#### 7. Check the changed settings

If the settings have been changed, the new installation information is displayed.

Type "y" and press Enter to start the installation.

To change the settings again, type "c" and press Enter.

#### 8. Finish the installation

When the installation is complete, a message is displayed showing the status.

If installation was successful, the following message is displayed:

Installed successfully.



If an error occurs during the installation, read the error message and remove the cause of the error, and then reexecute the install.sh command.

#### 2.2.3 Installation in Silent Mode

Installation in silent mode can be performed only when the installation method is one of the following:

- New installation
- Multi-version installation



See

Refer to the FUJITSU Enterprise Postgres product website for information on installation in silent mode, such as the installation parameters and error messages.

The installation procedure is described below.

#### 1. Change to the superuser

Run the following command to switch to the superuser on the system.

```
$ su -
Password:*****
```

#### 2. Mount the DVD drive

Insert the Server Assistant program DVD into the DVD drive, and then execute the following command:

#### Example

```
# mount -t iso9660 -r -o loop /dev/dvd /media/dvd
```

Here /dev/dvd is the device name for the DVD drive (which may vary depending on your environment), and /media/dvd is the mount point (which may need to be created before calling the command).



If the DVD was mounted automatically using the automatic mount daemon (autofs), "noexec" is set as the mount option, so the installer may fail to start. In this case, use the mount command to remount the DVD correctly, and then run the installation. Note that the mount options of a mounted DVD can be checked by executing the mount command without any arguments.

#### 3. Create an installation parameters CSV file

Create an installation parameters CSV file that uses the following specification format.

```
sectionName, parameterName, value sectionName, parameterName, value
```



The templates for the installation parameters CSV file is "mountPoint/sample/sample.csv"

#### 4. Run the installation

Execute the following command:

#### Example

```
# cd /media/dvd
# ./silent.sh /home/work/inspara.csv
```

In the example above, /media/dvd is the DVD mount point, and /home/work/inspara.csv is the installation parameter CSV

If the silent installer ends in an error, a message is output to the log file and return values are returned.

## 2.3 Uninstallation

This section describes how to uninstall the Linux Server Assistant.



Before uninstalling the product, close the product program.

#### 2.3.1 Uninstallation in Interactive Mode

The uninstallation procedure is described below.



If an error occurs while the product is being uninstalled, refer to "Uninstall (middleware) Messages" in the FUJITSU Enterprise Postgres product website, and take the required action.

#### 1. Stop the program

Before starting the uninstallation, stop the following:

- Mirroring Controller arbitration process

Execute the mc\_arb command in stop mode to stop the Mirroring Controller arbitration process.

Example

```
$ mc_arb stop -M /mcarb_dir/arbiter1
```

#### 2. Change to the superuser

Run the following command to switch to the superuser on the system.

```
$ su -
Password:****
```

#### 3. Start Uninstall (middleware)

Execute the following command:

```
# /opt/FJSVcir/cimanager.sh -c
```

#### 4. Select the product

Enter the number for the product to be uninstalled, and press Enter.

#### 5. Start the uninstallation

To start the uninstallation, type "y" and press Enter.

To display the list of products again, type "b" and press Enter.

#### 6. Finish the uninstallation

If uninstallation is successful, the message below is displayed.

The installation directory may remain after uninstallation. If it is not required, delete it.

#### 2.3.2 Uninstallation in Silent Mode

The uninstallation procedure is described below.



Refer to the FUJITSU Enterprise Postgres product website for information on uninstallation in silent mode, such as the error messages.

......

#### 1. Stop the program

Before starting the uninstallation, stop the following:

- Mirroring Controller arbitration process

Execute the mc\_arb command in stop mode to stop the Mirroring Controller arbitration process.

Example

\$ mc\_arb stop -M /mcarb\_dir/arbiter1

#### 2. Change to the superuser

Run the following command to switch to the superuser on the system.

```
$ su -
Password:*****
```

#### 3. Run the uninstallation

Execute the following command:

The installation directory may remain after uninstallation. If it is not required, delete it.

#### Example

```
# /opt/fsepv<x>assistant/setup/suninst.sh
```

In the example above, /opt/fsepv<*x*>assistant is the installation directory, and "<*x*>" indicates the product version.

## Chapter 3 Installation and Uninstallation of the Windows Server Assistant

This chapter explains how to install and uninstall the Windows Server Assistant.

## 3.1 Operating Environment

This section describes the operating environment required in order to use the Windows Server Assistant.

### 3.1.1 Required Operating System

One of the following operating systems is required in order to use the Windows Server Assistant:

- Microsoft(R) Windows Server(R) 2012 Datacenter
- Microsoft(R) Windows Server(R) 2012 Standard
- Microsoft(R) Windows Server(R) 2012 Essentials
- Microsoft(R) Windows Server(R) 2012 Foundation
- Microsoft(R) Windows Server(R) 2012 R2 Datacenter
- Microsoft(R) Windows Server(R) 2012 R2 Standard
- Microsoft(R) Windows Server(R) 2012 R2 Essentials
- Microsoft(R) Windows Server(R) 2012 R2 Foundation
- Microsoft(R) Windows Server(R) 2016 Datacenter
- Microsoft(R) Windows Server(R) 2016 Standard
- Microsoft(R) Windows Server(R) 2016 Essentials
- Microsoft(R) Windows Server(R) 2019 Datacenter
- Microsoft(R) Windows Server(R) 2019 Standard
- Microsoft(R) Windows Server(R) 2019 Essentials



- The following features of Windows Server(R) 2016 and Windows Server(R) 2019 are not supported:
  - Server Core
  - Nano Server
  - Windows Server Container
- The TCP/IP protocol must be installed.

#### 3.1.2 Related Software

There is no software related to FUJITSU Enterprise Postgres.



- The following programs are installed during installation of FUJITSU Enterprise Postgres:

- Microsoft Visual C++ 2010 Redistributable version 10.0.30319.01

- Microsoft Visual C++ 2013 Redistributable version 12.0.40649.5
- Microsoft Visual C++ 2015-2019 Redistributable version 14.24.28127.4

Do not uninstall the above programs as they are required for running FUJITSU Enterprise Postgres.

The following table lists servers that can be connected to the Windows Server Assistant.

Table 3.1 Connectable servers

os	Software name
Linux	FUJITSU Software Enterprise Postgres Standard Edition 11 or later
Windows	FUJITSU Software Enterprise Postgres Advanced Edition 11 or later

#### 3.1.3 Excluded Software

There is no excluded software.

### 3.1.4 Required Patches

There are no required patches.

#### 3.1.5 Hardware Environment

The following hardware is required in order to use the Windows Server Assistant:

Memory

At least 256 MB of memory is required.

Mandatory hardware

None.

## 3.1.6 Disk Space Required for Installation

The disk space shown below is required for new installation of the Windows Server Assistant.

Windows Server Assistant installation destination:

256 MB

System folder:

138 MB

## 3.1.7 Supported System Environment

This section describes the supported system environment.

#### TCP/IP protocol

FUJITSU Enterprise Postgres supports version 4 and 6 (IPv4 and IPv6) of TCP/IP protocols.



Do not use link-local addresses if TCP/IP protocol version 6 addresses are used.

#### File system

You can install FUJITSU Enterprise Postgres only if the system folder is an NTFS volume.

## 3.2 Installation

This section describes how to install the Windows Server Assistant.



- The installation must be performed by a user with administrator privileges (a user ID that belongs to the Administrators group).

............

- Stop all applications before starting the installation.
- The Windows Installer service must be running.
- If antivirus software is installed, the system may crash, fail to start, or stop responding during installation or when starting after installation. Set the scan exclusions so that the installation directory and the resource allocation directory are not scanned for viruses.
- The remote desktop service is installed in application server mode, it is necessary to switch to install mode by executing the command shown below before installation. Also, after the installation is completed, execute the command shown below to switch back to execute mode.

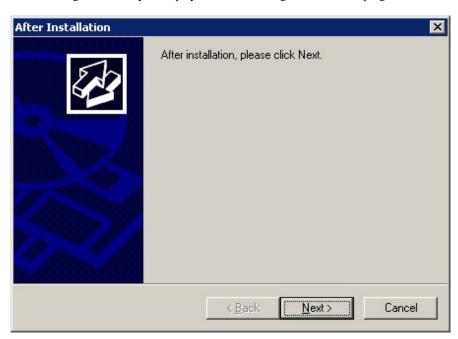
[Before the installation]

CHANGE USER /INSTALL

[After the installation]

CHANGE USER /EXECUTE

- The following window may be displayed when executing the installation program:



If this window is displayed, perform the following operations:

- 1. Perform the installation steps until the [InstallShield Wizard Complete] window is displayed.
- 2. At the window shown above, click [Next].

3. The window shown below is displayed. Click [Finish].





- If a [User Account Control] dialog box is displayed at the start of the installation, click [Yes] to continue processing: If [No] is clicked, permission to continue is denied and an [Error] dialog box will be displayed.

To continue the installation, click [Retry] at the [Error] dialog box. To end the installation, click [Cancel].

••••••••••

- If installation is suspended or processing terminates abnormally, the [Program Compatibility Assistant] dialog box may be displayed. Click [This program installed correctly] and continue operation.

#### 3.2.1 Pre-installation Tasks

This section describes the tasks to be performed before installing the Windows Server Assistant.

#### Check the disk space

Ensure that there is sufficient disk space to install the Windows Server Assistant.

Refer to "3.1.6 Disk Space Required for Installation" for information on disk space requirements.

#### hosts file settings

Using an editor, add the IP address and name of the server to be connected to the \System32\drivers\etc\hosts file in the folder where Windows is installed.

#### Check the installed product and determine the installation method

Using the operation shown below, start Uninstall (middleware), and check the installed products.

In Windows, click [All Programs] or [All apps], then [Fujitsu], and then [Uninstall (middleware)].

If the Windows Server Assistant is already installed, determine the installation method to use:

- Reinstallation
- Multi-version installation

#### Remove applied updates

If you perform reinstallation as the installation method, remove applied updates using the procedure shown below.



If a product is installed without removing applied updates, the following will occur:

- Performing reinstallation

If an update with the same update and version number is applied, an error informing you that the update has already been applied is displayed.

.....

#### 1. Display the applied updates

Execute the following command to display the applied updates:

C:\>uam showup

#### 2. Remove the updates

Execute the command below to remove the updates. If an update with the same update number was applied more than once, the updates are removed in order, starting from the highest version number.

C:\>uam remove -i update-number

#### 3.2.2 Installation in Interactive Mode

The installation procedure is described below.



In interactive mode installation, the default values of the installation information are set. The following setting can be changed for a new installation or a multi-version installation.

- Installation folder

#### 1. Stop the program

If the installation method is the following, the program must be stopped:

- Reinstallation

Before starting the installation, stop the following:

- Mirroring Controller arbitration process

 $Execute \ the \ mc\_arb \ command \ in \ stop \ mode \ to \ stop \ the \ Mirroring \ Controller \ arbitration \ process.$ 

Example

> mc\_arb stop -M D:\mcarb\_dir\arbiter1

#### 2. Insert the DVD

Insert the Server Assistant program DVD into the DVD drive.

#### 3. Run the installation

The installation menu is displayed. Click [Installation].

#### 4. Select the product to install

The [Installation product] window is displayed.

Select the product to install, and then click [Next].

If a selected product can only be reinstalled, refer to "6. Check the settings".



If the selected product has been installed, a window for selecting reinstallation or multi-version installation is displayed for each product. Follow the on-screen instructions to select the installation method.

#### 5. Confirm the contents for installation

The [Confirm installation] window will be displayed.

Click [Next] to start the installation.

To modify the settings, select [Modify], and then click [Next]. Follow the on-screen instructions.

#### 6. Check the settings

The [Confirm installation] window will be displayed. Click [Install] to start the installation.

To modify the settings again, click [Back].

#### 7. Completion of installation

The completion window is displayed. Click [Finish].

#### 3.2.3 Installation in Silent Mode

Installation in silent mode can be performed only when the installation method is one of the following:

- New installation
- Multi-version installation

The installation procedure is described below.



#### See

Refer to the FUJITSU Enterprise Postgres product website for information on installation in silent mode, such as the installation parameters and error messages.

#### 1. Insert the DVD

Insert the Server Assistant program DVD into the DVD drive.

The [Install Menu] window is displayed. Click [Finish].

#### 2. Create an installation parameters CSV file

Create an installation parameters CSV file that uses the following specification format.

```
sectionName, parameterName, value
sectionName, parameterName, value
:
```



The template for the installation parameters CSV file is "Z:\sample\sample.csv (Z is the drive into which the DVD is inserted.).

#### 3. Start the command prompt

In Windows, right-click [Command Prompt] and then select [Run as administrator].

#### 4. Run the installation

Execute the following command:

Z:\>silent.bat c:\temp\inspara.csv

Z: The drive into which the DVD is inserted.

c:\temp\inspara.csv: The installation parameter CSV file name.

If the silent installer ends in an error, a message is output to the log file and return values are returned.

#### 3.3 Uninstallation

This section describes how to uninstall the Windows Server Assistant.



- Before uninstalling the product, close the product program.
- Log in using an account that has administrator privileges and then execute the command, or switch to an account that has administrator privileges and then uninstall the product.



- If a [User Account Control] dialog box such as the following is displayed at the start of the uninstallation, click [Yes] to continue processing:

If [No] is clicked, permission to continue is denied and an [Error] dialog box will be displayed. To continue the uninstallation, click [Retry] in the [Error] dialog box. To end the operation, click [Cancel].

- If uninstallation is suspended or processing terminates abnormally, a dialog box of the Program Compatibility Assistant similar to the one shown below may be displayed. Click [This program uninstalled correctly] and continue operation.

#### 3.3.1 Uninstallation in Interactive Mode

The uninstallation procedure is described below.



If an error occurs while the product is being uninstalled, refer to "Uninstall (middleware) Messages" in the FUJITSU Enterprise Postgres product website, and take the required action.

#### 1. Stop the program

Before starting the uninstallation, stop the following:

Mirroring Controller arbitration process

1. Execute the mc\_arb command in stop mode to stop the Mirroring Controller arbitration process.

#### Example

```
> mc_arb stop -M D:\mcarb_dir\arbiter1
```

2. Unregister the Mirroring Controller arbitration process from the Windows service.

Execute the mc\_arb command in unregister mode to unregister the Mirroring Controller arbitration process from the Windows service.

#### Example

```
> mc_arb unregister -M D:\mcarb_dir\arbiter1
```

3. Delete registrations related to the event log.

If messages are output to the event log, DLLs are registered in accordance with "Preparatory Tasks for the Output of Error Logs to the Event Log (Windows)" in the Cluster Operation Guide (Database Multiplexing). Delete these registrations so that no unnecessary issues occur.

- Delete registration for the default event source name

```
Example)
```

The example below deletes a DLL of a 64-bit product registered under the default event source name. "< x >" indicates the product version.

```
> regsvr32 /u "c:\Program Files\Fujitsu\fsepv<x>assistant64\lib\mcarbevent.dll"
```

- Delete registration for other event source names

DLL registration is performed so that messages output to the event log are output to any event source named by the user.

Accordingly, it is necessary to delete the DLL registration. Do this for any event source name.

```
Example)
```

The example below deletes a DLL of a 64-bit product registered under the event source name "Mirroring Controller arbiterl". "<x>" indicates the product version.

```
> regsvr32 /u /i:"Mirroring Controller arbiter1" "c:\Program Files\Fujitsu
\fsepv<x>assistant64\lib\mcarbevent.dll"
```

- If installing multiple versions

If error logs are set to be output to the event log, use the DLL path name that you took note of previously as explained in "Preparatory Tasks for the Output of Error Logs to the Event Log (Windows)" in the Cluster Operation Guide (Database Multiplexing) to reregister the default event source name.



Ensure that you delete the DLLs before the uninstallation. If you perform the uninstallation without doing so, you may not be able to delete the DLLs at a later time.

#### 2. Start Uninstall (middleware)

In Windows, click [All Programs] or [All apps], then [Fujitsu], and then [Uninstall (middleware)].

#### 3. Select the software

Select the product to be uninstalled from [Software Name], and then click [Remove].

#### 4. Start the uninstallation

Click [Uninstall].

#### 5. Finish the uninstallation

The uninstallation completion window will be displayed. Click [Finish].

The installation folder may remain after uninstallation. If it is not required, delete it.

#### 6. Stop Uninstall (middleware)

Click [Close].

#### 3.3.2 Uninstallation in Silent Mode

The uninstallation procedure is described below.



See

Refer to the FUJITSU Enterprise Postgres product website for information on uninstallation in silent mode, such as the error messages.

#### 1. Stop the program

Before starting the uninstallation, stop the following:

Mirroring Controller arbitration process

1. Execute the mc\_arb command in stop mode to stop the Mirroring Controller arbitration process.

#### Example

```
> mc_arb stop -M D:\mcarb_dir\arbiter1
```

2. Unregister the Mirroring Controller arbitration process from the Windows service.

Execute the mc\_arb command in unregister mode to unregister the Mirroring Controller arbitration process from the Windows service.

#### Example

```
> mc_arb unregister -M D:\mcarb_dir\arbiter1
```

3. Delete registrations related to the event log.

If messages are output to the event log, DLLs are registered in accordance with "Preparatory Tasks for the Output of Error Logs to the Event Log (Windows)" in the Cluster Operation Guide (Database Multiplexing). Delete these registrations so that no unnecessary issues occur.

- Delete registration for the default event source name

```
Example)
```

The example below deletes a DLL of a 64-bit product registered under the default event source name. "<x>" indicates the product version.

```
> regsvr32 /u "c:\Program Files\Fujitsu\fsepv<x>assistant64\lib\mcarbevent.dll"
```

- Delete registration for other event source names

DLL registration is performed so that messages output to the event log are output to any event source named by

Accordingly, it is necessary to delete the DLL registration. Do this for any event source name.

```
Example
```

The example below deletes a DLL of a 64-bit product registered under the event source name "Mirroring Controller arbiterl". "<x>" indicates the product version.

> regsvr32 /u /i:"Mirroring Controller arbiter1" "c:\Program Files\Fujitsu \fsepv<x>assistant64\lib\mcarbevent.dll"

#### - If installing multiple versions

If error logs are set to be output to the event log, use the DLL path name that you took note of previously as explained in "Preparatory Tasks for the Output of Error Logs to the Event Log (Windows)" in the Cluster Operation Guide (Database Multiplexing) to reregister the default event source name.



Ensure that you delete the DLLs before the uninstallation. If you perform the uninstallation without doing so, you may not be able to delete the DLLs at a later time.

#### 2. Start the command prompt

In Windows, right-click [Command Prompt] and then select [Run as administrator].

#### 3. Run the uninstallation

Execute the following command:

The installation folder may remain after uninstallation. If it is not required, delete it.

#### Example

X:> installFolder\suninst.bat

X: Drive on which the product is installed

## Chapter 4 Setup of the Server Assistant

The Server Assistant is a feature that is installed and used on the arbitration server, so its setup is performed as the arbitration server setup.



Refer to "Setting Up Database Multiplexing Mode" in the Cluster Operation Guide (Database Multiplexing) for information on setting up and operating the Mirroring Controller arbitration server.

## Appendix A Estimating Memory Requirements

This appendix explains how to estimate the memory.

## A.1 Server Assistant Memory Requirements

This section describes the formula for estimating memory requirements for the Server Assistant.

Use the following formula to obtain a rough estimate of memory requirements:

Memory usage of the Server Assistant

= Peak memory usage of the Mirroring Controller arbitration

processes

+ Peak memory usage of the Mirroring Controller commands

Peak memory usage of the Mirroring Controller arbitration processes=100 MB

Peak memory usage of the Mirroring Controller commands=50 MB \* Number of commands executed simultaneously

L

## Appendix B Procedure when Modifying the Oracle JRE Installation

This appendix describes the procedure to follow when modifying the Oracle JRE installation.

The Oracle JRE, of which the installation destination is specified in the JAVA\_HOME environment variable when installing FUJITSU Enterprise Postgres, is used by the Server Assistant.

Therefore, when updating or reinstalling Oracle JRE after installing FUJITSU Enterprise Postgres, it is necessary to restart the Mirroring Controller arbitration process, therefore follow the procedure below to modify the Oracle JRE installation:

1. Stop the Mirroring Controller arbitration process.

Refer to the Cluster Operation Guide (Database Multiplexing) for details.

- 2. Modify the Oracle JRE installation.
- 3. Change the installation environment to be used by Mirroring Controller.

Set the JAVA\_HOME environment variable to the installation destination of Oracle JRE 8, and use the mc\_update\_jre\_env command to change the installation environment to be used by the Server Assistant. This procedure must be executed by the superuser.

#### Example

/opt/fsepv<x>assistant/bin is the installation directory where the Server Assistant is installed.

```
$ su -
Password:*****
# export JAVA_HOME="oracleJre8InstallDir"
# /opt/fsepv<x>assistant/bin/mc_update_jre_env
```

4. Start the Mirroring Controller arbitration process.

Refer to the Cluster Operation Guide (Database Multiplexing) for details.

## <u>Index</u>

[C]
Check the disk space5,15
Check the installed product and determine the installation method
[D]
Disk Space Required for Installation4,13
[E]
Estimating Memory Requirements
Excluded Software4,13
[F]
Features that can be Installed
[H]
Hardware Environment
hosts file settings
ſij
Installation
Installation and Uninstallation of the Linux Server Assistant3
Installation and Uninstallation of the Windows Server Assistant
Installation in Interactive Mode1,6,16
Installation in Silent Mode
Installation Procedure
Installation Types
[M]
Multi-Version Installation
[N]
New Installation1
[O]
Operating Environment
Operating environment
[P]
Pre-installation Tasks
Procedure when Wodifying the Oracle JRE instanation
[R]
Reinstallation1
Related Software4,12
Remove applied updates6,16
Required Operating System3,12
Required Patches
[S]
Server Assistant Memory Requirements23
Setup of the Server Assistant
Supported System Environment
· · · · · · · · · · · · · · · · · ·
[T]
TCP/IP Protocol
TCP/IP protocol

[U]	
Uninstallation	2,9,1
Uninstallation in Interactive Mode	9,1
Uninstallation in Silent Mode	11,20



# FUJITSU Enterprise Postgres 12



# DevSecOps

Application Development Guide > Operation Guide > Security Operation Guide > Cluster Operation Guide > Connection Manager User Guide >



## FUJITSU Enterprise Postgres 12



# Application Development Guide

## Preface

#### Purpose of this document

This is a guide for the developers of FUJITSU Enterprise Postgres applications.

#### Intended readers

This document is intended for developers of applications that use FUJITSU Enterprise Postgres. Of the interfaces provided by FUJITSU Enterprise Postgres, this guide describes the PostgreSQL extended interface.

Readers of this document are also assumed to have general knowledge of:

- L
- PostgreSQL
- SQL
- Linux
- w
- PostgreSQL
- SQL
- Windows

#### Structure of this document

This document is structured as follows:

#### Chapter 1 Overview of the Application Development Function

Provides an overview of FUJITSU Enterprise Postgres application development.

#### Chapter 2 JDBC Driver

Explains how to use JDBC drivers.

#### Chapter 3 ODBC Driver

Explains how to use ODBC drivers.

#### Chapter 4 .NET Data Provider

Explains how to use .NET Data Provider.

#### Chapter 5 C Library (libpq)

Explains how to use C applications.

#### Chapter 6 Embedded SQL in C

Explains how to use embedded SQL in C.

#### Chapter 7 Embedded SQL in COBOL

Explains how to use embedded SQL in COBOL.

#### Chapter 8 SQL References

Explains the SQL statements which were extended in FUJITSU Enterprise Postgres development.

#### Chapter 9 Compatibility with Oracle Databases

Explains features that are compatible with Oracle databases.

#### Chapter 10 Application Connection Switch Feature

Explains the application connection switch feature.

#### Chapter 11 Performance Tuning

Explains how to tune application performance.

#### Chapter 12 Scan Using a Vertical Clustered Index (VCI)

Explains how to perform scan using a Vertical Clustered Index (VCI).

#### Appendix A Precautions when Developing Applications

Provides some points to note about application development.

#### Appendix B Conversion Procedures Required due to Differences from Oracle Database

Explains how to convert from an Oracle database to FUJITSU Enterprise Postgres, within the scope noted in "Compatibility with Oracle Databases" from the following perspectives.

#### Appendix C Tables Used by the Features Compatible with Oracle Databases

Explains the tables used by the features compatible with Oracle databases.

#### Appendix D ECOBPG - Embedded SQL in COBOL

Explains application development using embedded SQL in COBOL.

#### Appendix E Quantitative Limits

This appendix explains limitations.

#### Appendix F Reference

Provides a reference for each interface.

#### **Export restrictions**

Exportation/release of this document may require necessary procedures in accordance with the regulations of your resident country and/or US export control laws.

#### Issue date and version

Edition 1.0: July 2020

#### Copyright

Copyright 2015-2020 FUJITSU LIMITED

## Contents

Chapter 1 Overview of the Application Development Function	1
1.1 Support for National Characters.	
1.1 Literal	2
1.1.2 Data Type	3
1.1.3 Functions and Operator	3
1.2 Integration with Visual Studio.	3
1.2.1 Relationship between .NET Framework and FUJITSU Enterprise Postgres	4
1.2.2 Automatic Application Generation	
1.3 Compatibility with Oracle Database	
1.4 Application Connection Switch Feature	
1.4.1 Integration with Database Multiplexing	6
1.5 Notes on Application Compatibility	7
1.5.1 Checking Execution Results	7
1.5.2 Referencing System Catalogs	7
1.5.3 Using Functions	
Chapter 2 JDBC Driver.	
2.1 Development Environment.	
2.1.1 Combining with JDK or JRE	
2.2 Setup	
2.2.1 Environment Settings.	
2.2.2 Message Language and Encoding System Used by Applications Settings.	
2.2.3 Settings for Encrypting Communication Data.	
2.3 Connecting to the Database.	
2.3.1 Using the DriverManager Class.	
2.3.2 Using the PGConnectionPoolDataSource Class	
2.3.3 Using the PGXADataSource Class.	
2.4 Application Development	
2.4.1 Relationship between the Application Data Types and Database Data Types	
2.4.2 Statement Caching Feature.	
2.4.3 Creating Applications while in Database Multiplexing Mode	
2.4.3.1 Errors when an Application Connection Switch Occurs and Corresponding Actions	10
Chapter 3 ODBC Driver	17
3.1 Development Environment	17
3.2 Setup	17
3.2.1 Registering ODBC Drivers	
3.2.2 Registering ODBC Data Sources(for Windows(R))	19
3.2.2.1 Registering Using GUI.	19
3.2.2.2 Registering Using Commands	21
3.2.3 Registering ODBC Data Sources(for Linux)	24
3.2.4 Message Language and Encoding System Used by Applications Settings	26
3.3 Connecting to the Database	27
3.4 Application Development	27
3.4.1 Compiling Applications (for Windows (R))	27
3.4.2 Compiling Applications (for Linux)	28
3.4.3 Creating Applications While in Database Multiplexing Mode	28
3.4.3.1 Errors when an Application Connection Switch Occurs and Corresponding Actions	29
Chapter 4 .NET Data Provider	30
4.1 Development Environment	
4.2 Setup	
4.2.1 Setting Up the Visual Studio Integration Add-On.	
4.2.2 Setting Up .NET Data Provider	
4.2.3 Setting Up .NET Data Provider Type Plugins	
4.2.4 Setting Up Noggal for Entity Framework	32

4.2.5 Message Language Settings	32
4.3 Connecting to the Database	
4.3.1 Using NpgsqlConnection	33
4.3.2 Using NpgsqlConnectionStringBuilder	33
4.3.3 Using the ProviderFactory Class	33
4.3.4 Connection String	34
4.4 Application Development	36
4.4.1 Data Types	36
4.4.2 Relationship between Application Data Types and Database Data Types	
4.4.3 Creating Applications while in Database Multiplexing Mode	
4.4.3.1 Errors when an Application Connection Switch Occurs and Corresponding Actions	
4.4.4 Notes	
4.5 Uninstallation	
4.5.1 Uninstalling Npgsql	
4.5.2 Uninstalling .NET Data Provider Type Plugins	
4.5.3 Uninstalling Npgsql for Entity Framework	46
Chapter 5 C Library (libpq)	47
5.1 Development Environment	47
5.2 Setup	47
5.2.1 Environment Settings.	47
5.2.2 Message Language and Encoding System Used by Applications Settings	48
5.2.3 Settings for Encrypting Communication Data	49
5.3 Connecting with the Database	49
5.4 Application Development	50
5.4.1 Compiling Applications	50
5.4.2 Creating Applications while in Database Multiplexing Mode	50
5.4.2.1 Errors when an Application Connection Switch Occurs and Corresponding Actions	51
Chapter 6 Embedded SQL in C	52
6.1 Development Environment	52
6.2 Setup	52
6.2.1 Environment Settings	52
6.2.2 Message Language and Encoding System Used by Applications Settings	52
6.2.3 Settings for Encrypting Communication Data	52
6.3 Connecting with the Database	53
6.4 Application Development	55
6.4.1 Support for National Character Data Types	55
6.4.2 Compiling Applications	55
6.4.3 Bulk INSERT	57
6.4.4 DECLARE STATEMENT	61
6.4.5 Creating Applications while in Database Multiplexing Mode	63
6.4.5.1 Errors when an Application Connection Switch Occurs and Corresponding Actions	
6.4.6 Notes	64
Chapter 7 Embedded SQL in COBOL	65
7.1 Development Environment	
7.2 Setup	
7.2.1 Environment Settings	65
7.2.2 Message Language and Encoding System Used by Applications	
7.2.3 Settings for Encrypting Communication Data	
7.3 Connecting with the Database	
7.4 Application Development	
7.4.1 Support for National Character Data Types	
7.4.2 Compiling Applications.	
7.4.3 Bulk INSERT	
7.4.4 DECLARE STATEMENT	
7.4.5 Creating Applications while in Database Multiplexing Mode	

Chapter 8 SQL References	77
8.1 Expanded Trigger Definition Feature	
8.1.1 CREATE TRIGGER.	
8.1.2 How to Define Triggers in pgAdmin	
Chapter 9 Compatibility with Oracle Databases	79
9.1 Overview	
9.2 Precautions when Using the Features Compatible with Oracle Databases	
9.2.1 Notes on SUBSTR	
9.2.2 Notes when Integrating with the Interface for Application Development	
9.3 Queries.	80
9.3.1 Outer Join Operator (+)	80
9.3.2 DUAL Table	82
9.4 SQL Function Reference	83
9.4.1 DECODE	83
9.4.2 SUBSTR	85
9.4.3 NVL	86
9.5 Package Reference	87
9.5.1 DBMS_OUTPUT	87
9.5.1.1 Description.	88
9.5.1.2 Example	90
9.5.2 UTL_FILE	91
9.5.2.1 Registering and Deleting Directories.	
9.5.2.2 Description	92
9.5.2.3 Example	96
9.5.3 DBMS_SQL	97
9.5.3.1 Description	
9.5.3.2 Example	
Chapter 10 Application Connection Switch Feature	105
10.1 Connection Information for the Application Connection Switch Feature	
10.2 Using the Application Connection Switch Feature	
10.2.1 Using the JDBC Driver	
10.2.2 Using the ODBC Driver.	
10.2.3 Using a .NET Data Provider	
10.2.4 Using a Connection Service File	
10.2.5 Using the C Library (libpq)	
10.2.6 Using Embedded SQL	
10.2.7 Using the psql Command	
Chapter 11 Performance Tuning	117
11.1 Enhanced Query Plan Stability	
11.1.1 Optimizer Hints	
11.1.2 Locked Statistics	
Chapter 12 Scan Using a Vertical Clustered Index (VCI)	123
12.1 Operating Conditions	
12.2 Usage	
12.2.1 Designing	
12.2.2 Checking	
12.2.3 Evaluating	
12.3 Usage Notes.	
Appendix A Precautions when Developing Applications	129
A.1 Precautions when Using Functions and Operators	
A.1.1 General rules of Functions and Operators	
11.1.1 Ocherar raics of ranctions and Operators	

A.2 Notes when Using Temporary Tables	130
A.3 Implicit Data Type Conversions.	130
A.3.1 Function Argument	132
A.3.2 Operators.	132
A.3.3 Storing Values	133
A.4 Notes on Using Index	133
A.4.1 SP-GiST Index	
A.5 Notes on Using Multibyte Characters in Definition Names	133
Appendix B Conversion Procedures Required due to Differences from Oracle Database	
B.1 Outer Join Operator (Perform Outer Join)	
B.1.1 Comparing with the ^= Comparison Operator	
B.2 DECODE (Compare Values and Return Corresponding Results)	
B.2.1 Comparing Numeric Data of Character String Types and Numeric Characters	
B.2.2 Obtaining Comparison Result from more than 50 Conditional Expressions	
B.2.3 Obtaining Comparison Result from Values with Different Data Types	
B.3 SUBSTR (Extract a String of the Specified Length from Another String)	
B.3.1 Specifying a Value Expression with a Data Type Different from the One that can be Specified for Function Argument	
B.3.2 Extracting a String with the Specified Format from a Datetime Type Value	
B.3.3 Concatenating a String Value with a NULL value	
B.4 NVL (Replace NULL)	
B.4.1 Obtaining Result from Arguments with Different Data Types	
B.4.2 Operating on Datetime/Numeric, Including Adding Number of Days to a Particular Day	
B.4.3 Calculating INTERVAL Values, Including Adding Periods to a Date	
B.5 DBMS_OUTPUT (Output Messages)	
B.5.1 Outputting Messages Such As Process Progress Status	
B.5.2 Receiving a Return Value from a Procedure (PL/SQL) Block (For GET_LINES)	
B.5.3 Receiving a Return Value from a Procedure (PL/SQL) Block (For GET_LINE)	
B.6 UTL_FILE (Perform File Operation)	
B.6.1 Registering a Directory to Load and Write Text Files	
B.6.2 Checking File Information	
B.6.3 Copying Files	
B.6.4 Moving/Renaming Files	
B.7 DBMS_SQL (Execute Dynamic SQL)	
B.7.1 Searching Using a Cursor	154
Appendix C Tables Used by the Features Compatible with Oracle Databases	
C.1 UTL_FILE.UTL_FILE_DIR	160
Appendix D ECOBPG - Embedded SQL in COBOL	
D.1 Precautions when Using Functions and Operators.	161
D.2 Managing Database Connections	162
D.2.1 Connecting to the Database Server	162
D.2.2 Choosing a Connection	
D.2.3 Closing a Connection	164
D.3 Running SQL Commands	164
D.3.1 Executing SQL Statements	164
D.3.2 Using Cursors	165
D.3.3 Managing Transactions.	165
D.3.4 Prepared Statements.	165
D.4 Using Host Variables	
D.4.1 Overview	166
D.4.2 Declare Sections	166
D.4.3 Retrieving Query Results	
D.4.4 Type Mapping	
D.4.5 Handling Nonprimitive SQL Data Types	
D.4.6 Indicators	
D.5 Dynamic SOI	176

D.5.1 Executing Statements without a Result Set	176
D.5.2 Executing a Statement with Input Parameters	177
D.5.3 Executing a Statement with a Result Set	177
D.6 Using Descriptor Areas	178
D.6.1 Named SQL Descriptor Areas	
D.7 Error Handling	180
D.7.1 Setting Callbacks	
D.7.2 sqlca	
D.7.3 SQLSTATE vs. SQLCODE	182
D.8 Preprocessor Directives	
D.8.1 Including Files.	
D.8.2 The define and undef Directives.	186
D.8.3 ifdef, ifndef, else, elif, and endif Directives	
D.9 Processing Embedded SQL Programs	
D.10 Large Objects	187
D.11 Embedded SQL Commands	187
D.11.1 ALLOCATE DESCRIPTOR	188
D.11.2 CONNECT	189
D.11.3 DEALLOCATE DESCRIPTOR	191
D.11.4 DECLARE	191
D.11.5 DESCRIBE	192
D.11.6 DISCONNECT	193
D.11.7 EXECUTE IMMEDIATE	
D.11.8 GET DESCRIPTOR	
D.11.9 OPEN	196
D.11.10 PREPARE	197
D.11.11 SET AUTOCOMMIT	198
D.11.12 SET CONNECTION	
D.11.13 SET DESCRIPTOR	199
D.11.14 TYPE	200
D.11.15 VAR	201
D.11.16 WHENEVER	202
D.12 PostgreSQL Client Applications.	203
D.12.1 ecobpg	203
Appendix E Quantitative Limits	205
ppendix F Reference	210
F.1 JDBC Driver	
F.2 ODBC Driver	
F.2.1 List of Supported APIs	
F.3 .NET Data Provider	
F.4 C Library (libpq)	
F.5 Embedded SQL in C	213
aday	21.4

## Chapter 1 Overview of the Application Development Function

The interface for application development provided by FUJITSU Enterprise Postgres is perfectly compatible with PostgreSQL.

Along with the PostgreSQL interface, FUJITSU Enterprise Postgres also provides the following extended interfaces:

#### - Support for National Characters

In order to secure portability from mainframes and databases of other companies, FUJITSU Enterprise Postgres provides data types that support national characters. The national characters are usable from the client application languages.

Refer to "1.1 Support for National Characters" for details.

#### - Integration with Visual Studio

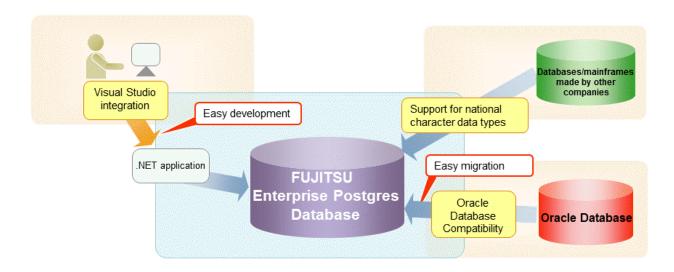
By integrating with Visual Studio, applications can be created using a standard framework for the building of a database server

Refer to "1.2 Integration with Visual Studio" for details.

#### - Compatibility with Oracle Databases

Compatibility with Oracle databases is offered. Use of the compatible features means that the revisions to existing applications can be isolated, and migration to open interfaces is made simpler.

Refer to "1.3 Compatibility with Oracle Database" for details.



#### - Application connection switch feature

The application connection switch feature is provided to enable automatic connection to the target server when there are multiple servers with redundant configurations.

Refer to "1.4 Application Connection Switch Feature" for details.

- Performance tuning

The following features are provided to control SQL statement query plans:

- Optimizer hints
- Locked statistics

Refer to "11.1 Enhanced Query Plan Stability" for details.

- Scanning using a Vertical Clustered Index (VCI)

Scans becomes faster during aggregation of many rows by providing the features below:

- Vertical clustered index (VCI)
- In-memory data

This feature can only be used in Advanced Edition.

Refer to "Chapter 12 Scan Using a Vertical Clustered Index (VCI)" for details.

## 1.1 Support for National Characters

NCHAR type is provided as the data type to deal with national characters.

The NCHAR type can be used with FUJITSU Enterprise Postgres pgAdmin.



- NCHAR can only be used when the character set of the database is UTF-8.
- NCHAR can be used in the places where CHAR can be used (function arguments, etc.).
- For applications handling NCHAR type data in the database, the data format is the same as CHAR type. Therefore, applications handling data in NCHAR type columns can also be used to handle data stored in CHAR type columns.



Note the following in order to cast NCHAR type data as CHAR type.

- When comparing NCHAR type data where the length differs, ASCII spaces are used to fill in the length of the shorter NCHAR type data so that it can be processed as CHAR type data.
- Depending on the character set, the data size may increase by between 1.5 and 2 times.

#### 1.1.1 Literal

#### Syntax

```
{ N | n }'[national character [ ...]]'
```

#### General rules

National character string literals consist of an 'N' or 'n', and the national character is enclosed in single quotation marks ('). Example: N'ABCDEF'

The data type is national character string type.

### 1.1.2 Data Type

#### Syntax

```
{ NATIONAL CHARACTER | NATIONAL CHAR | NCHAR } [ VARYING ][(length) ]
```

The data type of the NCHAR type column is as follows:

Data type specification format	Explanation
NATIONAL CHARACTER(n)	National character string with a fixed length of <i>n</i> characters
NATIONAL CHAR(n)	This will be the same as (1) if ( <i>n</i> ) is omitted.
NCHAR(n)	n is a whole number larger than 0.
NATIONAL CHARACTER VARYING(n)	National character string with a variable length with a maximum of <i>n</i> characters
NATIONAL CHAR VARYING(n) NCHAR VARYING(n)	Any length of national character string can be accepted when this is omitted.
	n is a whole number larger than 0.

#### General rules

NCHAR is the national character string type data type. The length is the number of characters.

The length of the national character string type is as follows:

- When VARYING is not specified, the length of national character strings is fixed and will be the specified length.
- When VARYING is specified, the length of national character strings will be variable. In this case, the lower limit will be 0 and the upper limit will be the value specified for length.
- NATIONAL CHARACTER, NATIONAL CHAR, and NCHAR each have the same meaning.

When the national character string to be stored is shorter than the declared upper limit, the NCHAR value is filled with spaces, whereas NCHAR VARYING is stored as is.

The upper limit for character storage is approximately 1GB.

### 1.1.3 Functions and Operator

#### Comparison operator

When a NCHAR type or NCHAR VARYING type is specified in a comparison operator, comparison is only possible between NCHAR types or NCHAR VARYING types.

#### String functions and operators

All of the string functions and operators that can be specified by a CHAR type can also be specified by a NCHAR type. The behavior of these string functions and operators is also the same as with CHAR type.

Pattern matching (LIKE, SIMILAR TO regular expression, POSIX regular expression)

The patterns specified when pattern matching with NCHAR types and NCHAR VARYING types specify the percent sign (%) and the underline (\_).

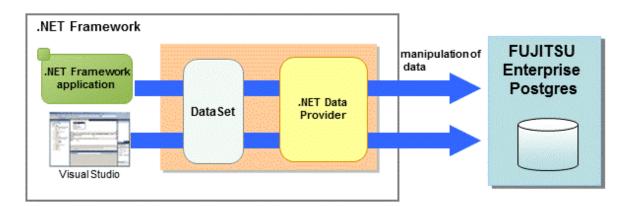
The underline (\_) means a match with one national character. The percent sign (%) means a match with any number of national characters 0 or over.

### 1.2 Integration with Visual Studio

When developing an application to access database server resources, you can create applications and build database server environments integrated with Microsoft Visual Studio.

# 1.2.1 Relationship between .NET Framework and FUJITSU Enterprise Postgres

FUJITSU Enterprise Postgres provides .NET Data Provider, which is an interface for ADO.NET of .NET Framework. This enables you to select FUJITSU Enterprise Postgres as the connection destination database of ADO.NET and use the intuitive and efficient application development features of Visual Studio.



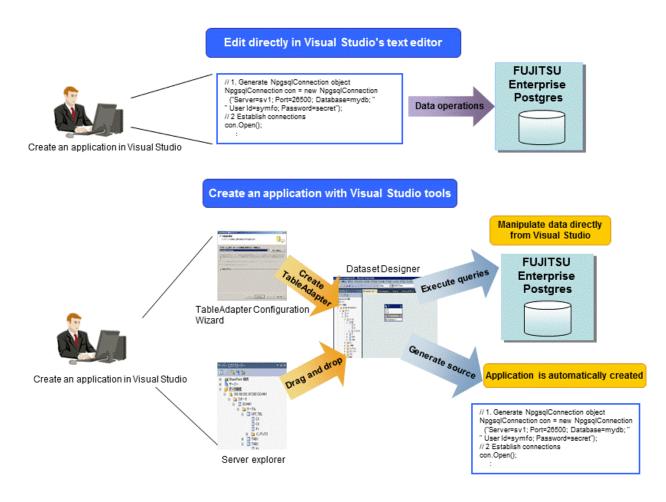
The following provides an overview of application development integrated with Visual Studio.

Edit directory in Visual Studio's text editor

By using a component specified in Visual Studio, applications to access database resources can be created manually.

Create an application with Visual Studio tools

By using basic drag-and-drop operations in the tools provided in Visual Studio, programs to access database resources can be generated automatically.

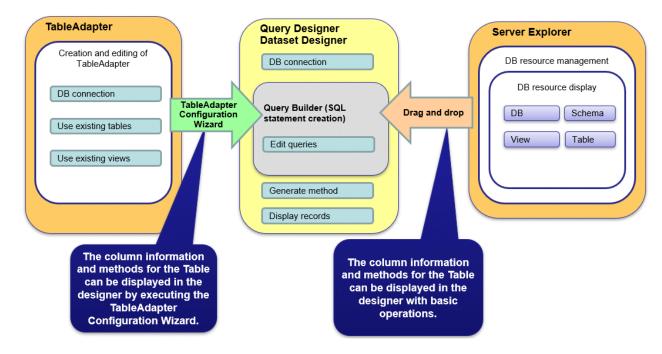


### 1.2.2 Automatic Application Generation

The Visual Studio tools used to automatically generate applications include TableAdapter and Server Explorer, which enable the following:

- Data manipulation of database resources with TableAdapter
- Management of database resources with Server Explorer

Whether you use TableAdapter or the Server Explorer, programs can be created with basic operations like drag and drop with the resources and tools that comprise Visual Studio.



The following features are available with TableAdapter and Server Explorer:

- Manipulation of database resources with TableAdapter
  - Generating queries using existing tables/views
  - Generating methods using existing tables/views
- Management of database resources with Server Explorer
  - Listing of database resources
  - Generating queries using existing tables/views
  - Generating methods using existing tables/views

### 1.3 Compatibility with Oracle Database

The following features have been extended in order to enhance compatibility with Oracle databases:

- Query (external join operator (+), DUAL table)
- Function (DECODE, SUBSTR, NVL)
- Built-in package (DBMS\_OUTPUT, UTL\_FILE, DBMS\_SQL)

Refer to "Chapter 9 Compatibility with Oracle Databases" for information on the features compatible with Oracle databases.

### 1.4 Application Connection Switch Feature

The application connection switch feature enables automatic connection to the target server when there are multiple servers with redundant configurations.

Refer to "Chapter 10 Application Connection Switch Feature" for information on the application connection switch feature.

### 1.4.1 Integration with Database Multiplexing

The application connection switch feature is provided to enable automatic connection to the appropriate server when there are multiple servers with redundant configurations.



Refer to the Cluster Operation Guide (Database Multiplexing) for information on database multiplexing.

### 1.5 Notes on Application Compatibility

FUJITSU Enterprise Postgres upgrades contain feature improvements and enhancements that may affect the applications.

Accordingly, note the points below when developing applications, to ensure compatibility after upgrade.

- Checking execution results
- Referencing system catalogs
- Using functions

### 1.5.1 Checking Execution Results

Refer to SQLSTATE output in messages to check the SQL statements used in applications and the execution results of commands used during development.



See

Refer to Messages for information on the message content and number.

Refer to "PostgreSQL Error Codes" under "Appendixes" in the PostgreSQL Documentation for information on SQLSTATE.

### 1.5.2 Referencing System Catalogs

System catalogs can be used to obtain information about the FUJITSU Enterprise Postgres system and database objects.

However, system catalogs may change when the FUJITSU Enterprise Postgres version is upgraded. Also, there are many system catalogs that return information that is inherent to FUJITSU Enterprise Postgres.

Accordingly, reference the information schema defined in standard SQL (information\_schema) wherever possible. Note also that queries specifying "\*" in the selection list must be avoided to prevent columns being added.



See

Refer to "The Information Schema" under "Client Interfaces" in the PostgreSQL Documentation for details.

The system catalog must be referenced to obtain information not found in the information schema. Instead of directly referencing the system catalog in the application, define a view for that purpose. Note, however, that when defining the view, the column name must be clearly specified after the view name.

An example of defining and using a view is shown below.



CREATE VIEW my\_tablespace\_view(spcname) AS SELECT spcname FROM pg\_tablespace; SELECT \* FROM my\_tablespace\_view V1, pg\_tables T1 WHERE V1.spcname = T1.tablespace;

If changes are made to a system catalog, the user will be able to take action by simply making changes to the view, without the need to make changes to the application.

The following shows an example of taking action by redefining a view as if no changes were made.

The pg\_tablespace system catalog is redefined in response to the column name being changed from spcname to spacename.



DROP VIEW my\_tablespace\_view;
CREATE VIEW my\_tablespace\_view(spcname) AS SELECT spacename FROM pg\_tablespace;

### 1.5.3 Using Functions

The default functions provided with FUJITSU Enterprise Postgres enable a variety of operations and manipulations to be performed, and information to be obtained, using SQL statements.

However, it is possible that internal FUJITSU Enterprise Postgres functions, such as those relating to statistical information or for obtaining system-related information, may change as FUJITSU Enterprise Postgres versions are upgraded.

Accordingly, when using these functions, define them as new functions and then use the newly-defined functions in the applications.

An example of defining and using a function is shown below.



CREATE FUNCTION my\_func(relid regclass) RETURNS bigint LANGUAGE SQL AS 'SELECT pg\_relation\_size(relid)';
SELECT my\_func(2619);

If changes are made to a function, the user will be able to take action by simply redefining the function, without the need to make changes to the application.

The following shows an example of taking action by redefining a function as if no changes were made.

The pg\_relation\_size function is redefined after arguments are added.



DROP FUNCTION my\_func(regclass);
CREATE FUNCTION my\_func(relid regclass) RETURNS bigint LANGUAGE SQL AS 'SELECT pg\_relation\_size(relid,\$\$main\$\$)';

## Chapter 2 JDBC Driver

This section describes how to use JDBC drivers.

### 2.1 Development Environment

This section describes application development using JDBC drivers and the runtime environment.

### 2.1.1 Combining with JDK or JRE

Refer to Installation and Setup Guide for Client for information on combining with JDK or JRE where JDBC drivers can operate.

### 2.2 Setup

This section describes the environment settings required to use JDBC drivers and how to encrypt communication data.

### 2.2.1 Environment Settings

Configuration of the CLASSPATH environment variable is required as part of the runtime environment for JDBC drivers.

The name of the JDBC driver file is as follows:

- If using JDK 6 or JRE 6

postgresql-jdbc4.jar

- If using JDK 7 or JRE 7

postgresql-jdbc41.jar

- If using JDK 8, JRE 8, JDK 11 or JRE 11

postgresql-jdbc42.jar

The examples below show how to set the CLASSPATH environment variable if JDK 6 or JRE 6 is used.

If JDK 7, JRE 7, JDK 8, JRE 8, JDK 11 or JRE 11 is used, only the name of the JDBC driver file will be different. The method for configuring the CLASSPATH environment variable is the same.

Note that "<x>" indicates the product version.

- Linux (64-bit)
  - Setting example (TC shell)

setenv CLASSPATH /opt/fsepv<x>client64/jdbc/lib/postgresql-jdbc4.jar:\${CLASSPATH}

- Setting example (bash)

 ${\tt CLASSPATH=/opt/fsepv<} x > {\tt client64/jdbc/lib/postgresql-jdbc4.jar: \$CLASSPATH: export CLASSPATH}; export CLASSPATH$ 

- Windows (32-bit)
  - Setting example

set CLASSPATH=C:\Program Files\Fujitsu\fsepv<x>client32\JDBC\lib\postgresqljdbc4.jar;%CLASSPATH%

- Windows (64-bit)
  - Setting example (when FUJITSU Enterprise Postgres Client 32-bit is installed)

```
set CLASSPATH=C:\Program Files (x86)\Fujitsu\fsepv<x>client32\JDBC\lib\postgresql-
jdbc4.jar;%CLASSPATH%
```

- Setting example (when FUJITSU Enterprise Postgres Client 64-bit is installed)

```
set CLASSPATH=C:\Program Files\Fujitsu\fsepv<x>client64\JDBC\lib\postgresql-
jdbc4.jar;%CLASSPATH%
```

# 2.2.2 Message Language and Encoding System Used by Applications Settings

If the JDBC driver is used, it will automatically set the encoding system on the client to UTF-8, so there is no need to configure this.



#### See

Refer to "Automatic Character Set Conversion Between Server and Client" in "Server Administration" in the PostgreSQL Documentation for information on encoding systems.

#### Language settings

You must match the language settings for the application runtime environment with the message locale settings of the database server.

Set language in the "user.language" system property.



Example of running a Java command with system property specified

```
java -Duser.language=en TestClass1
```

### 2.2.3 Settings for Encrypting Communication Data

When using the communication data encryption feature to connect to the database server, set as follows:

Settings for encrypting communication data for connection to the server

This section describes how to create applications for encrypting communication data.

Set the property of the SSL parameter to "true" to encrypt. The default for the SSL parameter is "false".

If ssl is set to "true", sslmode is internally treated as "verify-full".



Setting example 1

```
String url = "jdbc:postgresql://svl/test";
Properties props = new Properties();
props.setProperty("user", "fsepuser");
props.setProperty("password", "secret");
props.setProperty("ssl", "true");
```

.....

```
props.setProperty("sslfactory","org.postgresql.ssl.DefaultJavaSSLFactory");
Connection conn = DriverManager.getConnection(url, props);
```

- Setting example 2

```
String url = "jdbc:postgresql://svl/test?
user=fsepuser&password=secret&ssl=true&sslfactory=org.postgresql.ssl.DefaultJavaSSLF
actory";
Connection conn = DriverManager.getConnection(url);
```

To prevent spoofing of the database server, you need to use the keytool command included with Java to import the CA certificate to the Java keystore. In addition, specify "org.postgresql.ssl.DefaultJavaSSLFactory" for the sslfactory parameter. Refer to JDK documentation and the Oracle website for details.



There is no need to set the ssl parameter if the connection string of the DriverManager class is specified, or if the sslmode parameter is specified in the data source, such as when the application connection switch feature is used. If the ssl parameter is set, the value in the sslmode parameter will be enabled.



Refer to "Secure TCP/IP Connections with SSL" in "Server Administration" in the PostgreSQL Documentation for information on encrypting communication data.

### 2.3 Connecting to the Database

This section explains how to connect to a database.

- Using the DriverManager Class
- Using the PGConnectionPoolDataSource Class
- Using the PGXADataSource Class



Do not specify "V2" for the "protocol Version" of the connection string.

### 2.3.1 Using the DriverManager Class

To connect to the database using the DriverManager class, first load the JDBC driver, then specify the connection string as a URI in the API of the DriverManager class.

Load the JDBC driver

Specify org.postgresql.Driver.

Connection string

URI connection is performed as follows:

jdbc:postgresql://host:port/database?
user=user&password=password1&loginTimeout=loginTimeout&socketTimeout=socketTimeout

Argument	Description
host	Specify the host name for the connection destination.

Argument	Description	
port	Specify the port number for the database server.	
	The default is "27500".	
database	Specify the database name.	
user	Specify the username that will connect with the database.	
	If this is omitted, the username logged into the operating system that is executing the application will be used.	
password	Specify a password when authentication is required.	
loginTimeout	Specify the timeout for connections (in units of seconds).	
	Specify a value between 0 and 2147483647. There is no limit set if you set 0 or an invalid value.	
	An error occurs when a connection cannot be established within the specified time.	
socketTimeout	Specify the timeout for communication with the server (in units of seconds).	
	Specify a value between 0 and 2147483647. There is no limit set if you set 0 or an invalid value.	
	An error occurs when data is not received from the server within the specified time.	



#### Code examples for applications

```
import java.sql.*;
...
Class.forName("org.postgresql.Driver");
String url = "jdbc:postgresql://sv1:27500/mydb?
user=myuser&password=myuser01&loginTimeout=20&socketTimeout=20";
Connection con = DriverManager.getConnection(url);
```

### 2.3.2 Using the PGConnectionPoolDataSource Class

To connect to databases using data sources, specify the connection information in the properties of the data source.

#### Method description

Argument	Description	
setServerName	Specify the host name for the connection destination.	
setPortNumber	Specify the port number for the database server.	
	The default is "27500".	
setDatabaseName	Specify the database name.	
setUser	Specify the username of the database.	
	By default, the name used will be that of the user on the operating system that is executing the application.	
setPassword	Specify a password for server authentication.	
setLoginTimeout	Specify the timeout for connections (in units of seconds).	
	Specify a value between 0 and 2147483647. There is no limit set if you set 0 or an invalid value.	
	An error occurs when a connection cannot be established within the specified time.	

Argument	Description	
setSocketTimeout	Specify the timeout for communication with the server (in units of seconds).	
	Specify a value between 0 and 2147483647. There is no limit set if you set 0 or an invalid value.	
	An error occurs when data is not received from the server within the specified time.	



#### Code examples for applications

```
import java.sql.*;
import org.postgresql.ds.PGConnectionPoolDataSource;
...
PGConnectionPoolDataSource source = new PGConnectionPoolDataSource();
source.setServerName("svl");
source.setPortNumber(27500);
source.setDatabaseName("mydb");
source.setUser("myuser");
source.setPassword("myuser01");
source.setLoginTimeout(20);
source.setSocketTimeout(20);
...
Connection con = source.getConnection();
```

### 2.3.3 Using the PGXADataSource Class

To connect to databases using data sources, specify the connection information in the properties of the data source.

#### Method description

Argument	Description	
setServerName	Specify the host name for the connection destination.	
setPortNumber	Specify the port number for the database server.	
	The default is "27500".	
setDatabaseName	Specify the database name.	
setUser	Specify the username that will connect with the database.	
	If this is omitted, the name used will be that of the user on the operating system that is executing the application.	
setPassword	Specify a password when authentication by a password is required.	
setLoginTimeout	Specify the timeout for connections.	
	The units are seconds. Specify a value between 0 and 2147483647. There is no limit set if you set 0 or an invalid value.	
	An error occurs when a connection cannot be established within the specified time.	
setSocketTimeout	Specify the timeout for communication with the server.	
	The units are seconds. Specify a value between 0 and 2147483647. There is no limit set if you set 0 or an invalid value.	
	An error occurs when data is not received from the server within the specified time.	



#### Code examples for applications

```
import java.sql.*;
import org.postgresql.xa.PGXADataSource;
...
PGXADataSource source = new PGXADataSource();
source.setServerName("sv1");
source.setPortNumber(27500);
source.setDatabaseName("mydb");
source.setUser("myuser");
source.setPassword("myuser01");
source.setLoginTimeout(20);
source.setSocketTimeout(20);...
Connection con = source.getConnection();
```

### 2.4 Application Development

This section describes the data types required when developing applications that will be connected with FUJITSU Enterprise Postgres.

# 2.4.1 Relationship between the Application Data Types and Database Data Types

The following table shows the correspondence between data types in applications and data types in databases.

Data type on the server	Data type on the server  Java data type  Data types prescrib java.sql.Types	
character	String	java.sql.Types.CHAR
national character	String	java.sql.Types.NCHAR
character varying	String	java.sql.Types.VARCHAR
national character varying	String	java.sql.Types.NVARCHAR
text	String	java.sql.Types.VARCHAR
bytea	byte[]	java.sql.Types.BINARY
smallint	short	java.sql.Types.SMALLINT
integer	int	java.sql.Types.INTEGER
bigint	long	java.sql.Types.BIGINT
smallserial	short	java.sql.Types.SMALLINT
serial	int	java.sql.Types.INTEGER
bigserial	long	java.sql.Types.BIGINT
real	float	java.sql.Types.REAL
double precision	double	java.sql.Types.DOUBLE
numeric	java.math.BigDecimal	java.sql.Types.NUMERIC
decimal	java.math.BigDecimal	java.sql.Types.DECIMAL
money	String	java.sql.Types.OTHER
date	java.sql.Date	java.sql.Types.DATE
time with time zone	java.sql.Time	java.sql.Types.TIME

Data type on the server	Java data type	Data types prescribed by java.sql.Types
time without time zone	java.sql.Time	java.sql.Types.TIME
timestamp without time zone	java.sql.Timestamp	java.sql.Types.TIMESTAMP
timestamp with time zone	java.sql.Timestamp	java.sql.Types.TIMESTAMP
interval	org.postgresql.util.PGInterval	java.sql.Types.OTHER
boolean	boolean	java.sql.Types.BIT
bit	boolean	java.sql.Types.BIT
bit varying	org.postgresql.util.Pgobject	java.sql.Types.OTHER
oid	long	java.sql.Types.BIGINT
xml	java.sql.SQLXML	java.sql.Types.SQLXML
array	java.sql.Array	java.sql.Types.ARRAY
uuid	java.util.UUID	java.sql.Types.OTHER
point	org.postgresql.geometric.Pgpoint	java.sql.Types.OTHER
box	org.postgresql.geometric.Pgbox	java.sql.Types.OTHER
lseg	org.postgresql.geometric.Pglseg	java.sql.Types.OTHER
path	org.postgresql.geometric.Pgpath	java.sql.Types.OTHER
polygon	org.postgresql.geometric.PGpolygon	java.sql.Types.OTHER
circle	org.postgresql.geometric.PGcircle	java.sql.Types.OTHER
json	org.postgresql.util.PGobject	java.sql.Types.OTHER
Network address type (inet,cidr,macaddr, macaddr8)	org.postgresql.util.PGobject	java.sql.Types.OTHER
Types related to text searches (svector, tsquery)	org.postgresql.util.PGobject	java.sql.Types.OTHER
Enumerated type	org.postgresql.util.PGobject	java.sql.Types.OTHER
Composite type	org.postgresql.util.PGobject	java.sql.Types.OTHER
Range type	org.postgresql.util.PGobject	java.sql.Types.OTHER

Although the getString() method of the ResultSet object can be used for all server data types, it is not guaranteed that it will always return a string in the same format for the same data type.

Strings in a format compatible with the JDBC specifications can be obtained using the Java toString() method of the appropriate data type (for example, getInt(), getTimestamp()) to conform to the data type on the server.

### 2.4.2 Statement Caching Feature

The statement caching feature caches SQL statements for each individual connection. This means that when an SQL statement with an identical string is next executed, the analysis and creation of the statement can be skipped. This improves performance in cases such as when an SQL statement with an identical string is executed within a loop or method that is executed repeatedly. Furthermore, the statement caching feature can be combined with the connection pooling feature to further enhance performance.

#### Cache registration controls

You can configure whether to cache SQL statements using the setPoolable(boolean) method of the PreparedStatement class when the statement caching feature is enabled.

Values that can be configured are shown below:

false

SQL statements will not be cached, even when the statement caching feature is enabled.

true

SQL statements will be cached if the statement caching feature is enabled.

### 2.4.3 Creating Applications while in Database Multiplexing Mode

This section explains points to consider when creating applications while in database multiplexing mode.



#### See

- Refer to the Cluster Operation Guide (Database Multiplexing) for information on database multiplexing mode.
- Refer to "Application Development" in the Cluster Operation Guide (PRIMECLUSTER) for points to consider when creating applications using the failover feature integrated with the cluster software.

## 2.4.3.1 Errors when an Application Connection Switch Occurs and Corresponding Actions

If an application connection switch occurs while in database multiplexing mode, explicitly close the connection and then reestablish the connection or reexecute the application.

The table below shows errors that may occur during a switch, and the corresponding action to take.

State		Error information (*1)	Action
Server failure or FUJITSU Enterprise Postgres system failure	Failure occurs during access	57P01 08006 08007	After the switch is complete, reestablish the connection, or reexecute the application.
	Accessed during system failure	08001	
Switch to the standby server	Switched during access	57P01 08006 08007	
	Accessed during switch	08001	

<sup>\*1:</sup> Return value of the getSQLState() method of SQLException.

## Chapter 3 ODBC Driver

This section describes application development using ODBC drivers.

### 3.1 Development Environment

Applications using ODBC drivers can be developed using ODBC interface compatible applications, such as Access, Excel, and Visual Basic.

Refer to the manuals for the programming languages corresponding to the ODBC interface for information about the environment for development.

FUJITSU Enterprise Postgres supports ODBC 3.5.

### 3.2 Setup

You need to set up PsqlODBC, which is an ODBC driver, in order to use applications that use ODBC drivers with FUJITSU Enterprise Postgres. PsqlODBC is included in the FUJITSU Enterprise Postgres client package.

The following describes how to register the ODBC drivers and the ODBC data source.

### 3.2.1 Registering ODBC Drivers

When using the ODBC driver on Linux, register the ODBC driver using the following procedure:

1. Install the ODBC driver manager (unixODBC)



- FUJITSU Enterprise Postgres supports unixODBC Version 2.3 or later.

You can download unixODBC from the following site:

http://www.unixodbc.org/

- To execute unixODBC, you must first install libtool 2.2.6 or later.

You can download libtool from the following website:

http://www.gnu.org/software/libtool/

#### [Note]

- ODBC driver operation is supported.
- unixODBC operation is not supported.
- 2. Register the ODBC drivers

Edit the ODBC driver manager (unixODBC) odbcinst.ini file.



[location of the odbcinst.ini file]

unixOdbcInstallDir/etc/odbcinst.ini

Set the following content:

Definition name	Description	Setting value	
[Driver name]	ODBC driver name	Set the name of the ODBC driver.  Select the two strings below that correspond to the application type. Concatenate the strings with no spaces, enclose in "[]", and then specify this as the driver name.  Note  The placeholders shown below are enclosed in angle brackets '<>' to avoid confusion with literal text. Do not include the angle brackets in the string.  - Application architecture  "FUJITSUEnterprisePostgres - Encoding system used by the application  - In Unicode (only UTF-8 can be used)  "unicode"  - Other than Unicode  "ansi"  Example: The encoding system used by the application is Unicode:  "[FUJITSUEnterprisePostgres - FujitsuEnterprisePostgresClientVers>	
Descriptio n	Descriptio n of the ODBC driver	x64unicode]"  Specify a supplementary description for the current data source. Any description may be set.	
Driver64	Path of the ODBC driver (64- bit)	Set the path of the ODBC driver (64-bit).  - If the encoding system is Unicode:  fujitsuEnterprisePostgresClientInstallDir/ odbc/lib/psqlodbcw.so  - If the encoding system is other than Unicode:  fujitsuEnterprisePostgresClientInstallDir/ odbc/lib/psqlodbca.so	
FileUsage	Use of the data source file	Specify 1.	
Threading	Level of atomicity secured for connectio n pooling	Specify 2.	



Note that "<x>" indicates the product version.

[FUJITSU Enterprise Postgres12x64unicode]
Description = FUJITSU Enterprise Postgres 12 x64 unicode driver
Driver64 = /opt/fsepv<x>client64/odbc/lib/psqlodbcw.so
FileUsage = 1
Threading = 2

•••••

w

### 3.2.2 Registering ODBC Data Sources(for Windows(R))

This section describes how to register ODBC data sources.

There are the following two ways to register ODBC data sources on Windows(R).

### 3.2.2.1 Registering Using GUI

This section describes how to start the [ODBC Data Source Administrator] and register ODBC data sources.

Use the following procedure to register ODBC data sources:

1. Start the [ODBC Data Source Administrator].

Select [Start] >> [Control Panel] >> [Administrative Tools] >> [ODBC Data Source Administrator].

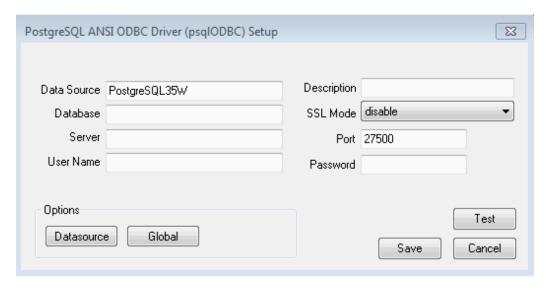


To register data sources for 32-bit applications in Windows(R) for 64-bit, execute the ODBC administrator (odbcad32.exe) for 32-bit, as shown below.

%SYSTEMDRIVE%\WINDOWS\SysWOW64\odbcad32.exe

- 2. When only the current user is to use the ODBC data source, select [User DSN]. When all users using the same computer are to use the ODBC data source, select [System DSN].
- 3. Click [Add].
- 4. Select one of the following drivers from the list of available ODBC drivers displayed in [Create New Data Source], and then click [Finish]. The notation "x" indicates the version of the FUJITSU Enterprise Postgres client feature.
  - FUJITSU Enterprise Postgres Unicode *x*Select this driver if using Unicode as the application encoding system.
  - FUJITSU Enterprise Postgres ANSI *x* Select this driver if using other than Unicode as the application encoding system.

5. The [PostgreSQL ANSI ODBC Driver (psqlODBC) Setup] window is displayed. Enter or select the required items, then click [Save].



#### Set the following content:

Definition name	Setting value
Data Source	Specify the data source name to be registered in the ODBC driver manager. The application will select the name specified here and connect with the FUJITSU Enterprise Postgres database. This parameter cannot be omitted. Specify the following characters up to 32 bytes.
	- National characters
	- Alphanumerics
	- "_", "<", ">", "+", "`", " ", "~", """ "&", """, "#", "\$", "%", "-", "^", ":", "/", ":",
Description	Specify a supplementary description for the current data source. Specify characters up to 255 bytes.
	- National characters
	- Alphanumerics
Database	Specify the database name to be connected.
SSLMode	Specify to encrypt communications. The default is "disable".
	The setting values for SSLMode are as follows:
	- disable: Connect without SSL
	- allow: Connect without SSL, and if it fails, connect using SSL
	- prefer: Connect using SSL, and if it fails, connect without SSL
	- require: Connect always using SSL
	- verify-ca: Connect using SSL, and use a certificate issued by a trusted CA (*1)
	- verify-full: Connect using SSL, and use a certificate issued by a trusted CA to verify if the server host name matches the certificate (*1)

Definition name	Setting value
Server	Specify the host name of the database server to connect to, using up to 63 bytes. This parameter cannot be omitted.
Port	Specify the port number to be used for remote access. The default value is "27500".
Username(*2)	Specify the user that will access the database.
Password(*2)	Specify the password for the user that will access the database.

<sup>\*1:</sup> If specifying either "verify-ca" or "verify-full", use the system environment variable PGSSLROOTCERT of your operating system to specify the CA certificate file as shown below.

#### Example:

```
Variable name: PGSSLROOTCERT
Variable value: cACertificateFile
```

### 3.2.2.2 Registering Using Commands

This section describes how to use commands to register ODBC data sources.

Use the following tools from Microsoft to register ODBC data sources.

- ODBCConf.exe
- Add-OdbcDsn

Refer to the Microsoft Developer Network (MSDN) Library for information on how to use these tools.

### When using ODBCConf.exe

ODBCConf.exe is a tool supported on all Windows(R) platforms.

#### Specification format

```
ODBCConf.exe /A { dataSourceType "odbcDriverName" "optionName=value[| optionName=value...]"} [/Lv fileName]
```

Refer to the Microsoft MSDN library for information on the format and parameters.

#### Description

Set the following content:

Definition name	Setting value
Data source type	<ul> <li>Specify the data source type.</li> <li>"CONFIGSYSDSN": A system data source is created. This requires user admin rights. The data source can be used by all users of the same computer.</li> <li>"CONFIGDSN": A user data source is created. The data source can be used by the current user only.</li> <li>When CONFIGSYSDSN is specified as the data source type, it is necessary to execute the command in the command prompt in administrator mode.</li> </ul>

<sup>\*2:</sup> In consideration of security, specify the Username and the Password by the application.

Definition name	Setting value			
ODBC driver name	Specify an ODBC driver name that has already been registered on the system.  Specify one of the following.  Note  The placeholders shown below are enclosed in angle brackets '<>' to avoid confusion with literal text. Do not include the angle brackets in the string.			
	<ul> <li>"FUJITSU Enterprise Postgres Unicode</li></ul>			
Option name	The following items must be set:  - "DSN": Specify the data source name.  - "Servername": Specify the host name for the database server.  - "Port": Specify the port number for connection to the database  - "Database": Specify the database name.  Specify the following values as required:  - "UID": User ID  - "Password": Password  - "SSLMode": Specify to encrypt communications. The default is "disable".  Refer to the SSLMode explanation in the table under step 5 of "3.2.2.1  Registering Using GUI" for information on how to configure SSLMode.			
File Name	You can output process information to a file when creating a data source. This operand can be omitted.			



ODBCConf.exe /A {CONFIGSYSDSN "FUJITSU Enterprise Postgres Unicode 12" "DSN=odbcconf1| Servername=sv1|Port=27500|Database=db01|SSLMode=verify-ca"} /Lv log.txt

.....



In consideration of security, specify the UID and the Password by the application.

#### When using Add-OdbcDsn

Add-OdbcDsn is used in the PowerShell command interface.

#### Specification format

Add-OdbcDsn dataSourceName -DriverName "odbcDriverName" -DsnType dataSourceType - Platform osArchitecture -SetPropertyValue @("optionName=value" [,"optionName=value"...])

Refer to the Microsoft MSDN library for information on the format and parameters.

#### Description

Set the following content:

Definition name	Setting value			
Data source name	Specify any name for the data source name.			
ODBC driver name	Specify an ODBC driver name that has already been registered on the system. Specify one of the following.			
	<b>G</b> Note			
	The placeholders shown below are enclosed in angle brackets '<>' to avoid confusion with literal text. Do not include the angle brackets in the string.			
	- "FUJITSU Enterprise Postgres Unicode <fujitsuenterprisepostgresclientvers>"  Specify this driver name if using Unicode as the application encoding system.</fujitsuenterprisepostgresclientvers>			
	- "FUJITSU Enterprise Postgres ANSI <pre></pre>			
Data source	Specify the data source type.			
type	- "System": A system data source is created. Requires user admin rights. The data source can be used by all users of the same computer.			
	- "User": A user data source is created. The data source can be used by the current user only.			
	<b>€</b> Note			
	When System is specified as the data source type, it is necessary to execute the command in the administrator mode of the command prompt.			
OS	Specify the OS architecture of the system.			
architecture	- "32-bit": 32-bit system			
	- "64-bit": 64-bit system			
Option name	The following items must be set:			
	- "Servername": Specify the host name for the database server.			
	- "Port": Specify the port number for connection to the database			
	- "Database": Specify the database name.			
	Specify the following values as required:			
	- "SSLMode": Specify to encrypt communications. The default is "disable".  Refer to the SSLMode explanation in the table under step 5 of "3.2.2.1  Registering Using GUI" for information on how to configure SSLMode.			

Definition name	Setting value
	When using Add-OdbcDsn, the strings "UID" and "Password" cannot be set as option names. These can only be used when using ODBCConf.exe.



Add-OdbcDsn odbcps1 -DriverName "FUJITSU Enterprise Postgres Unicode 12" -DsnType System -Platform 32-bit -SetPropertyValue @("Servername=sv1", "Port=27500", "Database=db01", "SSLMode=verify-ca")

### 3.2.3 Registering ODBC Data Sources(for Linux)

This section describes how to register ODBC data sources on Linux.

1. Register the data sources

Edit the odbc.ini definition file for the data source.



Edit the file in the installation directory for the ODBC driver manager (unixODBC)

unixOdbcInstallDir/etc/odbc.ini

Or

Create a new file in the HOME directory



#### Point

If unixOdbcInstallDir is edited, these will be used as the shared settings for all users that log into the system. If created in the HOME directory ( $\sim$ /), the settings are used only by the single user.

Set the following content:

Definition name	Setting value	
[Data source name]	Set the name for the ODBC data source.	
Description	Set a description for the ODBC data source. Any description may be set.	
Driver	Set the following as the name of the ODBC driver. Do not change this value.	
	Select the two strings below that correspond to the application type. Concatenate the strings with no spaces and then specify this as the driver name.	
	<b>G</b> Note	
	The placeholders shown below are enclosed in angle brackets '<>' to avoid confusion with literal text. Do not include the angle brackets in the string.	

Definition name	Setting value			
	- Application architecture			
	"FUJITSU Enterprise Postgres < fujitsuEnterprisePostgresClientVers>x64"			
	- Encoding system used by the application			
	- In Unicode (only UTF-8 can be used)			
	"unicode"			
	- Other than Unicode			
	"ansi"			
	Example: The encoding system used by the application is Unicode:			
	"FUJITSU Enterprise			
	Postgres < fujitsuEnterprisePostgresClientVers>x64unicode"			
Database	Specify the database name to be connected.			
Servername	Specify the host name for the database server.			
Username	Specify the user ID that will connect with the database.			
Password	Specify the password for the user that will connect to the database.			
Port	Specify the port number for the database server.			
	The default is "27500".			
SSLMode	Specify the communication encryption method. The setting values for SSLMode are as follows:			
	- disable: Connect without SSL			
	- allow: Connect without SSL, and if it fails, connect using SSL			
	- prefer: Connect using SSL, and if it fails, connect without SSL			
	- require: Connect always using SSL			
	- verify-ca: Connect using SSL, and use a certificate issued by a trusted CA (*1)			
	<ul> <li>verify-full: Connect using SSL, and use a certificate issued by a trusted CA to verify if the server host name matches the certificate (*1)</li> </ul>			
ReadOnly	Specify whether to set the database as read-only.			
	- 1: Set read-only			
	- 0: Do not set read-only			

<sup>\*1:</sup> If specifying either "verify-ca" or "verify-full", use the environment variable PGSSLROOTCERT to specify the CA certificate file as shown below.

#### Example

export PGSSLROOTCERT=cACertificateFileStorageDir/root.crt



[MyDataSource]

Description = FUJITSU Enterprise Postgres

Driver = FUJITSU Enterprise Postgres12x64ansi

Database = db01 Servername = sv1 Port = 27500ReadOnly = 0



In consideration of security, specify the UserName and the Password by the application.

2. Configure the environment variable settings

To execute applications that use ODBC drivers, all of the following settings must be configured in the LD\_LIBRARY\_PATH environment variable:

- fujitsuEnterprisePostgresClientInstallDir/lib
- unixOdbcInstallDir(\*1)/lib
- libtoolInstallDir(\*1)/lib

\*1: If the installation directory is not specified when unixODBC and libtool are installed, they will be installed in /usr/local.

# 3.2.4 Message Language and Encoding System Used by Applications Settings

This section explains the language settings for the application runtime environment and the encoding settings for the application.

#### Language settings

You must match the language settings for the application runtime environment with the message locale settings of the database server.

Messages output by an application may include text from messages sent from the database server. In the resulting text, the text of the application message will use the message locale of the application, and the text of the message sent by the database server will use the message locale of the database server. If the message locales do not match, more than one language or encoding system will be used. Moreover, if the encoding systems do not match, characters in the resulting text can be garbled.

- Linux

Set the locale for messages (LC\_MESSAGES category) to match the message locale of the database server. This can be done in a few different ways, such as using environment variables. Refer to the relevant manual of the operating system for information on the setlocale function.



Example of specifying "en\_US.UTF-8" with the setlocale function

setlocale(LC\_ALL, "en\_US.UTF-8");

Specifying the locale of the LC\_ALL category propagates the setting to LC\_MESSAGE.

- Windows(R)

W

Align the locale of the operating system with the message locale of the database server.

#### **Encoding System Settings**

Ensure that the encoding system that is embedded in the application and passed to the database, and the encoding system setting of the runtime environment, are the same. The encoding system cannot be converted correctly on the database server.

Use one of the following methods to set the encoding system for the application:

- Set the PGCLIENTENCODING environment variable in the runtime environment.
- Set the client\_encoding keyword in the connection string.
- Use the PQsetClientEncoding function.



#### See

Refer to "Supported Character Sets" in "Server Administration" in the PostgreSQL Documentation for information on the strings that represent the encoding system that can be set.

For example, when using "Unicode" and "8 bit", set the string "UTF8".



#### Setting the "PGCLIENTENCODING" environment variable

An example of setting when the encoding of the client is "UTF8" (Bash)

> PGCLIENTENCODING=UTF8; export PGCLIENTENCODING

W

An example of setting when the encoding of the client is "UTF8"

> set PGCLIENTENCODING=UTF8



Text may be garbled when outputting results to the command prompt. Review the font settings for the command prompt if this occurs.

......

## 3.3 Connecting to the Database

Refer to the manual for the programming language corresponding to the ODBC interface, i.e. Access, Excel, or Visual Basic, for example.

### 3.4 Application Development

This section describes how to develop applications using ODBC drivers.

#### W

### 3.4.1 Compiling Applications (for Windows (R))

Refer to the manual for the programming language corresponding to the ODBC interface, i.e. Access, Excel, or Visual Basic, for example.



The cl command expects input to be a program that uses one of the following code pages, so convert the program to these code pages and then compile and link it (refer to the Microsoft documentation for details).

- ANSI console code pages (example: UTF8)
- UTF-16 little-endian with or without BOM (Byte Order Mark)
- UTF-16 big-endian with or without BOM
- UTF-8 with BOM

The cl command converts strings in a program to an ANSI console code page before generating a module, so the data sent to and received from the database server becomes an ANSI console code page. Therefore, set the coding system corresponding to the ANSI console code page as the coding system of the client.

Refer to "Character Set Support" in "Server Administration" in the PostgreSQL Documentation for information on how to set the client encoding system.

### 3.4.2 Compiling Applications (for Linux)

Specify the following options when compiling applications.

Table 3.1 Include file and library path

Architecture	Option	How to specify the option
64-bit	Path of the include file	-I unixOdbc64bitIncludeFileDir
	Path of the library	-L unixOdbc64bitLibraryDir

#### Table 3.2 ODBC library

Type of library	Library name
Dynamic library	libodbc.so



Specify -m64 when creating a 64-bit application.



The following are examples of compiling ODBC applications:

gcc -m64 -I/usr/local/include(\*1) -L/usr/local/lib(\*1) -lodbc testproc.c -o testproc

\*1: This is an example of building and installing from the source without specifying an installation directory for unixODBC. If you wish to specify a location, set the installation directory.

### 3.4.3 Creating Applications While in Database Multiplexing Mode

This section explains points to consider when creating applications while in database multiplexing mode.



- Refer to the Cluster Operation Guide (Database Multiplexing) for information on database multiplexing mode.

- Refer to "Application Development" in the Cluster Operation Guide (PRIMECLUSTER) for points to consider when creating applications using the failover feature integrated with the cluster software.

## 3.4.3.1 Errors when an Application Connection Switch Occurs and Corresponding Actions

If an application connection switch occurs while in database multiplexing mode, explicitly close the connection and then reestablish the connection or reexecute the application.

The table below shows errors that may occur during a switch, and the corresponding action to take.

State		Error information (*1)	Action
Server failure or FUJITSU Enterprise Postgres system failure	Failure occurs during access	57P01 08S01	After the switch is complete, reestablish the connection, or reexecute the application.
	Accessed during system failure	08001	
Switch to the standby server	Switched during access	57P01 08S01	
	Accessed during switch	08001	

<sup>\*1:</sup> Return value of SQLSTATE.

## Chapter 4 .NET Data Provider

This chapter describes how to configure for the purpose of creating .NET applications with Visual Studio.

### 4.1 Development Environment

.NET Data Provider can operate in the following environments:

.NET Framework environment for the development and running of applications	.NET Framework 4.8 .NET Framework 4.7/4.7.x .NET Framework 4.6/4.6.x .NET Framework 4.5.1 or later	
Integrated development environment for applications running in a .NET Framework environment	Visual Studio 2019 Visual Studio 2017 Visual Studio 2015	
Combinations when TableAdapter is used	Visual Studio 2019	.NET Framework 4.8 .NET Framework 4.7/4.7.x .NET Framework 4.6/4.6.x .NET Framework 4.5.1 or later
	Visual Studio 2017	.NET Framework 4.8 .NET Framework 4.7/4.7.x .NET Framework 4.6/4.6.x .NET Framework 4.5.1 or later
	Visual Studio 2015	.NET Framework 4.8 .NET Framework 4.7/4.7.x .NET Framework 4.6/4.6.x .NET Framework 4.5.1 or later
Available development languages	C# Visual Basic .NET	

### 4.2 Setup

This section explains how to set up .NET Data Provider and Npgsql for Entity Framework.

### 4.2.1 Setting Up the Visual Studio Integration Add-On

A user with administrator privileges can register Npgsql Development Tools for .NET as an add-on installing the VSIX package provided. Note that Visual Studio must already be installed in the system prior to installing the VSIX package.

Location of VSIX binaries

The Npgsql.vsix setup package is stored in the following location

 $fujitsu \textit{EnterprisePostgresClientInstallDir \DOTNET \Npgsql.vsix}$ 

Using Npgsql.vsix

Navigate to the Npgsql.vsix binary directory and double-click the package to install it.

> Npgsql.vsix



Refer to "4.5.1 Uninstalling Npgsql" for details.

### 4.2.2 Setting Up .NET Data Provider

FUJITSU Enterprise Postgres has utilized Microsoft best practices of VSIX technology to integrate .NET Data Provider with Visual Studio.

There is no need to explicitly add a reference to .NET Data Provider when using VSIX. This will be done automatically when a database object is created and added to a project through the Visual Studio Server Explorer.



Additional setup is needed if using ProviderFactory to connect to a database in a multi-version installation. Refer to "2.4 Registering .NET Data Provider" in the Installation and Setup Guide for Client.



The following name will be displayed in [References] in Visual Studio Solution Explorer once a database object has been created and added to the project. Note that the Npgsql reference is automatically added when the new database object is first compiled.

- Npgsql

### 4.2.3 Setting Up .NET Data Provider Type Plugins

FUJITSU Enterprise Postgres .NET Data Provider now comes packaged with 6 Type Plugins that provide additional support for more data type mappings (eg date time support via the Npgsql.NodaTime). The plugins modify how Npgsql maps the PostgreSQL values to CLR types.

The type plugins are available for installation as a local NuGet packages. There are 6 pacakages available for installation. "<x>" indicates the product version.

- Npgsql.NodaTime.<x>.0.0.nupkg
- Npgsql.Json.NET.<x>.0.0.nupkg
- Npgsql.NetTopologySuite.<x>.0.0.nupkg: Spatial type(\*1)
- Npgsql.GeoJSON.<x>.0.0.nupkg: Spatial type(\*1)
- Npgsql.LegacyPostgis.<x>.0.0.nupkg: Spatial type(\*1)
- Npgsql.RawPostgis.<x>.0.0.nupkg: Spatial type(\*1)

\*1: Please note that the Spatial Type plugins require the PostGIS extension installed on the server.

Also, Refre to "Type Plugins" and "Additional Notes on each Type Plugin" on the each type plugins.

To install any of the plugins please follow the procedure below:

#### Location of NuGet package

The Plugin NuGet package are stored in the following location:

fujitsuEnterprisePostgresClientInstallDir\DOTNET\Npgsql.\*.<x>.0.0.nupkg

#### Add a local package source

In Visual Studio, add a NuGet local package source if one does not exist.

- 1. Click [Tools] >> [Options] >> [NuGet Package Manager], and then select [Package Sources].
- 2. Click [+] in the upper-right corner, and then set [Name] to "Local Package Source".
- 3. Click [...] and navigate to the folder above. Select this folder, and then click [OK].

#### Install the NuGet package

In Visual Studio, install the NuGet package from the local package source.

- 1. Click [Tools] >> [NuGet Package Manager] >> [Manage NuGet Packages for Solution].
- 2. In the upper-right corner, select "Local Package Source" from [Package Source].
- Once the local package source is set, all available NuGet packages in this local location will be displayed. Select the plugin to be installed (eg "Npgsql.NodaTime", and then select the projects for which this package is to be installed.
- 4. Click [Install].

### 4.2.4 Setting Up Npgsql for Entity Framework

Npgsql for Entity Framework is supplied as a NuGet package file. To install it locally, follow the procedure below.

#### Location of NuGet package

The EntityFramework6.Npgsql NuGet package is stored in the following location:

 $fujits u \textit{EnterprisePostgresClientInstallDir} \\ \texttt{DOTNET} \\ \texttt{EntityFramework6.npgsql.3.2.1.nupkg}$ 

#### Add a local package source

In Visual Studio, add a NuGet local package source if one does not exist.

- 1. Click [Tools] >> [Options] >> [NuGet Package Manager], and then select [Package Sources].
- 2. Click [+] in the upper-right corner, and then set [Name] to "Local Package Source".
- 3. Click [...] and navigate to the folder above. Select this folder, and then click [OK].

#### Install the NuGet package

In Visual Studio, install the NuGet package from the local package source.

- 1. Click [Tools] >> [NuGet Package Manager] >> [Manage NuGet Packages for Solution].
- 2. In the upper-right corner, select "Local Package Source" from [Package Source].
- 3. Once the local package source is set, all available NuGet packages in this local location will be displayed. Select "EntityFramework6.Npgsql", and then select the projects for which this package is to be installed.
- 4. Click [Install].

### 4.2.5 Message Language Settings

You must match the language settings for the application runtime environment with the message locale settings of the database server.

Set language using the "System.Globalization.CultureInfo.CreateSpecificCulture" method.



Code example for changing the locale in a C# application

System.Threading.Thread.CurrentThread.CurrentUICulture =
System.Globalization.CultureInfo.CreateSpecificCulture("en");

### 4.3 Connecting to the Database

This section explains how to connect to a database.

- Using NpgsqlConnection

.....

- Using NpgsqlConnectionStringBuilder
- Using the ProviderFactory Class

### 4.3.1 Using NpgsqlConnection

Connect to the database by specifying the connection string.



Code examples for applications

```
using Npgsql;

NpgsqlConnection conn = new NpgsqlConnection("Server=sv1;Port=27500;Database=mydb;
Username=myuser;Password=myuser01; Timeout=20;CommandTimeout=20;");
```

Refer to "4.3.4 Connection String" for information on database connection strings.

### 4.3.2 Using NpgsqlConnectionStringBuilder

Generate connection strings by specifying the connection information in the properties of the NpgsqlConnectionStringBuilder object.



Code examples for applications

```
using Npgsql;

NpgsqlConnectionStringBuilder sb = new NpgsqlConnectionStringBuilder();
sb.Host = "sv1";
sb.Port = 27500;
sb.Database = "mydb";
sb.Username = "myuser";
sb.Password = "myuser01";
sb.Timeout = 20;
sb.CommandTimeout = 20;
NpgsqlConnection conn = new NpgsqlConnection(sb.ConnectionString);
```

Refer to "4.3.4 Connection String" for information on database connection strings.

### 4.3.3 Using the ProviderFactory Class

Obtain the DbConnection object from the provider factory.



Code examples for applications

```
using System.Data.Common;

DbProviderFactory factory = DbProviderFactories.GetFactory("FUJITSU.Npgsql");
DbConnection conn = factory.CreateConnection();
conn.ConnectionString = "Server=sv1;Port=27500;Database=mydb;
Username=myuser;Password=myuser01; Timeout=20;CommandTimeout=20;";
```

### 4.3.4 Connection String

Specify the following connection information to connect to the database.



- (1) Specify the host name or IP address of the server to be connected. This must be specified.
- (2) Specify the port number for the database server. The default is "27500".
- (3) Specify the database name to be connected.
- (4) Specify the username that will connect with the database.
- (5) Specify the password for the user that will connect to the database.
- (6) Refer to the following for information on how to specify other connection information.

The table below shows keywords that are available to specify in the connection string in .NET Data Provider (Npgsql):

Note that some settings require care if using an Oracle database-compatible feature (refer to "9.2.2 Notes when Integrating with the Interface for Application Development" for details).

Keyword	Default value	Description	
Host	None	Specify the host name or IP address of the server to be connected.	
		Specify up to 63 bytes when specifying a host name.	
		A host name or IP address must be specified.	
Port	27500	Specify the port number for the database server.	
Username	None	Specify the username that will connect with the database. Not required if using Integrated Security.	
Password	None	Specify the password for the username that will connect to the database. Not required if using Integrated Security.	
Database	Username	Specify the database name to be connected.	
Search Path		Specify the default schema name of the SQL statements used in the application.	
Timeout	15	Specify the timeout for connections.	
		Specify a value between 0 and 1024 (in seconds). The default is 15 seconds. An error occurs when a connection cannot be established within the specified time.	
Connection Idle Lifetime	300	Specify the time to wait before closing idle connections in the pool if the count of all connections exceeds Minimum Pool Size.	
Pooling	true	Specify whether to use connection pooling.	
		- Connection pooling is used if you specify true.	
		- Connection pooling is not used if you specify false.	
Maximum Pool Size	100	Maximum size of a connection pool. If the request exceeds this limit, it will wait until another connection closes and the pool is available.	
		Specify in the range between 0 and 1024.	

Keyword	Default value	Description	
Minimum Pool Size	1	Minimum size of a connection pool. When you specify Minimum Pool Size, NpgsqlConnection will pre-allocate connections with the specified number of servers.	
		Specify in the range from 0 to the value specified at Maximum Pool Size.	
SSL Mode	Disable	Specify one of the following values for the SSL connection control mode:	
		- Prefer: SSL is used for connection wherever possible.	
		- Require: An exception is thrown when SSL connection is not possible.	
		- Disable: SSL connection is not performed.	
Enlist	true	Specify whether to have connections participate in transactions with the transaction scope declared:	
		- Connections will participate in transactions when true is specified.	
		<ul> <li>Connections will not participate in transactions when false is specified.</li> </ul>	
Command Timeout	30	Specify the timeout for communication with the server.	
		Specify a value between 0 and 2147483647 (in seconds). There is no limit set if you set 0. An error occurs when data is not received from the server within the specified time.	
Integrated Security	false	Set this when using Windows Integrated Security.	
Trust Server Certificate	false	Whether to trust the server certificate without validating it.	
Use SSL Stream	false	Npgsql uses its own internal implementation of TLS/SSL. Turn this on to use .NET SslStream instead.	
Check Certificate Revocation	false	Whether to check the certificate revocation list during authentication.	
Persist Security Info	false	Gets or sets a Boolean value that indicates if security-sensitive information, such as the password, is not returned as part of the connection if the connection is open or has ever been in an open state.	
Kerberos Service Name	postgres	The Kerberos service name to be used for authentication.	
Include Realm		The Kerberos realm to be used for authentication.	
Connection Pruning Interval	10	How many seconds the pool waits before attempting to prune idle connections that are beyond idle lifetime	
Internal Command Timeout	-1	The time to wait (in seconds) while trying to execute an internal command before terminating the attempt and generating an error1 uses Command Timeout, 0 means no timeout.	
Keepalive	disabled	The number of seconds of connection inactivity before Npgsql sends a keepalive query.	
Tcp Keepalive Time	disabled	The number of milliseconds of connection inactivity before a TCP keepalive query is sent. Use of this option is discouraged, use Keepalive instead if possible. Supported only on Windows.	

Keyword	Default value	Description	
Tcp Keepalive Interval	Tcp Keepalive Time	The interval, in milliseconds, between when successive keep-alive packets are sent if no acknowledgement is received. Tcp Keepalive Time must be non-zero as well. Supported only on Windows.	
Application Name		Optional application name parameter to be sent to the backend during connection initiation.	
Client Encoding		Sets the client_encoding parameter.	
EF Template Database	template1	The database template to specify when creating a database in Entity Framework.	
Max Auto Prepare	0	The maximum number of SQL statements that can be automatically prepared at any given point. Beyond this number the least-recently-used statement will be recycled. Zero disables automatic preparation.	
Auto Prepare Min Usages	5	The minimum number of usages an SQL statement is used before it is automatically prepared.	
Use Perf Counters	false	Makes Npgsql write performance information about connection use to Windows Performance Counters. Supported only on Windows.	
		Using performance counters first involves setting them up on your Windows system. To do this you will need to install Npgsql's MSI and ensure that the Performance Counters option is installed.	
		In addition, you will need to pass "Use Perf Counters=true" on your connection string. Once you start your Npgsql application with this addition, you should start seeing real-time data in the Performance Monitor.	
Read Buffer Size	8192	Size of the internal buffer Npgsql uses when reading. Increasing may improve performance if transferring large values from the database.	
Write Buffer Size	8192	Size of the internal buffer Npgsql uses when writing. Increasing may improve performance if transferring large values to the database.	
Socket Receive Buffer Size	System dependent	Size of socket receive buffer.	
Socket Send Buffer Size	System dependent	Size of socket send buffer.	

## 4.4 Application Development

This section explains the range of support provided with Visual Studio integration.

### 4.4.1 Data Types

A variety of data types can be used with FUJITSU Enterprise Postgres.

The data types below are supported whether you automatically generate applications using tools in Visual Studio (Query Builder in TableAdapter and Server Explorer), or create applications yourself (with DataProvider).

Table 4.1 List of supported data types

Table 4.1 Elst of supported data types	Supported	
Data Types	Operation in the Visual Studio integration window	Fujitsu Npgsql .NET Data Provider
character	Y	Y

	Supported		
Data Types	Operation in the Visual Studio integration window	Fujitsu Npgsql .NET Data Provider	
character varying	Y	Y	
national character	Y	Y	
national character varying	Y	Y	
text	Y	Y	
bytea	N	N	
smallint	Y	Y	
integer	Y	Y	
bigint	Y	Y	
smallserial	N	N	
serial	N	N	
bigserial	N	N	
real	Y	Y	
double precision	Y	Y	
numeric	Y	Y	
decimal	Y	Y	
money	N	Y	
date	Y	Y	
time with time zone	Conditional (*1)	Conditional (*1)	
time without time zone	Conditional (*1)	Conditional (*1)	
timestamp without time zone	Y	Y	
timestamp with time zone	Y	Y	
interval	Conditional (*2)	Conditional (*2)	
boolean	Y	Y	
bit	N	N	
bit varying	N	N	
uuid	Y	Y	
inet	N	Conditional (*3)	
macaddr	N	Y	
macaddr8	N	Y	
cidr	N	Conditional (*4)	
Geometric data type (point,lseg,box,path,polygon,circle)	N	Y	
аггау	N	Y	
oid	N	N	
xml	N	Y	
json	N	N	
Types related to text searches(tsvector,tsquery)	N	N	

	Supported		
Data Types	Operation in the Visual Studio integration window	Fujitsu Npgsql .NET Data Provider	
Enumerated type	N	N	
Composite type	N	N	
Range type	N	N	

Y: Supported

N: Not supported

\*1: As shown below, "time with time zone" and "time without time zone" values display the date portion as additional information. However, the actual data comprises the time data only, so with the exception of this displayed format, there are no other resulting issues.

#### Example:

Composition of table (t1)

col1 (time with time zone)	col2 (time without time zone)	
1/01/0001 10:21:30 +08:00	10:21:30	
1/01/0001 23:34:03 +08:00	23:34:03	
1/01/0001 17:23:54 +08:00	17:23:54	

"time with time zone" values display a fixed value of "2/01/0001" in the date portion, while "time without time zone" values display just the time data.

# 4.4.2 Relationship between Application Data Types and Database Data Types

The data types available for SQL data types are as follows:

DbType	NpgsqlDbType	PostgreSQL type	Accepted C# types	.Net Framework Type
	Char	char	string, char[], char, IConvertible	System.String, System.Char[], System.Char

<sup>\*2:</sup> The format is d.hh:mm:ss, where d is an integer and hh:mm:ss is a maximum of 23.59.59 (23 hours, 59 minutes, and 59 seconds).

<sup>\*3:</sup> When updating inet types, only a single host is supported. The input format is addr/y where addr is an IPv4 or IPv6 address and y is the number of bits in the netmask. If y is omitted, the number of bits in the netmask is set to 32 for an IPv4 address and 128 for an IPv6 address. On display, the y portion is suppressed.

<sup>\*4:</sup> When updating cidr types, only a single host is supported.

DbType	NpgsqlDbType	PostgreSQL type	Accepted C# types	.Net Framework Type
	Varchar	varchar	string, char[], char, IConvertible	System.String, System.Char[], System.Char
String, StringFixedLength, AnsiString, AnsiStringFixedLe ngth	Text	text	string, char[], char, IConvertible	System.String, System.Char[], System.Char
Binary	Bytea	bytea	byte[], ArraySegment	System.Byte[]
Int16	Smallint	int2	short, IConvertible	System.Int16
Int32	Integer	int4	int, IConvertible	System.Int32
Int64	Bigint	int8	long, IConvertible	System.Int64
Single	Real	float4	float, IConvertible	System.Single
Double	Double	float8	double, IConvertible	System.Double
Decimal, VarNumeric	Numeric	numeric	decimal, IConvertible	System.Decimal
Date	Date	date	DateTime, NpgsqlDate, IConvertible	System.DateTime
	TimeTZ	timetz	DateTimeOffset, DateTime, TimeSpan	System.DateTimeOffs et, System.DateTime, System.TimeSpan
Time	Time	time	TimeSpan, string	System.Timespan
DateTime, DateTime2	Timestamp	timestamp	DateTime, DateTimeOffset, NpgsqlDateTime, IConvertible	System.DateTime
DateTimeOffset	TimestampTZ	timestamptz	DateTime, DateTimeOffset, NpgsqlDateTime, IConvertible	System.DateTime
	Interval	interval	TimeSpan, NpgsqlTimeSpan, string	System.TimeSpan
	Boolean	bool	bool, IConvertible	System.Boolean
	Bit	bit	BitArray, bool, string	System.Boolean, System.String
	Uuid	uuid	Guid, string	System.Guid
	Inet	inet	IPAddress, NpgsqlInet	ValueTuple <ipaddres s, int&gt;, IPAddress, NpgsqlInet</ipaddres 
	MacAddr	macaddr	PhysicalAddress	System.Net.NetworkI nformation.PhysicalA ddress

DbType	NpgsqlDbType	PostgreSQL type	Accepted C# types	.Net Framework Type
	MacAddr8	macaddr8	PhysicalAddress	System.Net.NetworkI nformation.PhysicalA ddress
	Box	box	NpgsqlBox	NpgsqlBox
	Circle	circle	NpgsqlCircle	NpgsqlCircle
	LSeg	lseg	NpgsqlLSeg	NpgsqlLSeg
	Path	path	NpgsqlPath	NpgsqlPath
	Point	point	NpgsqlPoint	NpgsqlPoint
	Polygon	polygon	NpgsqlPolygon	NpgsqlPolygon
	Array	array types	Array, IList, IList	System.Array

## 4.4.3 Creating Applications while in Database Multiplexing Mode

This section explains points to consider when creating applications while in database multiplexing mode.



#### See

- Refer to the Cluster Operation Guide (Database Multiplexing) for information on database multiplexing mode.
- Refer to "Application Development" in the Cluster Operation Guide (PRIMECLUSTER) for points to consider when creating applications using the failover feature integrated with the cluster software.

.....

# 4.4.3.1 Errors when an Application Connection Switch Occurs and Corresponding Actions

If an application connection switch occurs while in database multiplexing mode, explicitly close the connection and then reestablish the connection or reexecute the application.

The table below shows errors that may occur during a switch, and the corresponding action to take.

State		Error information	Action
Server failure or FUJITSU Enterprise Postgres	Failure occurs during access	57P01 (*1) Empty string (*1)	After the switch is complete, reestablish the connection, or reexecute the application.
system failure	Accessed during system failure	Empty string (*1)	
Switch to the standby server	Switched during access	57P01 (*1) Empty string (*1)	
	Accessed during switch	Empty string (*1)	

<sup>\*1:</sup> This is the return value of the PostgresException attribute SqlState.

#### **4.4.4 Notes**

#### Notes on TableAdapter

- If [SELECT which returns a single value] is selected when adding a query to a TableAdapter, it will not be possible to execute the SQL statement displayed on the window - therefore, correct the SQL statement.

#### **Type Plugins**

- These type libraries include:
  - NodaTime the recommended way to interact with PostgrSQL date/time types
  - Json.NET allows Npgsql to use the NewtonSoft Json.NET library when reading and writing JSON data (both json and jsonb)
  - NetTopologySuite allows Npgsql to map PostGIS spatial types directly to the NetTopology suite types (the leading spatial library in .NET)
  - GeoJSON allows Npgsql to read and write PostGIS spatial types as GeoJSON types via the GeoJSON.NET library
  - LegacyPostgis implements the previously supported PostGIS spatial types as a plugin and is limited to geometry and XY only. Gerography and XYZ, XYM and XYZM is not supported in this plugin and if required, use the NetTopologySuite plugin
  - RawPostgis allows raw byte access to PostGIS types
- Setup the plugin in your application simply by adding a dependency on the plugin (this should have been done automatically when installed to the project) and set it up. See the following code snippet for an example of setting up the Npgsql.NodaTime plugin:

```
using Npgsql;
  // Place this at the beginning of your program to use NodaTime everywhere
(recommended)
  NpgsqlConnection.GlobalTypeMapper.UseNodaTime();
  // Or to temporarily use NodaTime on a single connection only:
  conn.TypeMapper.UseNodaTime();
```

Once the plugin is setup, you can read and write NodaTime objects as per the code snippet below:

```
// Write NodaTime Instant to PostgreSQL "timestamp without time zone"
    using (var cmd = new NpgsqlCommand(@"INSERT INTO mytable (my_timestamp) VALUES

(@p)", conn))
    {
        cmd.Parameters.Add(new NpgsqlParameter("p", Instant.FromUtc(2011, 1, 1, 10,

30)));
        cmd.ExecuteNonQuery();
    }

    // Read timestamp back from the database as an Instant
    using (var cmd = new NpgsqlCommand(@"SELECT my_timestamp FROM mytable", conn))
    using (var reader = cmd.ExecuteReader())
    {
        reader.Read();
        var instant = reader.GetFieldValue<Instant>(0);
    }
}
```

- To apply the type plugin updates, do one of the following:
  - After uninstalling the type plugin (Refer to "4.5.2 Uninstalling .NET Data Provider Type Plugins"), setup the type plugin (Refer to "4.2.3 Setting Up .NET Data Provider Type Plugins").
  - Remove the type plugin directory from the packages directory of the solution, and then restore it using the nuget restore command.
- When you deploy an application with a type plugin, the type plugin is included in the distribution. Therefore, after applying the type plugin updates, you must rebuild the application and deploy the updated application.

#### **Additional Notes on each Type Plugin**

Describe notes about each type plugin.

#### NodaTime

#### Mapping Table

Describes the mapping of PostgreSQL data types to NodaTime data types.

PostgreSQL Type	Default NodaTime Type	Additional NodaTime Type	Note
timestamp	Instant	LocalDateTime	It's common to store UTC timestamps in databases - you can simply do so and read/write Instant values. You also have the option of readin/writing LocalDateTime, which is a date/time with no information about timezones; this makes sense if you're storing the timezone in a different column and want to read both into a NodaTime ZonedDateTime
timestamp with time zone	Instant	ZonedDateTime, OffsetDateTime	This PostgreSQL type stores only a timestamp, assumed to be in UTC. If you read/write this as an Instant, it will be provided as stored with no timezone conversions whatsoever. If, however, you read/write as a ZonedDateTime or OffsetDateTime, the plugin will automatically convert to and from UTC according to your PostgreSQL session's timezone.
date	LocalDate		A simple date with no timezone or offset information.
time	LocalTime		A simple time-of-day, with no timezone or offset information.
time with time zone	OffsetTime		This is a PostgreSQL type that stores a time and an offset.
interval	Period		This is a human interval which does not have a fixed absolute length ("two months" can vary depending on the months in question), and so it is mapped to NodaTime's Period (and not Duration or TimeSpan).

#### Json.NET

Once the JSON plugin has been setup, users can transparently read and write CLR objects as JSON values and the plugin will automatically serialize/deserialize them

See the code snippet below:

```
// Write arbitrary CLR types as JSON
using (var cmd = new NpgsqlCommand(@"INSERT INTO mytable (my_json_column) VALUES (@p)",
conn))
{
    cmd.Parameters.Add(new NpgsqlParameter("p", NpgsqlDbType.Jsonb) { Value =
    MyClrType });
    cmd.ExecuteNonQuery();
}
```

```
// Read arbitrary CLR types as JSON
using (var cmd = new NpgsqlCommand(@"SELECT my_json_column FROM mytable", conn))
using (var reader = cmd.ExecuteReader())
{
   reader.Read();
   var someValue = reader.GetFieldValue<MyClrType>(0);
}
```

#### NetTopologySuite (spatial)

By default the plugin handles only ordinates provided by the DefaultCoordinateSequenceFactory of GeometryServiceProvider.Instance. If GeometryServiceProvider is initialized automatically the X and Y ordinates are handled. To change the behavior specify the handleOrdinates parameter like in the following example:

```
conn.TypeMapper.UseNetTopologySuite(handleOrdinates: Ordinates.XYZ);
```

To process the M ordinate, you must initialize GeometryServiceProvider.Instance to a new NtsGeometryServices instance with coordinateSequenceFactory set to a DotSpatialAffineCoordinateSequenceFactory. Or you can specify the factory when calling UseNetTopologySuite.

```
// Place this at the beginning of your program to use the specified settings everywhere
(recommended)
GeometryServiceProvider.Instance = new NtsGeometryServices(
    new DotSpatialAffineCoordinateSequenceFactory(Ordinates.XYM),
    new PrecisionModel(PrecisionModels.Floating),
    -1);

// Or specify settings for Npgsql only
conn.TypeMapper.UseNetTopologySuite(
    new DotSpatialAffineCoordinateSequenceFactory(Ordinates.XYM));
```

#### Reading and Writing Geometry Values

When reading PostGIS values from the database, Npgsql will automatically return the appropriate NetTopologySuite types: Point, LineString, and so on. Npgsql will also automatically recognize NetTopologySuite's types in parameters, and will automatically send the corresponding PostGIS type to the database. The following code demonstrates a roundtrip of a NetTopologySuite Point to the database:

```
var point = new Point(new Coordinate(ld, ld));
conn.ExecuteNonQuery("CREATE TEMP TABLE data (geom GEOMETRY)");
using (var cmd = new NpgsqlCommand("INSERT INTO data (geom) VALUES (@p)", conn))
{
    cmd.Parameters.AddWithValue("@p", point);
    cmd.ExecuteNonQuery();
}

using (var cmd = new NpgsqlCommand("SELECT geom FROM data", conn))
using (var reader = cmd.ExecuteReader())
{
    reader.Read();
    Assert.That(reader[0], Is.EqualTo(point));
}
```

You may also explicitly specify a parameter's type by setting NpgsqlDbType.Geometry.

#### Geography (geodetic) Support

PostGIS has two types:geometry (for Cartesian coordinates) and geography (for geodetic or spherical coordinates). You can read about the geometry/geography distinction in the PostGIS docs. In a nutshell, geography is much more

accurate when doing calculations over long distances, but is more expensive computationally and supports only a small subset of the spatial operations supported by geometry.

Npgsql uses the same NetTopologySuite types to represent both geometry and geography - the Point type represents a point in either Cartesian or geodetic space. You usually don't need to worry about this distinction because PostgreSQL will usually cast types back and forth as needed. However, it's worth noting that Npgsql sends Cartesian geometry by default, because that's the usual requirement. You have the option of telling Npgsql to send geography instead by specifying NpgsqlDbType.Geography:

```
using (var cmd = new NpgsqlCommand("INSERT INTO data (geog) VALUES (@p)", conn))
{
   cmd.Parameters.AddWithValue("@p", NpgsqlDbType.Geography, point);
   cmd.ExecuteNonQuery();
}
```

If you prefer to use geography everywhere by default, you can also specify that when setting up the plugin:

```
NpgsqlConnection.GlobalTypeMapper.UseNetTopologySuite(geographyAsDefault: true);
```

#### GeoJSON (spatial)

Using the GeoJSON plugin is the same as the NetTopologuSuite.

#### LegacyPostgis (spatial)

If you've used the internal PostGIS types in Npgsql 3.2 or earlier, the plugin works in the same way:

```
NpgsqlConnection.GlobalTypeMapper.UseLegacyPostgis();

// Write
var cmd = new NpgsqlCommand("INSERT INTO table (pg_point, pg_polygon) VALUES (@point, @polygon)", conn);
cmd.Parameters.AddWithValue("point", new PostgisPoint(3.5, 4.5));
cmd.ExecuteNonQuery();

// Read
var cmd = new NpgsqlCommand("SELECT * FROM table", conn);
var reader = cmd.ExecuteReader();
while (reader.Read()) {
   var point = reader.GetFieldValue<PostgisPoint>(0);
   var polygon = reader.GetFieldValue<PostgisPolygon>(1);
}
```

#### Notes on the Query Builder

- Prefix named parameters with "@".
- Uppercase object names cannot be used, even when enclosed in double quotation marks.
   To use uppercase object names enclosed in double quotation marks, include them in SQL statements and enter these in the [Generate the SQL statements] window rather than in the Query Builder.
- SQL statements cannot be correctly generated if the SQL statement specified in Filter matches any of the conditions below:
  - It uses PostgreSQL intrinsic operators such as << or ::.
  - It uses functions with keywords such as AS, FROM, IN, OVER.

Example: extract(field from timestamp), RANK() OVER

- It uses functions with the same names as those prescribed in SQL conventions, but that require different arguments.

#### **Notes on Server Explorer**

- The temporary table is not displayed.

#### Notes on metadata

- The CommandBehavior.KeyInfo argument must be specified if executing ExecuteReader before obtaining metadata using GetSchemaTable.

.............



```
NpgsqlDataReader ndr=cmd.ExecuteReader(CommandBehavior.KeyInfo);
DataTable dt = dr.GetSchemaTable();
```

#### Notes on automatically generating update-type SQL statements

- If the SQL statement includes a query (which cannot be updated) that matches any of the conditions below, an update-type SQL statement will be generated (note that it may not be possible to execute this SQL statement in some cases):
  - It includes derived tables
  - It includes the same column name as the select list

Update-type SQL statements will be automatically generated in the following cases:

- If update statements are obtained using NpgsqlCommandBuilder
- If data is updated using NpgsqlDataAdapter
- If data is updated using TableAdapter

#### Notes on distributed transactions

- Applications using transaction scope can use distributed transactions by linking with Microsoft Distributed Transaction Coordinator (MSDTC). In this case, note the following:
  - Ensure that the value of max\_prepared\_transactions is greater than max\_connection, so that "PREPARE TRANSACTION" can be issued for each transaction that simultaneously connects to the database server.
  - If each transaction in the transaction scope accesses the same resource using different connections, the database server will perceive it as requests from different applications, and a deadlock may occur. By configuring a timeout value for the transaction scope beforehand, the deadlock can be broken.

# 4.5 Uninstallation

This section explains how to uninstall Npgsql and Npgsql for Entity Framework.

# 4.5.1 Uninstalling Npgsql

To uninstall Npgsql, uninstall each of its components separately:

1. Uninstall DDEX.

DDEX provides the Visual Studio integration tools within the IDE through the Npgsql.VSIX package installation.

- 1. Open Visual Studio 2015.
- 2. Click [Tools], and then [Extensions and Updates].
- 3. Select the FUJITSU Npgsql PostgreSQL Integration extension, and then click [Uninstall].

4. In the confirmation dialog box "Are you sure you want to schedule FUJITSU Npgsql PostgreSQL Integration for uninstall?", click [Yes].

Note that the status at the bottom of the [Extensions and Updates] window will change to "Your changes will be scheduled. The modifications will begin when all Microsoft Visual Studio windows are closed".

- 5. Click [Close].
- 6. Close all Visual Studio instances currently open.

The VSIX Installer will automatically start.

- 7. Click [Modify] to continue with uninstallation of FUJITSU Npgsql PostgreSQL Integration.
- 8. Upon completion, a dialog box will be displayed click [Close].
- 2. Uninstall Npgsql GAC.

Npgsql.dll provides DBProviderFactory functionality for Npgsql.

- 1. Click [Control Panel], and then [Programs and Features].
- 2. Right-click the target program below in the list, and click [Uninstall]. (The notation "x.y.z" indicates the version of the FUJITSU Enterprise Postgres.)

Name: Fujitsu Npgsql .NET Data Provider GAC Register For .NET 4.5

Issuing company: FUJITSU LIMITED

Version: x.y.z

- 3. In the confirmation dialog box "Are you sure you want to uninstall Fujitsu Npgsql .NET Data Provider GAC Register For .NET 4.5?", click [Yes].
- 4. Upon completion, the uninstall window will close, and the uninstallation target version of Fujitsu Npgsql .NET Data Provider GAC Register For .NET 4.5 will no longer be listed.

### 4.5.2 Uninstalling .NET Data Provider Type Plugins

The .NET Data Provider type plugins are installed per project. To uninstall it, follow the procedure below:

- 1. In Visual Studio, open a project for which Npgsql the Plugin to be removed is installed.
- 2. Click [Tools] >> [NuGet Package Manager] >> [Manage NuGet Packages for Solution].
- 3. Select all the projects that have Npgsql Plugin(s) installed, and then click [Uninstall].

  Alternatively, if the plugin packages have been removed and are no longer installed, in the solution explorer open the Dependencies/NuGet node and delete the plugins that require uninstallation.

# 4.5.3 Uninstalling Npgsql for Entity Framework

Npgsql for Entity Framework is installed per project. To uninstall it, follow the procedure below:

- 1. In Visual Studio, open a project for which Npgsql for Entity Framework is installed.
- 2. Click [Tools] >> [NuGet Package Manager] >> [Manage NuGet Packages for Solution].
- 3. Select all the projects that have Npgsql for Entity Framework installed, and then click [Uninstall]. Alternatively, if the Entity Framework package has been removed and is no longer installed, in the solution explorer open the Dependencies/NuGet node and delete the Entity Framework package

# Chapter 5 C Library (libpq)

This chapter describes how to use C libraries.

# 5.1 Development Environment

Install FUJITSU Enterprise Postgres Client Package for the architecture to be developed and executed.



Linux

See

Refer to Installation and Setup Guide for Client for information on the C compiler required for C application development.

## 5.2 Setup

This section describes the environment settings required to use C libraries and how to encrypt data for communication.

### 5.2.1 Environment Settings

To execute an application that uses libpq, set the environment variable as shown below.

L

- Required for compile/link
  - LD\_LIBRARY\_PATH fujitsuEnterprisePostgresClientInstallDir/lib
- Required for execution of the application
  - PGLOCALEDIR fujitsuEnterprisePostgresClientInstallDin/share/locale



"<x>" indicates the product version.

> LD\_LIBRARY\_PATH=/opt/fsepv<x>client64/lib:\$LD\_LIBRARY\_PATH;export LD\_LIBRARY\_PATH
> PGLOCALEDIR=/opt/fsepv<x>client64/share/locale;export PGLOCALEDIR

W

#### Windows (R)

- Required for compile/link
  - LIB fujitsuEnterprisePostgresClientInstallDir\lib
  - INCLUDE fujitsuEnterprisePostgresClientInstallDit\include
- Required for execution of the application
  - PATH fujitsuEnterprisePostgresClientInstallDir\lib
  - PGLOCALEDIR fujitsuEnterprisePostgresClientInstallDir\share\locale



When the 32-bit version client package is installed on a 64-bit operating system. "<x>" indicates the product version.

- > SET PATH=%ProgramFiles(x86)%\Fujitsu\fsepv<x>client32\lib;%PATH%
- > SET LIB=%ProgramFiles(x86)%\Fujitsu\fsepv<x>client32\lib;%LIB%
- > SET INCLUDE=%ProgramFiles(x86)%\Fujitsu\fsepv<x>client32\include;%INCLUDE%
- > SET PGLOCALEDIR=%ProgramFiles(x86)%\Fujitsu\fsepv<x>client32\share\locale

# 5.2.2 Message Language and Encoding System Used by Applications Settings

This section explains the language settings for the application runtime environment and the encoding settings for the application.

#### Language settings

You must match the language settings for the application runtime environment with the message locale settings of the database server.

Messages output by an application may include text from messages sent from the database server. In the resulting text, the text of the application message will use the message locale of the application, and the text of the message sent by the database server will use the message locale of the database server. If the message locales do not match, more than one language or encoding system will be used. Moreover, if the encoding systems do not match, characters in the resulting text can be garbled.

- Linux

Set the locale for messages (LC\_MESSAGES category) to match the message locale of the database server. This can be done in a few different ways, such as using environment variables. Refer to the relevant manual of the operating system for information on the setlocale function.



Example of specifying "en\_US.UTF-8" with the setlocale function

setlocale(LC\_ALL, "en\_US.UTF-8");

Specifying the locale of the LC\_ALL category propagates the setting to LC\_MESSAGE.

Windows(R)

w

Align the locale of the operating system with the message locale of the database server.

#### **Encoding System Settings**

Ensure that the encoding system that is embedded in the application and passed to the database, and the encoding system setting of the runtime environment, are the same. The encoding system cannot be converted correctly on the database server.

Use one of the following methods to set the encoding system for the application:

- Set the PGCLIENTENCODING environment variable in the runtime environment.
- Set the client\_encoding keyword in the connection string.
- Use the PQsetClientEncoding function.

Refer to "Supported Character Sets" in "Server Administration" in the PostgreSQL Documentation for information on the strings that represent the encoding system that can be set.

For example, when using "Unicode" and "8 bit", set the string "UTF8".



Text may be garbled when outputting results to the command prompt. Review the font settings for the command prompt if this occurs.

......

## 5.2.3 Settings for Encrypting Communication Data

Set in one of the following ways when performing remote access using communication data encryption:

When setting from outside with environment variables

Specify "require", "verify-ca", or "verify-full" in the PGSSLMODE environment variable.

In addition, the parameters for the PGSSLROOTCERT and PGSSLCRL environment variables need to be set to prevent spoofing of the database server.



Refer to "Environment Variables" in "Client Interfaces" in the PostgreSQL Documentation for information on environment variables.

When specifying in the connection URI

Specify "require", "verify-ca", or "verify-full" in the "sslmode" parameter of the connection URI.

In addition, the parameters for the sslcert, sslkey, sslrootcert, and sslcrl need to be set to prevent spoofing of the database server.



See

Refer to "Secure TCP/IP Connections with SSL" in "Server Administration" in the PostgreSQL Documentation for information on encrypting communication data.

# 5.3 Connecting with the Database



Point

Use the connection service file to specify the connection destination. In the connection service file, a name (service name) is defined as a set, comprising information such as connection destination information and various types of tuning information set for connections. By using the service name defined in the connection service file when connecting to databases, it is no longer necessary to modify applications when the connection information changes.

Refer to "Client Interfaces", "The Connection Service File" in the PostgreSQL Documentation for details.



Refer to "Database Connection Control Functions" in "Client Interfaces" in the PostgreSOL Documentation.

# 5.4 Application Development



Refer to "libpq - C Library" in "Client Interfaces" in the PostgreSQL Documentation for information on developing applications.

However, if you are using the C library, there are the following differences to the PostgreSQL C library (libpq).

## **5.4.1 Compiling Applications**

Specify the following paths when compiling applications.

Refer to your compiler documentation for information on how to specify the path.

- Linux

Table 5.1 Include file and library path

Type of path	Path name	
Path of the include file	fujitsuEnterprisePostgresClientInstallDir/include	
Path of the library	fujitsuEnterprisePostgresClientInstallDin'lib	

L

Table 5.2 C Library (libpq library)

Type of library	Library name
Dynamic library	libpq.so
Static library	libpq.a

W

Windows(R)

If the include file and the library path have been set in the environment variable, there is no need to specify the paths shown below for the compile.

w

Table 5.3 Include file and library path

i		rable ele illerade ille alla library pari.		
	Type of path	Path name		
	Path of the include file	fujitsuEnterprisePostgresClientInstallDir\include		
	Path of the library	fujitsuEnterprisePostgresClientInstallDit\lib		

W

Table 5.4 C Library (libpg library)

raisie ein e maraily (	1.0   0   1.1.0   1.1.1   1.1.1
Type of library	Library name
Library for links	libpq.lib
Dynamic library	libpq.dll

# 5.4.2 Creating Applications while in Database Multiplexing Mode

This section explains points to consider when creating applications while in database multiplexing mode.



- Refer to the Cluster Operation Guide (Database Multiplexing) for information on database multiplexing mode.

......

- Refer to "Application Development" in the Cluster Operation Guide (PRIMECLUSTER) for points to consider when creating applications using the failover feature integrated with the cluster software.

# 5.4.2.1 Errors when an Application Connection Switch Occurs and Corresponding Actions

If an application connection switch occurs while in database multiplexing mode, explicitly close the connection and then reestablish the connection or reexecute the application.

The table below shows errors that may occur during a switch, and the corresponding action to take.

State		Error information	Action	
Server failure or FUJITSU Enterprise Postgres system failure	Failure occurs during access	PGRES_FATAL_ERRO R(*1) 57P01(*2) NULL(*2)	After the switch is complete, reestablish the connection, or reexecute the	
	Accessed during system failure	CONNECTION_BAD(* 3)	application.	
Switch to the standby server	Switched during access	PGRES_FATAL_ERRO R(*1) 57P01(*2) NULL(*2)		
	Accessed during switch	CONNECTION_BAD(* 3)		

<sup>\*1:</sup> Return value of PQresultStatus().

<sup>\*2:</sup> Return value of PQresultErrorField() PG\_DIAG\_SQLSTATE.

<sup>\*3:</sup> Return value of PQstatus().

# Chapter 6 Embedded SQL in C

This chapter describes application development using embedded SQL in C.

# **6.1 Development Environment**

Install FUJITSU Enterprise Postgres Client Package for the architecture to be developed and executed.



See

Refer to Installation and Setup Guide for Client for information on the C compiler required for C application development.



C++ is not supported. Create a library by implementing embedded SQL in C, and call it from C++.

## 6.2 Setup

### 6.2.1 Environment Settings

When using embedded SQL in C, the same environment settings as when using the C library (libpq) are required.

Refer to "5.2.1 Environment Settings" in "C Library (libpq)" for information on the environment settings for the library for C.

Additionally, set the following path for the precompiler ecpg in the PATH environment variable:



Linux

fujitsuEnterprisePostgresClientInstallDir/bin



Windows(R)

fujitsuEnterprisePostgresClientInstallDir\bin

# 6.2.2 Message Language and Encoding System Used by Applications Settings

The message language and the encoding System Settings Used by Applications settings are the same as when using the library for C.

 $However, in embedded \, SQL, the \, PQ setClient Encoding \, function \, cannot \, be \, used \, in \, the \, encoding \, system \, settings. \, In \, embedded \, SQL, \, use \, the \, SET \, command \, to \, specify \, the \, encoding \, system \, in \, client\_encoding.$ 

Refer to "5.2.2 Message Language and Encoding System Used by Applications Settings" in "C Library (libpq)" for information on the settings for the library for C.

# 6.2.3 Settings for Encrypting Communication Data

When encrypting the communication data, the same environment settings as when using the C library (libpq) are required.

Refer to "5.2.3 Settings for Encrypting Communication Data" in "C Library (libpq)" for information on the environment settings for the C library.

# Point

- It is recommended to use a connection service file to specify connection destinations. In the connection service file, a name (service name) is defined as a set, comprising information such as connection destination information and various types of tuning information set for connections. By using the service name defined in the connection service file when connecting to databases, it is no longer necessary to modify applications when the connection information changes. Refer to "The Connection Service File" in "Client Interfaces" in the PostgreSQL Documentation for information.
- If using a connection service file, perform either of the procedures below:

6.3 Connecting with the Database

- Set the service name as a string literal or host variable, as follows: tcp:postgresql://?service=my\_service
- Set the service name in the environment variable PGSERVICE, and use CONNECT TO DEFAULT

Use the CONNECT statement shown below to create a connection to the database server.

#### **Format**

EXEC SQL CONNECT TO target [AS connection-name] [USER user-name];

#### target

Write in one of the following formats:

- dbname@host:port
- tcp:postgresql://host:port/dbname[?options]
- unix:postgresql://host[:port][/dbname][?options] (Definition method when using the UNIX domain socket)
- SQL string literal containing one of the above formats
- Reference to a character variable containing one of the above formats
- DEFAULT

#### user-name

Write in one of the following formats:

- username
- username/password
- username IDENTIFIED BY password
- username USING password

#### Description of the arguments

Argument	Description	
dbname	Specify the database name.	
host	Specify the host name for the connection destination.	
port	Specify the port number for the database server.	
	The default is "27500".	
connection-name	Specify connection names to identify connections when multiple connections are to be processed within a single program.	

Argument	Description	
username	Specify the user that will connect with the database.	
	If this is omitted, the name used will be that of the user on the operating system that is executing the application.	
password	Specify a password when authentication is required.	
options	Specify the following parameter when specifying a time for timeout. Connect parameters with & when specifying more than one. The following shows the values specified for each parameter.	
	- connect_timeout	
	Specify the timeout for connections.	
	Specify a value between 0 and 2147483647 (in seconds). There is no limit set if you set 0 or an invalid value. If "1" is specified, the behavior will be the same as when "2" was specified. An error occurs when a connection cannot be established within the specified time.	
	- keepalives	
	This enables keepalive.	
	Keepalive is disabled if 0 is specified. Keepalive is enabling when any other value is specified. The default is keepalive enabled. Keepalive causes an error to occur when it is determined that the connection with the database is disabled.	
	- keepalives_idle	
	Specify the time until the system starts sending keepalive messages when communication with the database is not being performed.	
	- Linux	
	Specify a value between 1 and 32767 (in seconds). The default value of the system is used if this is not specified.	
	- Windows(R)	
	Specify a value between 1 and 2147483647 (in seconds). 7200 will be set as default if a value outside this range is specified or if nothing is specified.	
	- keepalives_interval	
	Specify the interval between resends when there is no response to keepalive messages.	
	- Linux	
	Specify a value between 1 and 32767 (in seconds). The default value of the system is used if this is not specified.	
	- Windows(R)	
	Specify a value between 1 and 2147483647 (in seconds). 1 will be set as default if a value outside this range is specified or if nothing is specified.	
	- keepalives_count	
	Specify the number of resends for keepalive messages.	
	- Linux	
	Specify a value between 1 and 127. The default value of the system is used if this is not specified.	
	- Windows(R)	
	The system default value is used irrespective of what is specified for this parameter.	
	- tcp_user_timeout	

Argument	Description
	After establishing the connection, when sending from the client to the server, if the TCP resend process operates, specify the time until it is considered to be disconnected.
	- Linux
	Specify a value between 0 and 2147483647 (in millseconds). The default value of the system is used if 0. 0 will be set as default if nothing is specified.
	- Windows(R)
	Cannot be specified.





If a value other than 0 is specified for the tcp\_user\_timeout parameter, the waiting time set by the tcp\_keepalives\_idle parameter and tcp\_keepalives\_interval parameter will be invalid and the waiting time specified by the tcp\_user\_timeout parameter will be used.

#### Code examples for applications

EXEC SQL CONNECT TO tcp:postgresql://sv1:27500/mydb?
connect\_timeout=20&keepalives\_idle=20&keepalives\_interval=5&keepalives\_count=2&keepalives=
1 USER myuser/myuser01;

# 6.4 Application Development

Refer to "ECPG - Embedded SQL in C" in "Client Interfaces" in the PostgreSQL Documentation for information on developing applications.

However, when using embedded SQL in C, there are the following differences to the embedded SQL (ECPG) in PostgreSQL C

# 6.4.1 Support for National Character Data Types

This section describes how to use the national character data types using the SQL embedded C preprocessor.

The following explains the C language variable types corresponding to the NCHAR type:

Specify the number of characters specified for the NCHAR type multiple by 4, plus 1 for the length of the host variable.

Data Type	Host variable type
NATIONAL CHARACTER(n)	NCHAR variable name [nx4+1]
NATIONAL CHARACTER VARYING(n)	NVARCHAR variable name [nx4+1]



See

Refer to "Handling Character Strings" in "Client Interfaces" in the PostgreSQL documentation for information on using character string types.

# **6.4.2 Compiling Applications**

Append the extension "pgc" to the name of the source file for the embedded SQL in C.

When the pgc file is precompiled using the ecpg command, C source files will be created, so use the C compiler for the compile.

#### Precompiling example

ecpg testproc.pgc

If an optimizer hint block comment is specified for the SQL statement, specify the following option in the ecpg command:

#### --enable-hint

Enables the optimizer hint block comment (hereafter, referred to as the "hint clause"). If this option is not specified, the hint clause will be removed as a result of the ecpg precompile and be disabled.

The SQL statements that can be specified in the hint clause are SELECT, INSERT, UPDATE, and DELETE.

The locations in which the hint clause can be specified are immediately after one of the SELECT, INSERT, UPDATE, DELETE, or WITH keywords. A syntax error will occur if any other location is specified.

#### Example of specifying the hint clause

EXEC SQL SELECT /\*+ IndexScan(prod ix01) \*/ name\_id INTO :name\_id FROM prod WHERE id
= 1;

Refer to "11.1.1 Optimizer Hints" for information on optimizer hints.



Take the following points into account when using embedded SQL source files:

- Multibyte codes expressed in SJIS or UTF-16 cannot be included in statements or host variable declarations specified in EXEC SQL.
- Do not use UTF-8 with a byte order mark (BOM), because an error may occur during compilation if the BOM character is incorrectly recognized as the source code.
- Multibyte characters cannot be used in host variable names.
- It is not possible to use a TYPE name that contains multibyte characters, even though it can be defined.

Specify the following paths when compiling a C application output with precompiling.

Refer to your compiler documentation for information on how to specify the path.

Linux

Table 6.1 Include file and library path

Type of Path	Path name
Path of the include file	fujitsuEnterprisePostgresClientInstallDin/include
Path of the library	fujitsuEnterprisePostgresClientInstallDin'lib

Table 6.2 C Library

Type of library	Library name	Note
Dynamic library	libecpg.so	
	libpgtypes.so	When using the pgtypes library
Static library	libecpg.a	
	libpgtypes.a	When using the pgtypes library

W

#### Windows(R)

If the include file and the library path have been set in the environment variable, there is no need to specify the paths shown below for the compile.

W

Table 6.3 Include file and library path

Type of Path	Path name	
Path of the include file	fujitsuEnterprisePostgresClientInstallDir include	
Path of the library	fujitsuEnterprisePostgresClientInstallDir lib	

W

Table 6.4 C Library

Type of library	Library name	Note	
Library for links	libecpg.lib		
	libpgtypes.lib	When using the pgtypes library	
Dynamic library	libecpg.dll		
	libpgtypes.dll	When using the pgtypes library	



- The libecpg library in Windows(R) is created by "release" and "multithreaded" options. When using the ECPGdebug function included in this library, compile using the "release" and "multithreaded" flags in all programs that use this library. When you do this, use the "dynamic" flag if you are using libecpg.dll, and use the "static" flag if you are using libecpg.lib.

Refer to "Library Functions" in "Client Interfaces" in the PostgreSQL Documentation for information on the ECPGdebug function.

- The cl command expects input to be a program that uses one of the following code pages, so convert the program to these code pages and then compile and link it (refer to the Microsoft documentation for details).
  - ANSI console code pages (example: Shift-JIS for Japanese)
  - UTF-16 little-endian with or without BOM (Byte Order Mark)
  - UTF-16 big-endian with or without BOM
  - UTF-8 with BOM

The cl command converts strings in a program to an ANSI console code page before generating a module, so the data sent to and received from the database server becomes an ANSI console code page. Therefore, set the coding system corresponding to the ANSI console code page as the coding system of the client.

Refer to "Character Set Support" in "Server Administration" in the PostgreSQL Documentation for information on how to set the client encoding system.

(Example: To use environment variables in Japanese, set SJIS in PGCLIENTENCODING.)

#### 6.4.3 Bulk INSERT

This section describes the bulk INSERT.

#### **Synopsis**

```
EXEC SQL [ AT conn ] [ FOR { numOfRows | ARRAY_SIZE } ]
INSERT INTO tableName [ ( colName [, ...] ) ]
{ VALUES ( { expr | DEFAULT } [, ...] ) [, ...] | query }
```

```
[ RETURNING * | outputExpr [ [ AS ] outputName ] [, ...]
INTO outputHostVar [ [ INDICATOR ] indicatorVar ] [, ...] ];
```

#### Description

Bulk INSERT is a feature that inserts multiple rows of data in bulk.

By specifying the array host variable that stored the data in the VALUES clause of the INSERT statement, the data for each element in the array can be inserted in bulk. This feature is used by specifying the insertion count in the FOR clause immediately before the INSERT statement.

#### **FOR Clause**

Specify the insertion count using numOfRows or ARRAY\_SIZE in the FOR clause. The FOR clause can be specified only in the INSERT statement, not in other update statements.

#### numOfRows and ARRAY\_SIZE

Insertion processing will be executed only for the specified count. However, if the count is 1, it will be assumed that the FOR clause was omitted when the application is executed. In this case, proceed according to the INSERT specification in the PostgreSQL Documentation.

Specify the FOR clause as an integer host variable or as a literal.

Specify ARRAY\_SIZE to insert all elements of the array in the table. When specifying ARRAY\_SIZE, specify at least one array in *expr*.

If two or more arrays were specified in expr, it will be assumed that ARRAY\_SIZE is the minimum number of elements in the array.

numOfRows or ARRAY\_SIZE must exceed the minimum number of elements in all arrays specified in expr, outputHostVar, and indicatorVal.

The following example shows how to specify the FOR clause.

#### expr

Specify the value to be inserted in the table. Array host variables, host variable literals, strings, and pointer variables can be specified. Structure type arrays and pointer variable arrays cannot be specified.

Do not use pointer variables and ARRAY\_SIZE at the same time. The reason for this is that the number of elements in the area represented by the pointer variable cannot be determined.

#### query

A query (SELECT statement) that supplies the rows to be inserted. The number of rows returned by *query* must be 1. If two or more rows are returned, an error will occur. This cannot be used at the same time as ARRAY\_SIZE.

#### output Host Var, indicator Val

These must be array host variables or pointer variables.

#### **Error Messages**

Given below are the error messages that are output when bulk INSERT functionality is not used correctly.

#### Message

#### invalid statement name "FOR value should be positive integer"

#### Cause

The value given for *numOfRows* is less than or equal to 0.

#### Solution

Specify a value that is more than or equal to 1 for numOfRows.

#### Message

#### invalid statement name "Host array variable is needed when using FOR ARRAY\_SIZE"

#### Cause

A host array is not specified in the values clause when using the ARRAY\_SIZE keyword.

#### Solution

At least one host array variable should be included in the values clause

#### Message

#### **SELECT...INTO** returns too many rows

#### Cause

The number of rows returned by the 'SELECT ... INTO' query in the INSERT statement is more than one.

#### Solution

When the value of *numOfRows* is more than one, the maximum number of rows that can be returned by the 'SELECT ... INTO' query in the INSERT statement is one.

#### Limitations

The limitations when using bulk INSERT are given below.

- Array of structures should not be used as an input in the 'VALUES' clause. Attempted use will result in junk data being inserted into the table.
- Array of pointers should not be used as an input in the 'VALUES' clause. Attempted use will result in junk data being inserted into the table.
- ECPG supports the use of 'WITH' clause in single INSERT statements. 'WITH' clause cannot be used in bulk INSERT statements.
- ECPG does not calculate the size of the pointer variable. So when a pointer variable is used that includes multiple elements, *numOfRows* should be less than or equal to the number of elements in the pointer. Otherwise, junk data will be inserted into the table.
- If an error occurs, all bulk INSERT actions will be rolled back, therefore, no rows are inserted. However, if the RETURNING clause was used, and the error occurred while obtaining the rows after the insertion was successful, the insertion processing will not be rolled back.

#### **Samples**

Given below are some sample usages of the bulk INSERT functionality.

#### Basic Bulk INSERT

```
int in_f1[4] = {1,2,3,4};
...
EXEC SQL FOR 3 INSERT INTO target (f1) VALUES (:in_f1);
```

The number of rows to insert indicated by the FOR clause is 3, so the data in the first 3 elements of the host array variable are inserted into the table. The contents of the target table will be:

```
f1
----
1
2
3
(3 rows)
```

Also a host integer variable can be used to indicate the number of rows that will be inserted in FOR clause, which will produce the same result as above:

```
int num = 3;
int in_f1[4] = {1,2,3,4};
...
EXEC SQL FOR :num INSERT INTO target (f1) VALUES (:in_f1);
```

#### Inserting constant values

Constant values can also be bulk INSERTed into the table as follows:

```
EXEC SQL FOR 3 INSERT INTO target (f1,f2) VALUES (DEFAULT, 'hello');
```

Assuming the 'DEFAULT' value for the 'f1' column is '0', the contents of the target table will be:

#### Using ARRAY\_SIZE

'FOR ARRAY\_SIZE' can be used to insert the entire contents of a host array variable, without explicitly specifying the size, into the table.

```
int in_f1[4] = {1,2,3,4};
...
EXEC SQL FOR ARRAY_SIZE INSERT INTO target (f1) VALUES (:in_f1);
```

In the above example, four rows are inserted into the table.



If there are multiple host array variables specified as input values, then the number of rows inserted is same as the smallest array size. The example given below demonstrates this usage.

```
int in_f1[4] = {1,2,3,4};
char in_f3[3][10] = {"one", "two", "three"};
...
EXEC SQL FOR ARRAY_SIZE INSERT INTO target (f1,f3) VALUES (:in_f1,:in_f3);
```

In the above example, the array sizes are 3 and 4. Given that the smallest array size is 3, only three rows are inserted into the table. The table contents are given below.

```
3 | three
(3 rows)
```

...........

Using Pointers as Input

Pointers that contain multiple elements can be used in bulk INSERT.

```
int *in_pf1 = NULL;
in_pf1 = (int*)malloc(4*sizeof(int));
in_pf1[0]=1;
in_pf1[1]=2;
in_pf1[2]=3;
in_pf1[3]=4;
...
EXEC SQL FOR 4 INSERT INTO target (f1) values (:in_pf1);
```

The above example will insert four rows into the target table.

#### Using SELECT query

When using bulk INSERT, the input values can be got from the results of a SELECT statement. For example,

```
EXEC SQL FOR 4 INSERT INTO target(f1) SELECT age FROM source WHERE name LIKE 'foo';
```

Assuming that the 'SELECT' query returns one row, the same row will be inserted into the target table four times.



If the 'SELECT' query returns more than one row, the INSERT statement will throw an error.

```
EXEC SQL FOR 1 INSERT INTO target(f1) SELECT age FROM source;
```

In the above example, all the rows returned by the 'SELECT' statement will be inserted into the table. In this context '1' has the meaning of 'returned row equivalent'.

#### Using RETURNING clause

Bulk INSERT supports the same RETURNING clause syntax as normal INSERT. An example is given below.

```
int out_f1[4];
int in_f1[4] = {1,2,3,4};
...
EXEC SQL FOR 3 INSERT INTO target (f1) VALUES (:in_f1) RETURNING f1 INTO :out_f1;
```

After the execution of the above INSERT statement, the 'out\_f1' array will have 3 elements with the values of '1','2' and '3'.

#### 6.4.4 DECLARE STATEMENT

This section describes the DECLARE STATEMENT.

#### **Synopsis**

```
EXEC SQL [ AT connName] DECLARE statementName STATEMENT;
```

#### Description

DECLARE STATEMENT is an embedded SQL command that declares an identifier for a prepared statement. The declared identifier can be used as an identifier in a prepared statement for the following SQL commands:

- EXECUTE
- DECLARE
- DESCRIBE
- PREPARE (Embedded SQL commands)
- PREPARE (SQL commands)

You can associate an identifier in a prepared statement with a connection by executing a DECLARE STATEMENT that specifies the connection. If you specify an identifier associated with the connection in a later SQL command, the SQL command is executed using the connection associated with the identifier. The association between the connection and the identifier in the prepared statement is shared throughout the process.

Only one connection can be associated with a prepared statement of the same name. If you make multiple associations across files, subsequent DECLARE STATEMENT are ignored. If you associate the same file more than once, precompiling the file fails.

If you use the identifier associated with a connection in a SQL command, do not use the AT clause. If a connection different from the connection linked to the identifier using the AT clause is selected, a runtime error will occur.

#### **Parameters**

#### connName

A database connection name established by the CONNECT command.

If AT clause is omitted, no association is made between the connection and the identifier. DECLARE STATEMENT is executed, but has no effect.

#### statementName

Specify the identifier of the prepared statement. You can specify either a SQL identifier or a host variable.

#### **Examples**

Dynamic SQL statement

DECLARE STATEMENT is primarily used to execute dynamic SQL statements.

```
EXEC SQL BEGIN DECLARE SECTION;

char dbname[128];

EXEC SQL END DECLARE SECTION;

...

EXEC SQL CONNECT TO postgres AS con1;

EXEC SQL CONNECT TO another_database AS con2;

EXEC SQL AT con1 DECLARE sql_stmt STATEMENT;

EXEC SQL DECLARE cursor_name CURSOR FOR sql_stmt;

EXEC SQL PREPARE sql_stmt FROM "SELECT current_database()";

EXEC SQL OPEN cursor_name;

EXEC SQL FETCH cursor_name INTO :dbname;

EXEC SQL CLOSE cursor_name;
```

In the example above, the connection 'con1' is associated with the prepared statement identifier 'sql\_stmt'. So the current connection is 'con2', but the embedded SQL commands that follow are executed on 'con1'.

#### PREPARE AS statement

The following is an example using a PREPARE statement in the SQL command:

```
EXEC SQL AT db1 DECLARE stmt STATEMENT;
EXEC SQL PREPARE stmt (int) AS
```

```
SELECT * FROM employee WHERE id = $1;
EXEC SQL EXECUTE stmt USING 1;
```

The above SELECT statement is executed on the connection 'db1'.

#### Compatibility

DECLARE STATEMENT is not specified in the SQL standard.

#### See Also

EXECUTE, DECLARE, DESCLIBE, PREPARE (Embedded SQL commands), PREPARE (SQL commands)



#### See

- Refer to "Embedded SQL Commands" in "Client Interfaces" in the PostgreSQL documentation for information on the embedde SQL commands.

......

- Refer to "SQL Commands" in "Reference" in the PostgreSQL documentation for information on the SQL commands.

# 6.4.5 Creating Applications while in Database Multiplexing Mode

This section explains points to consider when creating applications while in database multiplexing mode.



#### See

- Refer to the Cluster Operation Guide (Database Multiplexing) for information on database multiplexing mode.
- Refer to "Application Development" in the Cluster Operation Guide (PRIMECLUSTER) for points to consider when creating applications using the failover feature integrated with the cluster software.

# 6.4.5.1 Errors when an Application Connection Switch Occurs and Corresponding Actions

If an application connection switch occurs while in database multiplexing mode, explicitly close the connection and then reestablish the connection or reexecute the application.

The table below shows errors that may occur during a switch, and the corresponding action to take.

State		Error information (*1)	Action
Server failure or FUJITSU Enterprise Postgres system failure	Failure occurs during access	57P01 57P02 YE000 26000 40001	After the switch is complete, reestablish the connection, or reexecute the application.
	Accessed during node/system failure	08001	
Switch to the standby server	Switched during	57P01	
	access	57P02	
		YE000	
		26000	

State		Error information (*1)	Action
		40001	
	Accessed during switch	08001	

<sup>\*1:</sup> Return value of SQLSTATE.

# **6.4.6 Notes**

### Notes on creating multithreaded applications

In embedded SQL in C, DISCONNECT ALL disconnects all connections within a process, and therefore it is not thread-safe in all operations that use connections. Do not use it in multithreaded applications.

# Chapter 7 Embedded SQL in COBOL

This chapter describes application development using embedded SQL in COBOL.

# 7.1 Development Environment

Install FUJITSU Enterprise Postgres Client Package for the architecture to be developed and executed.



See

Refer to the Installation and Setup Guide for Client for information on the COBOL compiler required for COBOL application development.

## 7.2 Setup

## 7.2.1 Environment Settings

When using embedded SQL in COBOL, the same environment settings as when using the C library (libpq) are required. Refer to "5.2.1 Environment Settings" in "C Library (libpq)" for information on the environment settings for the library for C.

Additionally, set the following path for the precompiler ecobpg in the PATH environment variable:



Linux

fujitsuEnterprisePostgresClientInstallDir/bin



Windows(R)

fujitsuEnterprisePostgresClientInstallDir\bin

# 7.2.2 Message Language and Encoding System Used by Applications

The settings for the message language and the encoding system used by applications should be the same as those required when using the library for C.

However, in embedded SQL, the PQsetClientEncoding function cannot be used in the encoding system settings. In embedded SQL, use the SET command to specify the encoding system in client\_encoding.

Refer to "5.2.2 Message Language and Encoding System Used by Applications Settings" in "C Library (libpq)" for information on the settings for the library for C.

# 7.2.3 Settings for Encrypting Communication Data

When encrypting the communication data, the same environment settings as when using the C library (libpq) are required.

Refer to "5.2.3 Settings for Encrypting Communication Data" in "C Library (libpq)" for information on the environment settings for the C library.

# 7.3 Connecting with the Database

Use the CONNECT statement shown below to create a connection to the database server.

#### **Format**

EXEC SQL CONNECT TO target [AS connection-name] [USER user-name]END-EXEC.

#### target

Write in one of the following formats:

- dbname@host:port
- tcp:postgresql://host:port/dbname[?options]
- unix:postgresql://host[:port][/dbname][?options]
   (Definition method when using the UNIX domain socket)
- SQL string literal containing one of the above formats
- Reference to a character variable containing one of the above formats
- DEFAULT



If target is DEFAULT, the AS clause and USER clause cannot be specified.

#### user-name

Write in one of the following formats:

- username
- username/password
- username IDENTIFIED BY password
- username USING password

#### **Description of the arguments**

Argument	Description
dbname	Specify the database name.
host	Specify the host name for the connection destination.
port	Specify the port number for the database server.
	The default is "27500".
connection-name	Specify connection names to identify connections when multiple connections are to be processed within a single program.
username	Specify the user that will connect with the database.
	If this is omitted, the name used will be that of the user on the operating system that is executing the application.
password	Specify a password when authentication is required.
options	Specify the following parameter when specifying a time for timeout. Connect parameters with & when specifying more than one. The following shows the values specified for each parameter.
	- connect_timeout
	Specify the timeout for connections.
	Specify a value between 0 and 2147483647 (in seconds). There is no limit set if you set 0 or an invalid value. If "1" is specified, the behavior will be the same as when "2" was specified. An error occurs when a connection cannot be established within the specified time.
	- keepalives
	This enables keepalive.

Argument	Description	
	Keepalive is disabled if 0 is specified. Keepalive is enabled when any other value is specified. The default is keepalive enabled. Keepalive causes an error to occur when it is determined that the connection with the database is disabled.	
	- keepalives_idle	
	Specify the time until the system starts sending keepalive messages when communication with the database is not being performed.	
	- Linux	
	Specify a value between 1 and 32767 (in seconds). The default value of the system is used if this is not specified.	
	- Windows(R)	
	Specify a value between 1 and 2147483647 (in seconds). 7200 will be set as default if a value outside this range is specified or if nothing is specified.	
	- keepalives_interval	
	Specify the interval between resends when there is no response to keepalive messages.	
	- Linux	
	Specify a value between 1 and 32767 (in seconds). The default value of the system is used if this is not specified.	
	- Windows(R)	
	Specify a value between 1 and 2147483647 (in seconds). 1 will be set as default if a value outside this range is specified or if nothing is specified.	
	- keepalives_count	
	Specify the number of resends for keepalive messages.	
	- Linux	
	Specify a value between 1 and 127. The default value of the system is used if this is not specified.	
	- Windows(R)	
	The system default value is used irrespective of what is specified for this parameter.	

#### Code examples for applications

EXEC SQL CONNECT TO tcp:postgresql://sv1:27500/mydb? connect\_timeout=20&keepalives\_idle=20&keepalives\_interval=5&keepalives\_count=2&keepalives=1 USER myuser/myuser01 END-EXEC.

# 7.4 Application Development

Refer to "Appendix D ECOBPG - Embedded SQL in COBOL" for information on developing applications.

# 7.4.1 Support for National Character Data Types

This section describes how to use the national character data types using the SQL embedded COBOL preprocessor.

The table below lists the COBOL variable types supporting the national character data types. The number of characters specified for the national character data type must be specified for the length of the host variable.

National character data type	COBOL variable type	
CHARACTER(n)	<i>varName</i> PIC N( <i>n</i> )	

National character data type	COBOL variable type	
NATIONAL CHARACTER(n)		
CHARACTER VARYING(n)	varName PIC N(n) VARYING	
NATIONAL CHARACTER VARYING(n)		

To use COBOL variable types that support national character data types, it is necessary to specify the ECOBPG\_NCHAR environment variable.

ECOBPG\_NCHAR={ UTF16LE | UTF16BE | UTF32LE | UTF32BE | SJIS }

In SQL embedded COBOL, specify the encoding of the COBOL variable types that support national character data types.

- UTF16LE: UTF-16 little-endian

- UTF16BE: UTF-16 big-endian

- UTF32LE: UTF-32 little-endian

- UTF32BE: UTF-32 big-endian

- SJIS: Shift JIS

If this environment variable is omitted, the encoding will be determined according to the encoding system of the client.

- If UTF8 is used: UTF16 (endians will be encoded in accordance with endians of the client system)

- If SJIS is used: SJIS

If encoding is specified for the translation option when compiling with NetCOBOL, the encoding specified for the national character data types should be used for the environment variable ECOBPG\_NCHAR.

The list below shows NetCOBOL translation options and their corresponding environment variable ECOBPG\_NCHAR values.

NetCOBOL translation options	Environment variable ECOBPG_NCHAR	
ENCODE (UTF-8,UTF16,LE)	UTF-16LE	
RCS (UTF-16,LE)	UIT-10LE	
ENCODE (UTF-8,UTF-16,BE)	UTF-16BE	
RCS (UTF-16,BE)	UIF-10BE	
ENCODE (UTF-8,UTF-32,LE)	UTF-32LE	
ENCODE (UTF-8,UTF-32,BE)	UTF-32BE	
ENCODE (SJIS,SJIS)	SJIS	
Not specified	No need to specify	

Also, if the post-compiling encoding for an application differs from the locale of the execution environment, then the client encoding must be used for the application.

The list below shows the values supported for the combinations of application encoding, locale of the execution environment, and client encodings.

Application encoding	Locale used when executing an application	Client encoding
UTF-8	UTF-8	UTF-8
	SJIS	UTF-8
SJIS	UTF-8	SJIS

Application encoding	Locale used when executing an application	Client encoding
	SJIS	SJIS

Refer to "7.2.2 Message Language and Encoding System Used by Applications" for information on how to set client encoding systems.

The following example shows host variable declaration of a national character data type.

```
01 DATA1 PIC N(10).
01 DATA2 PIC N(10) VARYING.
```



- Halfwidth characters should not be used for the national character data type COBOL variable.
- The national character data type column attribute obtained by applications should be the CHAR type.
- Encoding cannot be specified using the ENCODING clause, which is a feature of NetCOBOL.

### 7.4.2 Compiling Applications

Append the extension "pco" to the name of the source file for the embedded SQL in COBOL.

When the pco file is precompiled using the ecobpg command, COBOL source files will be created, so use the COBOL compiler for the compile.

#### Precompiling example

```
ecobpg testproc.pco
```

For COBOL code notation, "fixed" or "variable" format can be specified as an ecobpg command option. If not specified, "fixed" format is used.

Refer to "D.1 Precautions when Using Functions and Operators" and "D.12.1 ecobpg" for information on COBOL code notation and how to specify options for ecobpg.

If an optimizer hint block comment is specified for the SQL statement, specify the following option in the ecobpg command:

#### --enable-hint

Enables the optimizer hint block comment (hereafter, referred to as the "hint clause"). If this option is not specified, the hint clause will be removed as a result of the ecobpg precompile and be disabled.

The SQL statements that can be specified in the hint clause are SELECT, INSERT, UPDATE, and DELETE.

The locations in which the hint clause can be specified are immediately after one of the SELECT, INSERT, UPDATE, DELETE, or WITH keywords. A syntax error will occur if it is specified in any other location.

Example of specifying the hint clause

```
EXEC SQL SELECT /*+ IndexScan(prod ix01) */ name_id
INTO :name_id FROM prod WHERE id = 1 END-EXEC.
```

Refer to "11.1.1 Optimizer Hints" for information on optimizer hints.

If the encoding used for embedded SQL source files differs from that of the locale when precompiling was executed, set the encoding for the embedded SQL source files by specifying the following option for ecobpg.

#### -E-encode

```
Specify "UTF8", "SJIS", or "EUC_JP".
```

If this option is omitted, the encoding is processed based on the locale.

#### Path of the library file

The ecobpg command defines a group item "sqlca\_t" to handle errors, which is defined in the library file stored in the following path:

Linux

Library file name	The storage destination of library file
SQLCA-COBOL.cob	fujitsuEnterprisePostgresClientInstallDin'include

Windows(R)

Library file name	The storage destination of library file
SQLCA-COBOL.cob	fujitsuEnterprisePostgresClientInstallDir\include

When the ecobpg command generates a COBOL file, it inserts a COPY statement with no options to copy the library file. Therefore, specify the path of the storage destination of library file when compiling. How to specify the path must conform to your compiler's specifications.

There is also a library file with the same contents without the extension"cob".



Refer "D.7.2 sqlca" for information on the sqlca\_t.

#### Libraries to use

The applications generated by ecobpg connect to PostgreSQL through the ECPG library. The ECPG library internally loads the libpq library.

#### Path of library

Refer to "6.4.2 Compiling Applications" for information on the location and name of the ECPG library. And Refer to "Chapter 5 C Library (libpq)" for information on the location and name of the libpq library.

The COBOL compiler provides the how to link various libraries, so be sure to specify the path and libraries according to the specifications of your compiler.

#### **Entry information of subprogram**

If you use the ECPG library with a dynamic program structure, copy the entry information stored below. For details, follow the specifications of your compiler.

Linux

W

fujitsuEnterprisePostgresClientInstallDir/share/cobol\_entry.info

Windows(R)

fujitsuEnterprisePostgresClientInstallDir\share\cobol\_entry.info



The examples of compiling the applications that dynamically links the COBOL language library. Note that "<*x>*" indicates the product version.

L

Linux

- Linux 64-bit application:

```
cobol -M -o testproc -I/opt/fsepv<x>client64/include -L/opt/fsepv<x>client64/lib -
lecpg -lpq testproc.cob
```

W

#### Windows(R)

The examples of compiling on a 64-bit operating system.

- 64-bit application

```
> SET LIB=%ProgramFiles%\Fujitsu\fsepv<x>client64\lib;%LIB%
> cobol -I "%ProgramFiles%\Fujitsu\fsepv<x>client64\include" -M testproc.cob
> link testproc.obj F4AGCIMP.LIB LIBCMT.LIB LIBECPG.LIB LIBPQ.LIB /OUT:testproc.exe
```

- 32-bit application

#### [NetCOBOL V10.5 or earlier]

```
> SET LIB=%ProgramFiles(x86)%\Fujitsu\fsepv<x>client32\lib;%LIB%
> cobol32 -I "%ProgramFiles(x86)%\Fujitsu\fsepv<x>client32\include" -M testproc.cob
> link testproc.obj LIBC.LIB F3BICIMP.LIB LIBECPG.LIB LIBPQ.LIB /OUT:testproc.exe
```

#### [NetCOBOL V11.0 or later]

```
> SET LIB=%ProgramFiles(x86)%\Fujitsu\fsepv<x>client32\lib;%LIB%
> cobol32 -I "%ProgramFiles(x86)%\Fujitsu\fsepv<x>client32\include" -M testproc.cob
> link testproc.obj MSVCRT.LIB F3BICIMP.LIB LIBECPG.LIB LIBPQ.LIB /OUT:testproc.exe
```

......

### 7.4.3 Bulk INSERT

This section describes the bulk INSERT.

#### **Synopsis**

```
EXEC SQL [ AT conn ] [ FOR { numOfRows | ARRAY_SIZE } ]
   INSERT INTO tableName [ ( colName [, ...] ) ]
   { VALUES ( { expr | DEFAULT } [, ...] ) [, ...] | query }
   [ RETURNING * | outputExpr [ [ AS ] outputName ] [, ...]
   INTO outputHostVar [ [ INDICATOR ] indicatorVar ] [, ...] ] END-EXEC
```

#### Description

Bulk INSERT is a feature that inserts multiple rows of data in bulk.

By specifying the array host variable that stored the data in the VALUES clause of the INSERT statement, the data for each element in the array can be inserted in bulk. This feature is used by specifying the insertion count in the FOR clause immediately before the INSERT statement.

#### **FOR Clause**

Specify the insertion count using *numOfRows* or ARRAY\_SIZE in the FOR clause. The FOR clause can be specified only in the INSERT statement, not in other update statements.

#### numOfRows and ARRAY\_SIZE

Insertion processing will be executed only for the specified count. However, if the count is 1, it will be assumed that the FOR clause was omitted when the application is executed. In this case, proceed according to the INSERT specification in the PostgreSQL Documentation.

Specify the FOR clause as an integer host variable or as a literal.

Specify ARRAY\_SIZE to insert all elements of the array in the table. When specifying ARRAY\_SIZE, specify at least one array in *expr*.

If two or more arrays were specified in *expr*, it will be assumed that ARRAY\_SIZE is the minimum number of elements in the array.

numOfRows or ARRAY\_SIZE must exceed the minimum number of elements in all arrays specified in expr, outputHostVar, and indicatorVal.

The following example shows how to specify the FOR clause.

```
01 NUMBER-OF-ROWS PIC S9(9) COMP VALUE 10.
01 GROUP-ITEM.
05 ID1 PIC S9(9) OCCURS 25.
05 NAME PIC X(10) OCCURS 25.
* will process 10 rows
EXEC SQL FOR :NUMBER-OF-ROWS
INSERT INTO prod (name, id) VALUES (:NAME, :ID1) END-EXEC
* will process 25 rows
EXEC SQL FOR ARRAY_SIZE
INSERT INTO prod (name, id) VALUES (:NAME, :ID1) END-EXEC
```

#### expr

Specify the value to be inserted in the table. Array host variables, host variable literals, strings, and pointer variables can be specified. Structure type arrays and pointer variable arrays cannot be specified.

Do not use pointer variables and ARRAY\_SIZE at the same time. The reason for this is that the number of elements in the area represented by the pointer variable cannot be determined.

#### query

A query (SELECT statement) that supplies the rows to be inserted. The number of rows returned by *query* must be 1. If two or more rows are returned, an error will occur. This cannot be used at the same time as ARRAY\_SIZE.

#### outputHostVar and indicatorVal

These must be array host variables or pointer variables.

#### **Error Messages**

The messages below are output if an error occurs when the bulk INSERT is used.

#### Message

#### The value for the FOR clause must be a positive integer.

#### Cause

The value given for *numOfRows* is less than or equal to 0.

#### Solution

Specify a value that is more than or equal to 1 for numOfRows.

#### Message

#### Array host variable is needed when using FOR ARRAY\_SIZE.

#### Cause

An array host variable is not specified in the VALUES clause.

#### Solution

Specify more than one array host variable in the VALUES clause.

#### Message

#### The SELECT..INTO query returned too many rows in row number %d.

#### Cause

The "SELECT ... INTO" query in the INSERT statement returned more than one row.

#### Solution

If *numOfRows* is more than one, the maximum number of rows that can be returned in the "SELECT ... INTO" query in the INSERT statement is one.

#### Limitations

The limitations when using bulk INSERT are given below.

- Array of structures should not be used as an input in the 'VALUES' clause.
- Array of pointers should not be used as an input in the 'VALUES' clause.
- ECOBPG supports the use of 'WITH' clause in single INSERT statements. 'WITH' clause cannot be used in bulk INSERT statements.
- If an error occurs, all bulk INSERT actions will be rolled back, therefore, no rows are inserted. However, if the RETURNING clause was used, and the error occurred while obtaining the rows after the insertion was successful, the insertion processing will not be rolled back.

#### **Samples**

Given below are some sample usages of the bulk INSERT functionality.

#### Basic Bulk INSERT

```
01 GROUP-ITEM.
05 IN-F1 PIC S9(9) OCCURS 4.

MOVE 1 TO IN-F1(1)

MOVE 2 TO IN-F1(2)

MOVE 3 TO IN-F1(3)

MOVE 4 TO IN-F1(4)

...

EXEC SQL FOR 3 INSERT INTO target (f1) VALUES (:IN-F1) END-EXEC
```

The number of rows to insert indicated by the FOR clause is 3, so the data in the first 3 elements of the host array variable are inserted into the table. The contents of the target table will be:

```
f1
----
1
2
3
(3 rows)
```

Also a host integer variable can be used to indicate the number of rows that will be inserted in FOR clause, which will produce the same result as above:

```
01 NUM PIC S9(9) COMP VALUE 3.
01 GROUP-ITEM.
```

```
05 IN-F1 PIC S9(9) OCCURS 4.

MOVE 1 TO IN-F1(1)

MOVE 2 TO IN-F1(2)

MOVE 3 TO IN-F1(3)

MOVE 4 TO IN-F1(4)

...

EXEC SQL FOR :NUM INSERT INTO target (f1) VALUES (:IN-F1) END-EXEC
```

#### Inserting constant values

Constant values can also be bulk INSERTed into the table as follows:

```
EXEC SQL FOR 3 INSERT INTO target (f1,f2) VALUES (DEFAULT, 'hello') END-EXEC
```

Assuming the 'DEFAULT' value for the 'f1' column is '0', the contents of the target table will be:

```
f1 | f2

---+-----

0 | hello

0 | hello

0 | hello

(3 rows)
```

#### Using ARRAY\_SIZE

'FOR ARRAY\_SIZE' can be used to insert the entire contents of a host array variable, without explicitly specifying the size, into the table.

```
01 GROUP-ITEM.
05 IN-F1 PIC S9(9) OCCURS 4.

MOVE 1 TO IN-F1(1)

MOVE 2 TO IN-F1(2)

MOVE 3 TO IN-F1(3)

MOVE 4 TO IN-F1(4)

...

EXEC SQL FOR ARRAY_SIZE INSERT INTO target (f1) VALUES (:IN-F1) END-EXEC
```



If there are multiple host array variables specified as input values, then the number of rows inserted is same as the smallest array size. The example given below demonstrates this usage.

```
01 GROUP-ITEM.
05 IN-F1 PIC S9(9) OCCURS 4.
05 IN-F3 PIC X(10) OCCURS 3.

MOVE 1 TO IN-F1(1)

MOVE 2 TO IN-F1(2)

MOVE 3 TO IN-F1(3)

MOVE 4 TO IN-F1(4)

MOVE "one" TO IN-F3(1)

MOVE "two" TO IN-F3(2)

MOVE "three" TO IN-F3(3)

...

EXEC SQL FOR ARRAY_SIZE INSERT INTO target (f1,f3) VALUES (:IN-F1,:IN-F3) END-EXEC
```

In the above example, the array sizes are 3 and 4. Given that the smallest array size is 3, only three rows are inserted into the table. The table contents are given below.

```
2 | two
3 | three
(3 rows)
```

# Using SELECT query

The result of a SELECT query can be used to insert values.

```
EXEC SQL FOR 4 INSERT INTO target(f1) SELECT age FROM source WHERE name LIKE 'foo' END-EXEC
```

In the example above, assuming that the SELECT query returns one row, the same row will be inserted into the table four times



If "2" or more is specified for the FOR clause, the INSERT statement returns an error when two or more rows of query results are returned.

If "1" is specified for the FOR clause, all rows returned by the SELECT query will be inserted into the table.

```
EXEC SQL FOR 1 INSERT INTO target(f1) SELECT age FROM source END-EXEC
```

In the example above, "1" specified for the FOR clause indicates all returned rows.

# Using RETURNING clause

Bulk INSERT supports the same RETURNING clause syntax as normal INSERT. An example is given below.

```
01 GROUP-ITEM.
05 IN-F1 PIC S9(9) OCCURS 4.
05 OUT-F1 PIC S9(9) OCCURS 4.
MOVE 1 TO IN-F1(1)
MOVE 2 TO IN-F1(2)
MOVE 3 TO IN-F1(3)
MOVE 4 TO IN-F1(4)
...
EXEC SQL FOR 3 INSERT INTO target (f1) VALUES (:IN-F1) RETURNING f1 INTO :OUT-F1 END-EXEC
```

After the execution of the above INSERT statement, the 'out\_f1' array will have 3 elements with the values of '1', '2' and '2'

# 7.4.4 DECLARE STATEMENT

Refer to "6.4.4 DECLARE STATEMENT" in "Embedded SQL in C".

# 7.4.5 Creating Applications while in Database Multiplexing Mode

This section explains points to consider when creating applications while in database multiplexing mode.



- Refer to the Cluster Operation Guide (Database Multiplexing) for information on database multiplexing mode.
- Refer to "Application Development" in the Cluster Operation Guide (PRIMECLUSTER) for points to consider when creating applications using the failover feature integrated with the cluster software.

# 7.4.5.1 Errors when an Application Connection Switch Occurs and Corresponding Actions

If an application connection switch occurs while in database multiplexing mode, explicitly close the connection and then reestablish the connection or reexecute the application.

The table below shows errors that may occur during a switch, and the corresponding action to take.

State		Error information (*1)	Action
Server failure or FUJITSU Enterprise Postgres system failure	Failure occurs during access	57P01 57P02 YE000 26000 40001	After the switch is complete, reestablish the connection, or reexecute the application.
	Accessed during system failure	08001	
Switch to the standby server	Switched during access	57P01 57P02 YE000 26000 40001	
	Accessed during switch	08001	

<sup>\*1:</sup> Return value of SQLSTATE.

# Chapter 8 SQL References

This chapter explains the SQL statement features expanded by FUJITSU Enterprise Postgres.

# 8.1 Expanded Trigger Definition Feature

This section explains the expanded trigger definition feature.

# 8.1.1 CREATE TRIGGER

In addition to features of PostgreSQL, triggers can be created with OR REPLACE option and DO option.

# **Synopsis**

```
CREATE [ OR REPLACE ] [ CONSTRAINT ] TRIGGER name { BEFORE | AFTER | INSTEAD OF } { event [ OR ... ] }

ON table_name

[ FROM referenced_table_name ]

[ NOT DEFERRABLE | [ DEFERRABLE ] [ INITIALLY IMMEDIATE | INITIALLY DEFERRED ] ]

[ REFERENCING { OLD | NEW } TABLE [ AS ] transition_relation_name } [ ... ] ]

[ FOR [ EACH ] { ROW | STATEMENT } ]

[ WHEN ( condition ) ]

{ EXECUTE { FUNCTION | PROCEDURE } function_name ( arguments )

| DO [ LANGUAGE lang_name ] code }
```

# **Description**

Refer to the PostgreSQL Documentation for information about CREATE TRIGGER. This section describes OR REPLACE option and DO option.

A trigger which is created with OR REPLACE option and DO option will be associated with the specified table or view and will execute the specified code by the specified procedural language of DO (unnamed code block) when certain events occur.

#### **Parameters**

# OR REPLACE

If the specified trigger is not defined in the table, it defines a new trigger.

If the specified trigger is already defined in the table, the named trigger replaces existing trigger.

code

When the certain events occur, it executes the code in a specified procedural language. The unnamed code block does not require a prior definition like a function. Syntax is same as procedural language.

# lang\_name

The name of the language that the function is implemented in. Can be SQL, C, internal, or the name of a user-defined procedural language. The default is 'plpgsql'.

plpgsql is supported in CREATE TRIGGER.



- A normal trigger cannot be replaced by a constraint trigger.
- A constraint trigger cannot be replaced by a normal trigger.
- A trigger defined with DO option cannot be replaced by a trigger defined with EXECUTE PROCEDURE option.
- A trigger defined with EXECUTE PROCEDURE option cannot be replaced by a trigger defined with DO option.

#### **Examples**

It executes the code block that is specified by DO before the table is updated. (Example that LANGUAGE is plpgsql)

CREATE TRIGGER check\_update

BEFORE UPDATE ON accounts

FOR EACH ROW

DO \$\$BEGIN RETURN NEW; END;\$\$;



When a trigger created with DO option, a new function is created internally. The name of function is "schema name". "on table name"\_"trigger name"\_TRIGPROC(serial number).

......

# 8.1.2 How to Define Triggers in pgAdmin

The expanded features of the trigger definition can also be used in pgAdmin.



See

Refer to "pgAdmin Help" for information on how to define triggers using pgAdmin.

# Chapter 9 Compatibility with Oracle Databases

This chapter describes the environment settings and functionality offered for features that are compatible with Oracle databases.

# 9.1 Overview

Features compatible with Oracle databases are provided. These features enable you to easily migrate to FUJITSU Enterprise Postgres and reduce the costs of reconfiguring applications.

The table below lists features compatible with Oracle databases.

Table 9.1 Features compatible with Oracle databases

Cat	egory	Feature		Feature	
		Item	Overview		
SQL	Queries	Outer join operator (+)	Operator for outer joins		
		DUAL table	Table provided by the system		
	Functions	DECODE	Compares values, and if they match, returns a corresponding value		
		SUBSTR	Extracts part of a string using characters to specify position and length		
		NVL	Returns a substitute value when a value is NULL		
Pac	ckage	DBMS_OUTPUT	Sends messages to clients		
		UTL_FILE (*1)	Enables text file operations		
		DBMS_SQL	Enables dynamic SQL execution		

<sup>\*1:</sup> It is not available on Windows (R).



#### See

In addition to the above, refer to the file below for information on the Orafce function.

- L
- Linux:

fujitsuEnterprisePostgresInstallDir/share/doc/extension/README.asciidoc



- Windows(R):
  - fujitsuEnterprisePostgresInstallDir\doc\extension\README.asciidoc

# 9.2 Precautions when Using the Features Compatible with Oracle Databases

This section provides notes on using the features compatible with oracle databases.

# 9.2.1 Notes on SUBSTR

SUBSTR is implemented in FUJITSU Enterprise Postgres and Oracle databases using different external specifications.

For this reason, when using SUBSTR, define which specification is to take precedence. In the default configuration of FUJITSU Enterprise Postgres, the specifications of FUJITSU Enterprise Postgres take precedence.

When using the SUBSTR function compatible with Oracle databases, set "oracle" and "pg\_catalog" in the "search\_path" parameter of postgresql.conf. You must specify "oracle" before "pg\_catalog" when doing this.

```
search_path = '"$user", public, oracle, pg_catalog'
```



- The search\_path parameter specifies the order in which schemas are searched. The SUBSTR function in Oracle databases is defined in the oracle schema.
- Refer to "Statement Behavior" in "Client Connection Defaults" in "Server Administration" in the PostgreSQL Documentation for information on search\_path.

# 9.2.2 Notes when Integrating with the Interface for Application Development

The SQL noted in "Table 9.1 Features compatible with Oracle databases" can be used in the interface for application development. However, outer join operators cannot be used when integrated with Visual Studio.

When integrated with Visual Studio or using the features compatible with Oracle databases from Fujitsu Npgsql .NET Data Provider, select one of the actions below for the SearchPath parameter, which is one of the pieces of information needed to connect to databases specified for individual connections.

- Do not specify the SearchPath parameter itself, or
- Specify both "public" and the schema name in the SQL statement.

Note that both "public" and the schema name in the SQL statement must be specified as the SearchPath parameter before "oracle" and "pg\_catalog" when using the Oracle database-compatible feature SUBSTR.

# 9.3 Queries

The following queries are supported:

- Outer Join Operator (+)
- DUAL Table

# 9.3.1 Outer Join Operator (+)

In the WHERE clause conditional expression, by adding the plus sign (+), which is the outer join operator, to the column of the table you want to add as a table join, it is possible to achieve an outer join that is the same as a joined table (OUTER JOIN).

## Syntax

#### SELECT statement

```
SELECT ... [WHERE [NOT] joinCond ...] ...
SELECT ... [WHERE srchCond ]... ] ...

Join condition
{ colSpec(+) = colSpec | colSpec = colSpec(+) }
```



Here we are dealing only with the WHERE clause of the SELECT statement. Refer to "SQL Commands" in "Reference" in the PostgreSQL Documentation for information on the overall syntax of the SELECT statement.

#### General rules

#### WHERE clause

- The WHERE clause specifies search condition or join conditions for the tables that are derived.

- Search conditions are any expressions that return BOOLEAN types as the results of evaluation. Any rows that do not
  meet these conditions are excluded from the output. When the values of the actual rows are assigned to variables and
  if the expression returns TRUE, those rows are considered to have met the conditions.
- Join conditions are comparison conditions that specify outer join operators. Join conditions in a WHERE clause return a table that includes all the rows that meet the join conditions, including rows that do not meet all the join conditions.
- Join conditions take precedence over search conditions. For this reason, all rows returned by the join conditions are subject to the search conditions.
- The following rules and restrictions apply to queries that use outer join operators. It is therefore recommended to use FROM clause joined tables (OUTER JOIN) rather than outer join operators:
  - Outer join operators can only be specified in the WHERE clause.
  - Outer join operators can only be specified for base tables or views.
  - To perform outer joins using multiple join conditions, it is necessary to specify outer join operators for all join conditions.
  - When combining join conditions with constants, specify outer join operators in the corresponding column specification. When not specified, they will be treated as search conditions.
  - The results column of the outer join of table t1 is not returned if table t1 is joined with table t2 by specifying an outer join operator in the column of t1, then table t1 is joined with table t3 by using search conditions.
  - It is not possible to specify columns in the same table as the left/right column specification of a join condition.
  - It is not possible to specify an expression other than a column specification for outer join operators, but they may be specified for the columns that compose the expression.

There are the following limitations on the functionality of outer join operators when compared with joined tables (OUTER JOIN). To use functionality that is not available with outer join operators, use joined tables (OUTER JOIN).

Table 9.2 Range of functionality with outer join operators

Functionality available with joined tables (OUTER JOIN)	Outer join operator
Outer joins of two tables	Y
Outer joins of three or more tables	Y (*1)
Used together with joined tables within the same query	N
Use of the OR logical operator to a join condition	N
Use of an IN predicate to a join condition	N
Use of subqueries to a join condition	N

Y: Available

N: Not available

\*1: The outer joins by outer join operators can return outer join results only for one other table. For this reason, to combine outer joins of table t1 and table t2 or table t2 and table t3, it is not possible to specify outer join operators simultaneously for table t2.



Table configuration

t1

col1	col2	col3

1001	AAAAA	1000
1002	BBBBB	2000
1003	CCCCC	3000

t2

col1	col2
1001	aaaaa
1002	bbbbb
1004	ddddd

Example 1: Return all rows in table t2, including those that do not exist in table t1.

This is the same syntax as the joined table (OUTER JOIN) of the FROM clause shown next.

```
SELECT *
FROM t1 RIGHT OUTER JOIN t2
ON t1.col1 = t2.col1;
```

Example 2: In the following example, the results are filtered to records above 2000 in t1.col3 by search conditions, and the records are those in table t2 that include ones that do not exist in table t1. After filtering with the join conditions, there is further filtering with the search conditions, so there will only be one record returned.

This is the same syntax as the joined table (OUTER JOIN) of the FROM clause shown next.

```
SELECT *

FROM t1 RIGHT OUTER JOIN t2

ON t1.col1 = t2.col1

WHERE t1.col3 >= 2000;
```

# 9.3.2 DUAL Table

DUAL table is a virtual table provided by the system. Use when executing SQL where access to a base table is not required, such as when performing tests to get result expressions such as functions and operators.



In the following example, the current system date is returned.

```
SELECT CURRENT_DATE "date" FROM DUAL;
date
------
2013-05-14
(1 row)
```

# 9.4 SQL Function Reference

The following SQL functions are supported:

- DECODE
- SUBSTR
- NVL

# **9.4.1 DECODE**

## Description

Compares values and if they match, returns a corresponding value.

#### Syntax

```
DECODE(expr, srch, result [, srch, result ]... [, default ])
```

#### General rules

- DECODE compares values of the value expression to be converted and the search values one by one. If the values match, a corresponding result value is returned. If no values match, the default value is returned if it has been specified. A NULL value is returned if a default value has not been specified.
- If the same search value is specified more than once, then the result value returned is the one listed for the first occurrence of the search value.
- The following data types can be used in result values and in the default value:
  - CHAR
  - VARCHAR
  - NCHAR
  - NCHAR VARYING
  - TEXT
  - INTEGER
  - BIGINT
  - NUMERIC
  - DATE
  - TIME WITHOUT TIME ZONE
  - TIMESTAMP WITHOUT TIME ZONE
  - TIMESTAMP WITH TIME ZONE
- The same data type must be specified for the values to be converted and the search values. However, note that different data types may also be specified if a literal is specified in the search value, and the value expression to be converted

contains data types that can be converted. When specifying literals, refer to "Table A.1 Data type combinations that contain literals and can be converted implicitly" in "A.3 Implicit Data Type Conversions" for information on the data types that can be specified.

- If the result values and default value are all literals, the data types for these values will be as shown below:
  - If all values are string literals, all will become character types.
  - If there is one or more numeric literal, all will become numeric types.
  - If there is one or more literal cast to the datetime/time types, all will become datetime/time types.
- If the result values and default value contain a mixture of literals and non-literals, the literals will be converted to the data types of the non-literals. When specifying literals, refer to "Table A.1 Data type combinations that contain literals and can be converted implicitly" in "A.3 Implicit Data Type Conversions" for information on the data types that can be converted.
- The same data type must be specified for all result values and for the default value. However, different data types can be specified if the data type of any of the result values or default value can be converted these data types are listed below:

Table 9.3 Data type combinations that can be converted by DECODE (summary)

		Other res	ult values	or default value
		Numeric type	Characte r type	Date/time type
Result value	Numeric type	Y	N	N
(any)	Character type	N	Y	N
	Date/time type	N	N	S (*1)

Y: Can be converted

S: Some data types can be converted

N: Cannot be converted

\*1: The data types that can be converted for date/time types are listed below:

Table 9.4 Result value and default value date/time data types that can be converted by DECODE

		Other result values or default value			value
		DATE	TIME WITHOUT TIME ZONE	TIMESTAMP WITHOUT TIME ZONE	TIMESTAM P WITH TIME ZONE
Result	DATE	Y	N	Y	Y
value (any)	TIME WITHOUT TIME ZONE	N	Y	N	N
	TIMESTAMP WITHOUT TIME ZONE	Y	N	Y	Y
	TIMESTAMP WITH TIME ZONE	Y	N	Y	Y

Y: Can be converted

N: Cannot be converted

- The data type of the return value will be the data type within the result or default value that is longest and has the highest precision.



In the following example, the value of col3 in table t1 is compared and converted to a different value. If the col3 value matches search value 1, the result value returned is "one". If the col3 value does not match any of search values 1, 2, or 3, the default value "other number" is returned.

# **9.4.2 SUBSTR**

#### Description

Extracts part of a string using characters to specify position and length.

### Syntax

```
SUBSTR(str, startPos [, len ])
```

#### General rules

- SUBSTR extracts and returns a substring of string str, beginning at position startPos, for number of characters len.
- When startPos is positive, it will be the number of characters from the beginning of the string.
- When startPos is 0, it will be treated as 1.
- When startPos is negative, it will be the number of characters from the end of the string.
- When len is not specified, all characters to the end of the string are returned. NULL is returned when len is less than 1.
- For *startPos* and *len*, specify a SMALLINT or INTEGER type. When specifying literals, refer to "Table A.1 Data type combinations that contain literals and can be converted implicitly" in "A.3 Implicit Data Type Conversions" for information on the data types that can be specified.
- The data type of the return value is TEXT.



- There are two types of SUBSTR. One that behaves as described above, and one that behaves the same as SUBSTRING. The search\_path parameter must be modified for it to behave the same as the specification described above.

- It is recommended to set search\_path in postgresql.conf. In this case, it will be effective for each instance. Refer to "9.2.1 Notes on SUBSTR" for information on how to configure postgresql.conf.
- The configuration of search\_path can be done at the user level or at the database level. Setting examples are shown below.
  - Example of setting at the user level

This can be set by executing an SQL command. In this example, user1 is used as the username.

```
ALTER USER user1 SET search_path = "$user",public,oracle,pg_catalog;
```

- Example of setting at the database level

This can be set by executing an SQL command. In this example, db1 will be used as the database name.

```
ALTER DATABASE db1 SET search_path = "$user",public,oracle,pg_catalog;
```

You must specify "oracle" before "pg\_catalog".

- If the change has not been implemented, SUBSTR is the same as SUBSTRING.



#### See

Refer to "SQL Commands" in "Reference" in the PostgreSQL Documentation for information on ALTER USER and ALTER DATABASE.



# Information

The general rules for SUBSTRING are as follows:

- The start position will be from the beginning of the string, whether positive, 0, or negative.
- When len is not specified, all characters to the end of the string are returned.
- An empty string is returned if no string is extracted or *len* is less than 1.



#### Sec

Refer to "String Functions and Operators" under "The SQL Language" in the PostgreSQL Documentation for information on SUBSTRING.



# Example

In the following example, part of the string "ABCDEFG" is extracted:

```
SELECT SUBSTR('ABCDEFG',3,4) "Substring" FROM DUAL;

Substring
------
CDEF
(1 row)

SELECT SUBSTR('ABCDEFG',-5,4) "Substring" FROM DUAL;

Substring
------
(1 row)
```

# 9.4.3 NVL

## Description

Returns a substitute value when a value is NULL.

# Syntax

```
NVL(expr1, expr2)
```

# General rules

- NVL returns a substitute value when the specified value is NULL. When *expr1* is NULL, *expr2* is returned. When *expr1* is not NULL, *expr1* is returned.

- Specify the same data types for *expr1* and *expr2*. However, if a constant is specified in *expr2*, and the data type can also be converted by *expr1*, different data types can be specified. When this happens, the conversion by *expr2* is done to suit the data type in *expr1*, so the value of *expr2* returned when *expr1* is a NULL value will be the value converted in the data type of *expr1*.
- When specifying literals, refer to "Table A.1 Data type combinations that contain literals and can be converted implicitly" in "A.3 Implicit Data Type Conversions" for information on the data types that can be converted.



In the following example, "IS NULL" is returned if the value of col1 in table t1 is a NULL value.

```
SELECT col2, NVL(col1,'IS NULL') "nvl" FROM t1;

col2 | nvl

----+----

aaa | IS NULL

(1 row)
```

# 9.5 Package Reference

A "package" is a group of features, brought together by schemas, that have a single functionality, and are used by calling from PL/pgSQL.

The following packages are supported:

- DBMS\_OUTPUT
- UTL\_FILE
- DBMS\_SQL

To call the different functionalities from PL/pgSQL, use the PERFORM statement or SELECT statement, using the package name to qualify the name of the functionality. Refer to the explanations for each of the package functionalities for information on the format for calling.

# 9.5.1 DBMS OUTPUT

#### Overview

Sends messages to clients such as psql from PL/pgSQL.

# Features

Features	Description
ENABLE	Enables features of this package.
DISABLE	Disables features of this package.
SERVEROUTPUT	Controls whether messages are sent.
PUT	Sends messages.
PUT_LINE	Sends messages with a newline character appended.
NEW_LINE	Sends a newline character.
GET_LINE	Retrieves a line from the message buffer.
GET_LINES	Retrieves multiple lines from the message buffer.

# Syntax

```
{ ENABLE( [buffSize ])
| DISABLE()
| SERVEROUTPUT(sendMsgs)
```

```
| PUT(str)
| PUT_LINE(str)
| NEW_LINE()
| GET_LINE()
| GET_LINES(maxLineNum)
```

# 9.5.1.1 Description

This section explains each feature of DBMS\_OUTPUT.

#### **ENABLE**

- ENABLE enables the use of PUT, PUT\_LINE, NEW\_LINE, GET\_LINE, and GET\_LINES.
- With multiple executions of ENABLE, the value specified last is the buffer size (in bytes). Specify a buffer size from 2000 to 1000000.
- The default value of the buffer size is 20000. If NULL is specified as the buffer size, 1000000 will be used.
- If ENABLE has not been executed, PUT, PUT\_LINE, NEW\_LINE, GET\_LINE, and GET\_LINES are ignored even if they are executed.



PERFORM DBMS\_OUTPUT.ENABLE(20000);

#### **DISABLE**

- DISABLE disables the use of PUT, PUT\_LINE, NEW\_LINE, GET\_LINE, and GET\_LINES.
- Remaining buffer information is discarded.



PERFORM DBMS\_OUTPUT.DISABLE();

# **SERVEROUTPUT**

- SERVEROUTPUT controls whether messages are sent.
- Specify TRUE or FALSE for sendMsgs.
- If TRUE is specified, when PUT, PUT\_LINE, or NEW\_LINE is executed, the message is sent to a client such as psql and not stored in the buffer.

.....

- If FALSE is specified, when PUT, PUT\_LINE, or NEW\_LINE is executed, the message is stored in the buffer and not sent to a client such as psql.



Refer to "Boolean Type" in "Data Types" in "The SQL Language" in the PostgreSQL Documentation for information on boolean type (TRUE/FALSE) values.



PERFORM DBMS\_OUTPUT.SERVEROUTPUT(TRUE);

### PUT

- PUT sets the message to be sent.
- The string is the message to be sent.
- When TRUE is specified for SERVEROUTPUT, the messages are sent to clients such as psql.
- When FALSE is specified for SERVEROUTPUT, the messages are retained in the buffer.
- PUT does not append a newline character. To append a newline character, execute NEW\_LINE.
- If a string longer than the buffer size specified in ENABLE is sent, an error occurs.



PERFORM DBMS\_OUTPUT.PUT('abc');

#### **PUT\_LINE**

- PUT\_LINE sets the message to be sent appended with a newline character.
- The string is the message that is sent.
- When TRUE is specified for SERVEROUTPUT, the messages are sent to clients such as psql.
- When FALSE is specified for SERVEROUTPUT, the messages are retained in the buffer.
- PUT\_LINE appends a newline character to the end of messages.
- If a string longer than the buffer size specified in ENABLE is sent, an error occurs.



PERFORM DBMS\_OUTPUT.PUT\_LINE('abc');

### NEW\_LINE

- NEW\_LINE appends a newline character to the message created with PUT.
- When TRUE is specified for SERVEROUTPUT, the messages are sent to clients such as psql.
- When FALSE is specified for SERVEROUTPUT, the messages are retained in the buffer.



PERFORM DBMS\_OUTPUT.NEW\_LINE();

# **GET\_LINE**

- GET\_LINE retrieves a line from the message buffer.
- Use a SELECT statement to obtain the retrieved line and status code returned by the operation, which are stored in the line and status columns.
- The line column stores the line retrieved from the buffer. The data type of line is TEXT.
- The status column stores the status code returned by the operation: 0-completed successfully; 1-failed because there are no more lines in the buffer. The data type of status is INTEGER.
- If GET\_LINE or GET\_LINES is executed and then PUT, PUT\_LINE, or NEW\_LINE is while messages that have not been retrieved from the buffer still exist, the messages not retrieved from the buffer will be discarded.



```
DECLARE

buff1 VARCHAR(20);

stts1 INTEGER;

BEGIN

SELECT line, status INTO buff1, stts1 FROM DBMS_OUTPUT.GET_LINE();
```

#### **GET\_LINES**

- GET\_LINES retrieves multiple lines from the message buffer.
- Use a SELECT statement to obtain the retrieved lines and the number of lines retrieved, which are stored in the lines and numlines columns.
- The lines column stores the lines retrieved from the buffer. The data type of lines is TEXT.
- The numlines column stores the number of lines retrieved from the buffer. The data type of numlines is INTEGER.
- maxLineNum is the maximum number of lines to retrieve from the buffer. The data type is INTEGER.
- If GET\_LINE or GET\_LINES is executed and then PUT, PUT\_LINE, or NEW\_LINE is executed while messages that have not been retrieved from the buffer still exist, the messages not retrieved from the buffer will be discarded.



```
DECLARE

buff VARCHAR(20)[10];

stts INTEGER := 10;

BEGIN

SELECT lines, numlines INTO buff,stts FROM DBMS_OUTPUT.GET_LINES(stts);
```

# 9.5.1.2 Example

A usage example of DBMS\_OUTPUT is shown below.

```
CREATE FUNCTION dbms_output_exe() RETURNS VOID AS $$

DECLARE

buff1  VARCHAR(20);

buff2  VARCHAR(20);

stts1  INTEGER;

stts2  INTEGER;

BEGIN

PERFORM DBMS_OUTPUT.DISABLE();
```

```
PERFORM DBMS_OUTPUT.ENABLE();
    PERFORM DBMS_OUTPUT.SERVEROUTPUT(FALSE);
    PERFORM DBMS_OUTPUT.PUT('DBMS_OUTPUT TEST 1');
    PERFORM DBMS_OUTPUT.NEW_LINE();
    PERFORM DBMS_OUTPUT.PUT_LINE('DBMS_OUTPUT TEST 2');
    SELECT line, status INTO buff1, stts1 FROM DBMS_OUTPUT.GET_LINE();
    SELECT line, status INTO buff2, stts2 FROM DBMS_OUTPUT.GET_LINE();
    PERFORM DBMS_OUTPUT.SERVEROUTPUT(TRUE);
    PERFORM DBMS_OUTPUT.PUT_LINE(buff1);
    PERFORM DBMS_OUTPUT.PUT_LINE(buff2);
END;

$$ LANGUAGE plpgsql;
SELECT dbms_output_exe();
DROP FUNCTION dbms_output_exe();
```

# 9.5.2 UTL\_FILE

#### Overview

Text files can be written and read using PL/pgSQL.

To perform these file operations, the directory for the operation target must be registered in the UTL\_FILE.UTL\_FILE\_DIR table beforehand. Use the INSERT statement as the database administrator or a user who has INSERT privileges to register the directory. Also, if the directory is no longer necessary, delete it from the same table. Refer to "9.5.2.1 Registering and Deleting Directories" for information on the how to register and delete the directory.



- When performing file operations, access privileges for the UTL\_FILEUTL\_FILE\_DIR table are required according
  to the operation. However, because the user performing file operations can access any file regardless of the operating
  system access privileges for the target file, set the access privileges for the table appropriately so that important files
  are not accessed by any user.
- It is not available on Windows (R).

Refer to "C.1 UTL\_FILE.UTL\_FILE\_DIR" for information on the UTL\_FILE.UTL\_FILE\_DIR table.

Declare the file handler explained hereafter as follows in PL/pgSQL:

```
DECLARE
f UTL_FILE.FILE_TYPE;
```

#### **Features**

Feature	Description
FCLOSE	Closes a file.
FCLOSE_ALL	Closes all files open in a session.
FCOPY	Copies a whole file or a contiguous portion thereof.
FFLUSH	Flushes the buffer.
FGETATTR	Retrieves the attributes of a file.
FOPEN	Opens a file.
FRENAME	Renames a file.
GET_LINE	Reads one line from a text file.
IS_OPEN	Checks if a file is open.
NEW_LINE	Writes newline characters.

Feature	Description
PUT	Writes a string.
PUT_LINE	Appends a newline character to a string and writes the string.
PUTF	Writes a formatted string.

## Syntax

```
{ FCLOSE(fileHandle)
| FCLOSE_ALL()
| FCOPY(srcDir, srcFileName, destDir, destFileName
       [ {,startLine | ,startLine ,endLine } ])
| FFLUSH(fileHandle)
| FGETATTR(dir, filename)
| FOPEN(dir, fileName, openMode [, maxLineSize ])
| FRENAME(srcDir, srcFileName, destDir, destFileName [ ,overwrite ])
| GET_LINES(fileHandle [,len ])
| IS_OPEN(fileHandle)
| NEW_LINE(fileHandle [, numOfNewLines ])
| PUT(fileHandle, str)
| PUT_LINE(fileHandle, str [, writeToFile ])
| PUT_LINE(fileHandle, fmt [, args ]... ])
```

# 9.5.2.1 Registering and Deleting Directories

# Registering the directory

1. Check if the directory is already registered (if it is, then step 2 is not necessary).

```
SELECT * FROM UTL_FILE.UTL_FILE_DIR WHERE dir='/home/fsep';
```

2. Register the directory.

```
INSERT INTO UTL_FILE.UTL_FILE_DIR VALUES('/home/fsep');
```

# **Deleting the directory**

```
DELETE FROM UTL_FILE.UTL_FILE_DIR WHERE dir='/home/fsep';
```

# 9.5.2.2 Description

This section explains each feature of UTL\_FILE.

# **FCLOSE**

- FCLOSE closes a file that is open.
- Specify an open file handle.
- The value returned is a NULL value.



```
f := UTL_FILE.FCLOSE(f);
```

# FCLOSE\_ALL

- FCLOSE\_ALL closes all files open in a session.
- Files closed with FCLOSE\_ALL can no longer be read or written.



PERFORM UTL\_FILE.FCLOSE\_ALL();

#### **FCOPY**

- FCOPY copies a whole file or a contiguous portion thereof. The whole file is copied if *startLine* and *endLine* are not specified.
- Specify the directory location of the source file.
- Specify the source file.
- Specify the directory where the destination file will be created.
- Specify the file name of the destination file.
- Specify the line number at which to begin copying. Specify a value greater than 0. If not specified, 1 is used.
- Specify the line number at which to stop copying. If not specified, the last line number of the file is used.



PERFORM UTL\_FILE.FCOPY('/home/fsep', 'regress\_fsep.txt', '/home/fsep', 'regress\_fsep2.txt');

# **FFLUSH**

- FFLUSH forcibly writes the buffer data to a file.
- Specify an open file handle.



PERFORM UTL\_FILE.FFLUSH(f);

# **FGETATTR**

- FGETATTR retrieves file attributes: file existence, file size, and information about the block size of the file.

.....

- Specify the directory where the file exists.
- Specify the file name.
- Use a SELECT statement to obtain the file attributes, which are stored in the fexists, file\_length, and blocksize columns.
- The fexists column stores a boolean (TRUE/FALSE) value. If the file exists, fexists is set to TRUE. If the file does not exist, fexists is set to FALSE. The data type of fexists is BOOLEAN.
- The file\_length column stores the length of the file in bytes. If the file does not exist, file\_length is NULL. The data type of file\_length is INTEGER.

- The blocksize column stores the block size of the file in bytes. If the file does not exist, blocksize is NULL. The data type of blocksize is INTEGER.

......



```
SELECT fexists, file_length, blocksize INTO file_flag, file_len, size FROM UTL_FILE.FGETATTR('/home/fsep', 'regress_fsep.txt');
```

#### **FOPEN**

- FOPEN opens a file.
- Specify the directory where the file exists.
- Specify the file name.
- Specify the mode for opening the file:
  - r: Read
  - w: Write
  - a: Add
- Specify the maximum string length (in bytes) that can be processed with one operation. If omitted, the default is 1024. Specify a value from 1 to 32767.
- Up to 50 files per session can be open at the same time.



```
f := UTL_FILE.FOPEN('/home/fsep','regress_fsep.txt','r',1024);
```

### **FRENAME**

- FRENAME renames a file.
- Specify the directory location of the source file.
- Specify the source file to be renamed.
- Specify the directory where the renamed file will be created.
- Specify the new name of the file.
- Specify whether to overwrite a file if one exists with the same name and in the same location as the renamed file. If TRUE is specified, the existing file will be overwritten. If FALSE is specified, an error occurs. If omitted, FALSE is set.



#### See

Refer to "Boolean Type" in "Data Types" in "The SQL Language" in the PostgreSQL Documentation for information on boolean type (TRUE/FALSE) values.

# 撑 Example

```
PERFORM UTL_FILE.FRENAME('/home/fsep', 'regress_fsep.txt', '/home/fsep', 'regress_fsep2.txt', TRUE);
```

.....

# **GET\_LINE**

- GET\_LINE reads one line from a file.
- Specify the file handle returned by FOPEN using the r (read) mode.
- Specify the number of bytes to read from the file. If not specified, the maximum string length specified at FOPEN will be used.
- The return value is the buffer that receives the line read from the file.
- Newline characters are not loaded to the buffer.
- An empty string is returned if a blank line is loaded.
- Specify the maximum length (in bytes) of the data to be read. Specify a value from 1 to 32767. If not specified, the maximum string length specified at FOPEN is set. If no maximum string length is specified at FOPEN. 1024 is set.
- If the line length is greater than the specified number of bytes to read, the remainder of the line is read on the next call.
- A NO\_DATA\_FOUND exception will occur when trying to read past the last line.



buff := UTL\_FILE.GET\_LINE(f);

# IS\_OPEN

- IS\_OPEN checks if a file is open.
- Specify the file handle.
- The return value is a BOOLEAN type. TRUE represents an open state and FALSE represents a closed state.



#### See

Refer to "Boolean Type" in "Data Types" in "The SQL Language" in the PostgreSQL Documentation for information on boolean type (TRUE/FALSE) values.

# 🔰 Example

```
IF UTL_FILE.IS_OPEN(f) THEN
    PERFORM UTL_FILE.FCLOSE(f);
END IF;
```

#### **NEW LINE**

- NEW\_LINE writes one or more newline characters.
- Specify an open file handle.
- Specify the number of newline characters to be written to the file. If omitted, "1" is used.



PERFORM UTL\_FILE.NEW\_LINE(f, 2);

.....

#### PUT

- PUT writes a string to a file.
- Specify the file handle that was opened with FOPEN using w (write) or a (append).
- Specify the string to be written to the file.
- The maximum length (in bytes) of the string to be written is the maximum string length specified at FOPEN.
- The return value is a TEXT type and is the buffer that receives the line loaded from the file.



PERFORM UTL\_FILE.PUT(f, 'ABC');

# **PUT\_LINE**

- PUT\_LINE appends a newline character to a string and writes the string.
- Specify the file handle that was opened with FOPEN w (write) or a (append).
- Specify whether to forcibly write to the file. If TRUE is specified, file writing is forced. If FALSE is specified, file writing is asynchronous. If omitted, FALSE will be set.

.....

- The maximum length of the string (in bytes) is the maximum string length specified at FOPEN.



PERFORM UTL\_FILE.PUT\_LINE(f, 'ABC', TRUE);

#### **PUTF**

- PUTF writes a formatted string.
- Specify the file handle that was opened with FOPEN w (write) or a (append).
- Specify the format, which is a string that includes the formatting characters \n and %s.
- The  $\n$  in the format is code for a newline character.
- Specify the same number of input values as there are %s in the format. Up to a maximum of five input values can be specified. The %s in the format are replaced with the corresponding input characters. If an input value corresponding to %s is not specified, it is replaced with an empty string.



PERFORM UTL\_FILE.PUTF(f, '[1=%s, 2=%s, 3=%s, 4=%s, 5=%s]\n', '1', '2', '3', '4', '5');

# 9.5.2.3 Example

The procedure when using UTL\_FILE, and a usage example, are shown below.

#### 1. Preparation

Before starting a new job that uses UTL\_FILE, register the directory in the UTL\_FILE.UTL\_FILE\_DIR table.

Refer to "9.5.2.1 Registering and Deleting Directories" for information on how to register the directory.

#### 2. Performing a job

Perform a job that uses UTL\_FILE. The example is shown below.

```
CREATE OR REPLACE FUNCTION gen_file(mydir TEXT, infile TEXT, outfile TEXT, copyfile
TEXT) RETURNS void AS $$
DECLARE
  v1 VARCHAR(32767);
  inf UTL_FILE.FILE_TYPE;
  otf UTL_FILE.FILE_TYPE;
BEGIN
  inf := UTL_FILE.FOPEN(mydir, infile,'r',256);
  otf := UTL_FILE.FOPEN(mydir, outfile,'w');
  v1 := UTL_FILE.GET_LINE(inf,256);
  PERFORM UTL_FILE.PUT_LINE(otf,v1,TRUE);
  v1 := UTL_FILE.GET_LINE(inf,256);
  PERFORM UTL_FILE.PUTF(otf,'%s\n',v1);
  v1 := UTL_FILE.GET_LINE(inf, 256);
  PERFORM UTL_FILE.PUT(otf,v1);
  PERFORM UTL_FILE.NEW_LINE(otf);
  PERFORM UTL_FILE.FFLUSH(otf);
  inf := UTL_FILE.FCLOSE(inf);
  otf := UTL_FILE.FCLOSE(otf);
  PERFORM UTL_FILE.FCOPY(mydir, outfile, mydir, copyfile, 2, 3);
  PERFORM UTL_FILE.FRENAME(mydir, outfile, mydir, 'rename.txt');
$$ LANGUAGE plpgsql;
SELECT gen_file('/home/fsep', 'input.txt', 'output.txt', 'copyfile.txt');
```

#### 3. Post-processing

If you remove a job that uses UTL\_FILE, delete the directory information from the UTL\_FILE.UTL\_FILE\_DIR table. Ensure that the directory information is not being used by another job before deleting it.

Refer to "9.5.2.1 Registering and Deleting Directories" for information on how to delete the directory.

# 9.5.3 DBMS\_SQL

#### Overview

Dynamic SQL can be executed from PL/pgSQL.

# **Features**

Feature	Description
BIND_VARIABLE	Sets values in the host variable within the SQL statement.
CLOSE_CURSOR	Closes the cursor.
COLUMN_VALUE	Retrieves the value of the column in the select list extracted with FETCH_ROWS.
DEFINE_COLUMN	Defines the column from which values are extracted and the storage destination.
EXECUTE	Executes SQL statements.

Feature	Description	
FETCH_ROWS	Positions the specified cursor at the next row and extracts values from the row.	
OPEN_CURSOR	Opens a new cursor.	
PARSE	Parses SQL statements.	



- In DBMS\_SQL, the data types supported in dynamic SQL are limited, and therefore the user must consider this. The supported data types are:

- INTEGER
- DECIMAL
- NUMERIC
- REAL
- DOUBLE PRECISION
- CHAR(\*1)
- VARCHAR(\*1)
- NCHAR(\*1)
- NCHAR VARYING(\*1)
- TEXT
- DATE
- TIMESTAMP WITHOUT TIME ZONE
- TIMESTAMP WITH TIME ZONE
- INTERVAL(\*2)
- SMALLINT
- BIGINT

\*1:

The host variables with CHAR, VARCHAR, NCHAR, and NCHAR VARYING data types are treated as TEXT, to match the string function arguments and return values. Refer to "String Functions and Operators" in "Functions and Operators" in "The SQL Language" in the PostgreSQL Documentation for information on string functions. When specifying the arguments of the features compatible with Oracle databases NVL and/or DECODE, use CAST to convert the data types of the host variables to ensure that data types between arguments are the same.

\*2:

When using COLUMN\_VALUE to obtain an INTERVAL type value specified in the select list, use an INTERVAL type variable with a wide range such as when no interval qualifier is specified, or with a range that matches that of the variable in the select list. If an interval qualifier variable with a narrow range is specified, then the value within the interval qualifier range will be obtained, but an error that the values outside the range have been truncated will not occur.



This example illustrates where a value expression that returns an INTERVAL value is set in the select list and the result is received with COLUMN\_VALUE. Note that the SQL statement operation result returns a value within the INTERVAL DAY TO SECOND range.

# [Bad example]

Values of MINUTE, and those after MINUTE, are truncated, because the variable(v\_interval) is INTERVAL DAY TO HOUR.

```
v_interval INTERVAL DAY TO HOUR;
...
    PERFORM DBMS_SQL.PARSE(cursor, 'SELECT CURRENT_TIMESTAMP - ''2010-01-01'' FROM
DUAL', 1);
...
    SELECT value INTO v_interval FROM DBMS_SQL.COLUMN_VALUE(cursor, 1, v_interval);
    result:1324 days 09:00:00
```

# [Good example]

By ensuring that the variable(v\_interval) is INTERVAL, the values are received correctly.

```
v_interval INTERVAL;
...
    PERFORM DBMS_SQL.PARSE(cursor, 'SELECT CURRENT_TIMESTAMP - ''2010-01-01'' FROM
DUAL', 1);
...
    SELECT value INTO v_interval FROM DBMS_SQL.COLUMN_VALUE(cursor, 1, v_interval);
    result:1324 days 09:04:37.530623
```

Syntax

```
{ BIND_VARIABLE(cursor, varName, val [, len ])
| CLOSE_CURSOR(cursor)
| COLUMN_VALUE(cursor, colPos, varName)
| DEFINE_COLUMN(cursor, colPos, varName [, len ])
| EXECUTE(cursor)
| FETCH_ROWS(cursor)
| OPEN_CURSOR([parm1 ])
| PARSE(cursor, sqlStmt, parm1 [, parm2, parm3, parm4 ])
}
```

# 9.5.3.1 Description

This section explains each feature of DBMS\_SQL.

# BIND\_VARIABLE

- BIND\_VARIABLE sets values in the host variable within the SQL statement.
- Specify the cursor number to be processed.
- Specify the name of the host variable within the SQL statement using a string for the host variable name.
- Specify the value set in the host variable. The data type of the host variable is the same as that of the value expression
   it is implicitly converted in accordance with its position within the SQL statement. Refer to "A.3 Implicit Data Type Conversions" for information on implicit conversions.
- If the value is a character type, the string length is the number of characters. If the string length is not specified, the size is the total length of the string.
- It is necessary to place a colon at the beginning of the host variable in SQL statements to identify the host variable.
   The colon does not have to be added to the host variable names specified at BIND\_VARIABLE. The following shows examples of host variable names specified with SQL statements and host variable names specified with BIND\_VARIABLE:

```
PERFORM DBMS_SQL.PARSE(cursor, 'SELECT emp_name FROM emp WHERE sal > :x', 1);
```

In this example, BIND\_VARIABLE will be as follows:

PERFORM DBMS\_SQL.BIND\_VARIABLE(cursor, ':x', 3500);

Or,

PERFORM DBMS\_SQL.BIND\_VARIABLE(cursor, 'x', 3500);

- The length of the host variable name can be up to 30 bytes (excluding colons).
- If the data type of the set value is string, specify the effective size of the column value as the fourth argument.



If the data type of the value to be set is not a string:

```
PERFORM DBMS_SQL.BIND_VARIABLE(cursor, ':NO', 1);
```

If the data type of the value to be set is a string:

PERFORM DBMS\_SQL.BIND\_VARIABLE(cursor, ':NAME', h\_memid, 5);

#### CLOSE\_CURSOR

- CLOSE\_CURSOR closes the cursor.
- Specify the cursor number to be processed.
- The value returned is a NULL value.



cursor := DBMS\_SQL.CLOSE\_CURSOR(cursor);

# COLUMN\_VALUE

- COLUMN\_VALUE retrieves the value of the column in the select list extracted with FETCH\_ROWS.
- Specify the cursor number to be processed.
- Specify the position of the column of the select list in the SELECT statement. The position of the first column is 1.
- Specify the destination variable name.
- Use a SELECT statement to obtain the values of the value, column\_error, and actual\_length columns.
- The value column returns the value of the column specified at the column position. The data type of the variable name must match that of the column. If the data type of the column in the SELECT statement specified in PARSE is not compatible with DBMS\_SQL, use CAST to convert to a compatible data type.
- The data type of the column\_error column is NUMERIC. If the column value could not be set correctly in the value column, a value other than 0 will be returned:

22001: The extracted string has been truncated

22002: The extracted value contains a NULL value

- The data type of the actual\_length column is INTEGER. If the extracted value is a character type, the number of characters will be returned (if the value was truncated, the number of characters prior to the truncation will be returned), otherwise, the number of bytes will be returned.



When retrieving the value of the column, the error code, and the actual length of the column value:

```
\label{eq:select_value} SELECT\ value,\ column\_error,\ actual\_length\ INTO\ v\_memid,\ v\_col\_err,\ v\_act\_len\ FROM\ DBMS\_SQL.COLUMN\_VALUE(cursor,\ 1,\ v\_memid);
```

When retrieving just the value of the column:

```
SELECT value INTO v_memid FROM DBMS_SQL.COLUMN_VALUE(cursor, 1, v_memid);
```

#### DEFINE\_COLUMN

- DEFINE\_COLUMN defines the column from which values are extracted and the storage destination.
- Specify the cursor number to be processed.
- Specify the position of the column in the select list in the SELECT statement. The position of the first column is 1.
- Specify the destination variable name. The data type should be match with the data type of the column from which the value is to be extracted. If the data type of the column in the SELECT statement specified in PARSE is not compatible with DBMS\_SQL, use CAST to convert to a compatible data type.
- Specify the maximum number of characters of character type column values.
- If the data type of the column value is string, specify the effective size of the column value as the fourth argument.



When the data type of the column value is not a string:

```
PERFORM DBMS_SQL.DEFINE_COLUMN(cursor, 1, v_memid);
```

When the data type of the column value is a string:

```
PERFORM DBMS_SQL.DEFINE_COLUMN(cursor, 1, v_memid, 10);
```

### **EXECUTE**

- EXECUTE executes SQL statements.
- Specify the cursor number to be processed.
- The return value is an INTEGER type, is valid only with INSERT statement, UPDATE statement, and DELETE statement, and is the number of rows processed. Anything else is invalid.



```
ret := DBMS_SQL.EXECUTE(cursor);
```

# FETCH\_ROWS

- FETCH\_ROWS positions at the next row and extracts values from the row.
- Specify the cursor number to be processed.
- The return value is an INTEGER type and is the number of rows extracted. 0 is returned if all are extracted.

- The extracted information is retrieved with COLUMN\_VALUE.



```
LOOP

IF DBMS_SQL.FETCH_ROWS(cursor) = 0 THEN

EXIT;

END IF;

...

END LOOP;
```

# OPEN\_CURSOR

- OPEN\_CURSOR opens a new cursor.
- The parameter is used for compatibility with Oracle databases only, and is ignored by FUJITSU Enterprise Postgres. An INTEGER type can be specified, but it will be ignored. If migrating from an Oracle database, specify 1.
- Close unnecessary cursors by executing CLOSE\_CURSOR.
- The return value is an INTEGER type and is the cursor number.



```
cursor := DBMS_SQL.OPEN_CURSOR();
```

# **PARSE**

- PARSE analyzes dynamic SQL statements.
- Specify the cursor number to be processed.
- Specify the SQL statement to be parsed.
- Parameters 1, 2, 3, and 4 are used for compatibility with Oracle databases only, and are ignored by FUJITSU Enterprise Postgres. If you are specifying values anyway, specify the following:
  - Parameter 1 is an INTEGER type. Specify 1.
  - Parameters 2 and 3 are TEXT types. Specify NULL.
  - Parameter 4 is a BOOLEAN type. Specify TRUE.

If migrating from an Oracle database, the specified values for parameters 2, 3, and 4 do not need to be changed.

- Add a colon to the beginning of host variables in SQL statements.
- The DDL statement is executed when PARSE is issued. EXECUTE is not required for the DDL statement.
- If PARSE is called again for opened cursors, the content in the data regions within the cursors is reset, and the SQL statement is parsed anew.



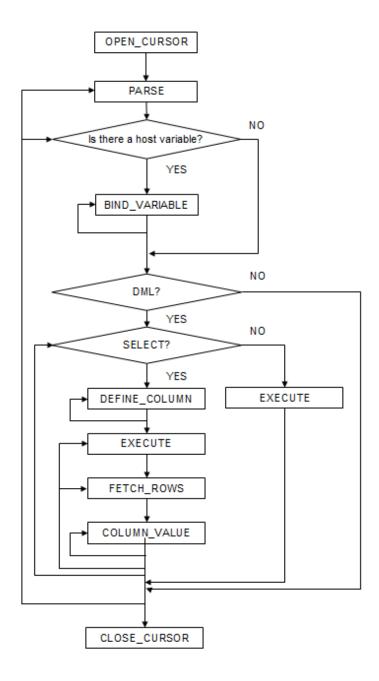
PERFORM DBMS\_SQL.PARSE(cursor, 'SELECT memid, memnm FROM member WHERE memid = :NO', 1);

# 9.5.3.2 Example

This section explains the flow of DBMS\_SQL and provides an example.

# Flow of DBMS\_SQL

Flow of DBMS\_SQL



# Example

```
CREATE FUNCTION smp_00()
RETURNS INTEGER
AS $$
DECLARE
str_sql VARCHAR(255);
cursor INTEGER;
h_smpid INTEGER;
```

```
v_smpid INTEGER;
   v_smpnm VARCHAR(20);
   v_smpage INTEGER;
   errcd INTEGER;
length INTEGER;
   ret
              INTEGER;
BEGIN
   str_sql := 'SELECT smpid, smpnm, smpage FROM smp_tbl WHERE smpid < :H_SMPID ORDER BY
smpid';
   h_smpid
            := 3;
              := 0;
   v_smpid
   v_smpnm
               := '';
   v_smpage
             := 0;
   cursor := DBMS_SQL.OPEN_CURSOR();
   PERFORM DBMS_SQL.PARSE(cursor, str_sql, 1);
   PERFORM DBMS_SQL.BIND_VARIABLE(cursor, ':H_SMPID', h_smpid);
   PERFORM DBMS_SQL.DEFINE_COLUMN(cursor, 1, v_smpid);
   PERFORM DBMS_SQL.DEFINE_COLUMN(cursor, 2, v_smpnm, 10);
   PERFORM DBMS_SQL.DEFINE_COLUMN(cursor, 3, v_smpage);
   ret := DBMS_SQL.EXECUTE(cursor);
   loop
       if DBMS_SQL.FETCH_ROWS(cursor) = 0 then
       end if;
       SELECT value,column_error,actual_length INTO v_smpid,errcd,length FROM
DBMS_SQL.COLUMN_VALUE(cursor, 1, v_smpid);
       RAISE NOTICE '----';
       RAISE NOTICE '----';
       RAISE NOTICE 'smpid = %', v_smpid;
RAISE NOTICE 'errcd = %', errcd;
RAISE NOTICE 'length = %', length;
       SELECT value,column_error,actual_length INTO v_smpnm,errcd,length FROM
DBMS_SQL.COLUMN_VALUE(cursor, 2, v_smpnm);
       RAISE NOTICE '-----
       RAISE NOTICE 'smpnm
RAISE NOTICE 'errcd
RAISE NOTICE 'length
                                = %', v_smpnm;
                                = %', errcd;
                                 = %', length;
       select value,column_error,actual_length INTO v_smpage,errcd,length FROM
DBMS_SQL.COLUMN_VALUE(cursor, 3, v_smpage);
       RAISE NOTICE '----';
       RAISE NOTICE 'smpage = %', v_smpage;
RAISE NOTICE 'errcd = %', errcd;
RAISE NOTICE 'length = %', length;
       RAISE NOTICE '';
   end loop;
   cursor := DBMS_SQL.CLOSE_CURSOR(cursor);
   RETURN 0;
END;
$$ LANGUAGE plpgsql;
```

# **Chapter 10 Application Connection Switch Feature**

The application connection switch feature enables automatic connection to the target server when there are multiple servers with redundant configurations.

When using this feature, specify the primary server and secondary server as the connected servers in the application connection information. A standby server can optionally be prioritized over the primary server as the target server.

If an application connection switch occurs, explicitly close the connection and then reestablish the connection or reexecute the application. Refer to "Errors when an Application Connection Switch Occurs and Corresponding Actions" of the relevant client interface for information on how to confirm the switch.

# 10.1 Connection Information for the Application Connection Switch Feature

To use the application connection switch feature, set the information shown below when connecting the database.

#### IP address or host name

Specify the IP address or host name that will be used to configure the database multiplexing system.

#### Port number

A port number used by each database server to listen for connections from applications.

In each client interface, multiple port numbers can be specified, however in the format shown below, for example:

host1,host2:port2

JDBC and .NET

If only one port number is specified, it will be assumed that host1: 27500 (the default value) and host2:port2 were specified.

Omit all port numbers, or specify only one per server.

Others

If only one port number is specified, it will be assumed that the same port is used for all the hosts.

### Target server

From the specified connection destination server information, specify the selection sequence of the servers to which the application will connect. The values specified for the target server have the meanings shown below. If a value is omitted, "any" will be assumed.

# Primary server

The primary server is selected as the connection target from the specified "IP addresses or host names". Specify this to perform tasks that can be performed only on the primary server, such as applications in line with updates, or management tasks such as REINDEX and VACUUM.

# Standby server (this value can be used only when the JDBC or .NET driver is used)

The standby server is selected as the connection target from the specified "IP addresses or host names". On standby server, the update will always fail. If the target server is not standby, the JDBC driver will throw an error stating that it is unable to find a server with the specified targetServerType.

# Priority given to a standby server

The standby server is selected preferentially as the connection target from the specified "IP addresses or host names". If there is no standby server, the application will connect to the primary server.

# Any

This method is not recommended in database multiplexing systems. This is because, although the connection destination server is selected in the specified sequence from the specified "IP addresses or host names", if the server that was successfully connected to first is the standby server, the write operations will always fail.

The table below shows the server selection order values to set for each driver:

Server selection order	JDBC and .NET drivers	Other drivers
Primary server	"master"	"read-write"
Standby server	"slave"	-
Priority given to a standby server	"preferSlave"	"prefer-read"
Any	"any"	"any"

### SSL server certificate Common Name (CN)

To perform SSL authentication by creating the same server certificate for each server in a multiplexing system, specify the SSL server certificate Common Name (CN) in this parameter. Accordingly, SSL authentication using the CN can be performed without having to consider the names of the multiple servers contained in the multiplexing system.

# 10.2 Using the Application Connection Switch Feature

This section explains how to set the connection destination server using the application connection switch feature.

Of the parameters used as connection information for each client interface, only the parameters specific to the application connection switch feature are explained here. Refer to "Setup" and "Connecting to the Database" for information on the other parameters of each client interface.

# 10.2.1 Using the JDBC Driver

Set the following information in the connection string of the DriverManager class, or in the data source.

Table 10.1 Information to be set

Argument	Explanation
host1 host2	Specify the IP address or host name.
port1 port2	Specify the port number for the connection. The port number can be omitted. If omitted, the default is 27500.
database_name	Specify the database name.
targetServerType	Specify the selection sequence of the servers to which the application will connect.  Refer to "Target server" for details.
sslmode	Specify this to encrypt communications. By default, this is disabled. The setting values for sslmode are as follows:
	disable: Connect without SSL
	require: Connect always with SSL
	verify-ca: Connect with SSL, using a certificate issued by a trusted CA (*1)
	verify-full: Connect with SSL, using a certificate issued by a trusted CA to verify if the server host name matches the certificate (*1)
sslservercertcn	This parameter is enabled only to perform SSL authentication (sslmode=verify-full).
	Specify the server certificate CN. If this is omitted, the value will be null, and the server certificate CN will be authenticated using the host name specified in host.

<sup>\*1:</sup> If specifying either "verify-ca" or "verify-full", the CA certificate file can be specified using connection string sslrootcert.

### When using Driver Manager

Specify the following URL in the API of the DriverManager class:

```
jdbc:postgresql://host1[:port1],host2[:port2]/dbName[?targetServerType={master | slave
| preferSlave | any}][&sslmode=verify-
full&sslrootcert=cACertificateFile&sslservercertcn=targetServerCertificateCN]
```

- If the target server is omitted, the default value "any" is used.
- When using IPV6, specify the host in the "[host]" (with square brackets) format.

### [Example]

```
jdbc:postgresql://[2001:Db8::1234]:27500,192.168.1.1:27500/dbName
```

#### When using the data source

Specify the properties of the data source in the following format:

```
source.setServerName("host1[:port1],host2[:port2]");
source.setTargetServerType("master");
source.setSslmode("verify-full");
source.setSslrootcert("cACertificateFile");
source.setSslservercertcn("targetServerCertificateCN");
```

- If the port number is omitted, the value specified in the portNumber property will be used. Also, if the portNumber property is omitted, the default is 27500.
- If the target server is omitted, the value will be "any".
- When using IPV6, specify the host in the "[host]" (with square brackets) format.

# [Example]

```
source.setServerName("[2001:Db8::1234]:27500,192.168.1.1:27500");
```



If using the connection parameter loginTimeout, the value will be applied for the time taken attempting to connect to all of the specified hosts.

# 10.2.2 Using the ODBC Driver

Set the following information in the connection string or data source.

Table 10.2 Information to be set

Parameter	Explanation
Servername	Specify IP address 1 and IP address 2, or the host name, using a comma as the delimiter. Based on ODBC rules, it is recommended to enclose the whole string containing comma delimiters with {}.  Format: {host1,host2}
Port	Specify the connection destination port numbers, using a comma as the delimiter.  Based on ODBC rules, it is recommended to enclose the whole string containing comma delimiters with {}.  Format: {port1,port2}

Parameter	Explanation	
	Specify the port number corresponding to the IP address or host specified for the nth Servername as the nth Port.	
	The port number can be omitted. If omitted, the default is 27500.	
	If <i>n</i> server names are specified, and m ports are specified then there will be error reported. The only exceptions are where m=n or m=1. In case only one port is specified, then the same is applied for all the hosts.	
target_session_attrs	Specify the selection sequence of the servers to which the application will connect.  Refer to "Target server" for details.	
SSLMode	Specify this to encrypt communications. By default, this is disabled. The setting values for SSLMode are as follows:	
	disable: Connect without SSL	
	allow: Connect without SSL, and if it fails, connect with SSL	
	prefer: Connect with SSL, and if it fails, connect without SSL	
	require: Connect always with SSL	
	verify-ca: Connect with SSL, using a certificate issued by a trusted CA (*1)	
	verify-full: Connect with SSL, using a certificate issued by a trusted CA to verify if the server host name matches the certificate (*1)	
SSLServerCertCN	This parameter is enabled only to perform SSL authentication (SSLMode=verifyfull).	
	Specify the server certificate CN. If this is omitted, the value will be null, and the server certificate CN will be authenticated using the host name specified in Servername.	

<sup>\*1:</sup> If specifying either "verify-ca" or "verify-full", use the system environment variable PGSSLROOTCERT of your operating system to specify the CA certificate file as shown below.

Example)

Variable name: PGSSLROOTCERT Variable value: *cACertificateFile* 

# When specifying a connection string

Specify the following connection string:

```
...;Servername={host1,host2};Port={port1,port2};[target_session_attrs={read-write | prefer-read | any}];[ SSLMode=verify-full;SSLServerCertCN=targetServerCertificateCN]...
```

- When using IPV6, specify the host in the "host" format.

#### [Example]

```
Servername={2001:Db8::1234,192.168.1.1};Port={27500,27500};
```

#### When using the data source

Specify the properties of the data source in the following format:

```
Servername={host1,host2}

Port={port1,port2}

target_session_attrs={read-write | prefer-read | any }

SSLMode=verify-full

SSLServerCertCN=targetServerCertificateCN
```

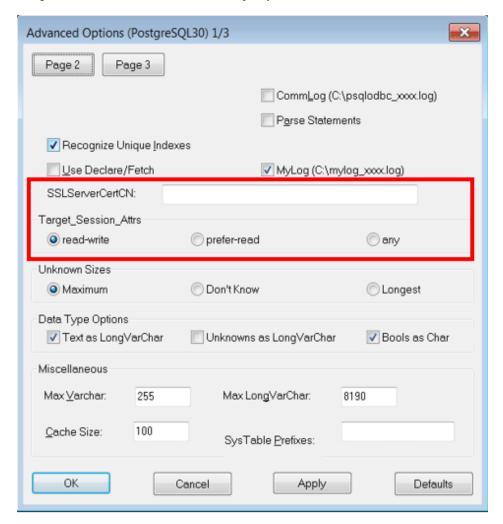
- When using IPV6, specify the host in the "host" format.

#### [Example]

```
Servername={2001:Db8::1234,192.168.1.1}
```

# Registering the data source using the ODBC Data Source Administrator

Using the ODBC Data Source Administrator, specify the items within the red border below:





If using the connection parameter login\_timeout, this value is applied for connections to each of the specified hosts. If both multiplexed database servers have failed, the connection will time out when a time equal to double the login\_timeout value elapses.

# 10.2.3 Using a .NET Data Provider

Set the following information in the connection string of NpgsqlConnection, or in the data source.

Table 10.3 Information to be set

Argument	Explanation
host1 host2	Specify the IP address or host name.
port1 port2	Specify the port number for the connection.
TargetServerType	Specify the selection sequence of the servers to which the application will connect.  Refer to "Target server" for details.

# When specifying a connection string

Specify the following connection string:

```
host1[:port1],host2[:port2];[TargetServerType={TargetServerType.master | TargetServerType.preferSlave | TargetServerType.any}];
```

- If the port number is omitted from the host string, the value specified for the Port keyword of the connection string will be used. Refer to "4.3.4 Connection String" for information on the Port keyword.
- When using IPV6, specify the host in the "[host]" (with square brackets) format.
- If the target server type is omitted, the value will be any.

#### [Example]

```
host=[2001:Db8::1234]:27500,192.168.1.1:27500;
```

When specifying the NpgsqlConnectionStringBuilder property, or adding a connection in TableAdapter Specify the Host property of the data source in the following format:

```
host1[:port1],host2[:port2]
```

- If the port number is omitted from the host string, the value specified in the Port property will be used. Also, if the Port property is omitted, the default is 27500.

Specify the TargetServerType property of the data source in the following format:

TargetServerType.master	TargetServerType.preferSlave	TargetServerType.any
		•

- If the target server type is omitted, the value will be any.



If using the connection parameter Timeout, this value is applied for connections to each of the specified hosts. If both multiplexed database servers have failed, the connection will time out when a time equal to double the Timeout value elapses.

# 10.2.4 Using a Connection Service File

Set the connection parameters as follows.

Table 10.4 Information to be set

Parameter	Explanation	
host	Specify the host names, using a comma as the delimiter.	
hostaddr	Specify IP address 1 and IP address 2, using a comma as the delimiter.	

Parameter	Explanation		
port	Specify the connection destination port numbers, using a comma as the delimiter.  Specify the port number for the server specified for the nth host or hostaddr at the nth port.  The port number can be omitted. If omitted, the default is 27500.  If <i>n</i> server names are specified, and m ports are specified then there will be		
	error reported. The only exceptions are where m=n or m=1. In case only one port is specified, then the same is applied for all the hosts.		
target_session_attrs	Specify the selection sequence of the servers to which the application will connect.  Refer to "Target server" for details.		
sslmode	Specify this to encrypt communications. By default, this is disabled. The setting values for sslmode are as follows:		
	disable: Connect without SSL		
	allow: Connect without SSL, and if it fails, connect with SSL		
	prefer: Connect with SSL, and if it fails, connect without SSL		
	require: Connect always with SSL		
	verify-ca: Connect with SSL, using a certificate issued by a trusted CA (*1		
	verify-full: Connect with SSL, using a certificate issued by a trusted CA to verify if the server host name matches the certificate (*1)		
sslservercertcn	This parameter is enabled only to perform SSL authentication (sslmode=verify-full).		
	Specify the server certificate CN. If this is omitted, the value will be null, and the server certificate CN will be authenticated using the host name specified in host.		

\*1: If specifying either "verify-ca" or "verify-full", use the system environment variable PGSSLROOTCERT (connection parameter sslrootcert) of your operating system to specify the CA certificate file as shown below.

Example)

Variable name: PGSSLROOTCERT Variable value: *cACertificateFile* 



If using the connection parameter connect\_timeout, this value is applied for connections to each of the specified hosts. If both multiplexed database servers have failed, the connection will time out when a time equal to double the connect\_timeout value elapses.



If using the C Library, embedded SQL or psql commands (including other client commands that specify connection destinations), it is recommended to use a connection service file to specify connection destinations.

In the connection service file, a name (service name) is defined as a set, comprising information such as connection destination information and various types of tuning information set for connections. By using the service name defined in the connection service file when connecting to databases, it is no longer necessary to modify applications when the connection information changes.

### 10.2.5 Using the C Library (libpq)

It is recommended that you use a connection service file. Refer to "10.2.4 Using a Connection Service File" for details.

If a connection service file will not be used, set the following information for the database connection control functions (PQconnectdbParams, PQconnectdb, and so on) or environment variables.

Table 10.5 Information to be set

Parameter (environment variable name)	Explanation			
host(PGHOST)	Specify the host names, using a comma as the delimiter.			
hostaddr(PGHOSTADDR)	Specify IP address 1 and IP address 2, using a comma as the delimiter.			
port(PGPORT)	Specify the connection destination port numbers, using a comma as the delimiter. Specify the port number for the server specified for the nth host or hostaddr as the nth port.			
	The port number can be omitted. If omitted, the default is 27500.			
	If <i>n</i> server names are specified, and m ports are specified then there will be error reported. The only exceptions are where m=n or m=1. In case only one port is specified, then the same is applied for all the hosts.			
target_session_attrs(PGTA RGETSESSIONATTRS)	Specify the selection sequence of the servers to which the application will connect.  Refer to "Target server" for details.			
sslmode(PGSSLMODE)	Specify this to encrypt communications. By default, this is disabled.  The setting values for sslmode are as follows:			
	disable: Connect without SSL			
	allow: Connect without SSL, and if it fails, connect with SSL			
	prefer: Connect with SSL, and if it fails, connect without SSL			
	require: Connect always with SSL			
	verify-ca: Connect with SSL, using a certificate issued by a trusted CA (*1)			
	verify-full: Connect with SSL, using a certificate issued by a trusted CA to verify if the server host name matches the certificate (*1)			
sslservercertcn(PGXSSLS ERVERCERTCN)	This parameter is enabled only to perform SSL authentication (sslmode=verifyfull).			
	Specify the server certificate CN. If this is omitted, the value will be null, and the server certificate CN will be authenticated using the host name specified in host.			

<sup>\*1:</sup> If specifying either "verify-ca" or "verify-full", use the system environment variable PGSSLROOTCERT (connection parameter sslrootcert) of your operating system to specify the CA certificate file as shown below.

Example)

Variable name: PGSSLROOTCERT Variable value: *cACertificateFile* 

#### When using URI

```
postgresql://host1[:port1],host2[:port2][,...]/database_name
[?target_session_attrs={read-write | prefer-read | any }]
```

- When using IPV6, specify the host in the "[host]" (with square brackets) format.

#### [Example]

```
postgresql://postgres@[2001:Db8::1234]:27500,192.168.1.1:27500/database_name
```

#### When using key-value

```
host=host1[,host2] port=port1[,port2] user=user1 password=pwd1 dbname=mydb
[target_session_attrs={read-write| prefer-read | any }]
```

- When using IPV6, specify the host in the "host" format.

#### [Example]

```
host=2001:Db8::1234,192.168.1.1 port=27500,27500
```



If using the connection parameter connect\_timeout, this value is applied for connections to each of the specified hosts. If both multiplexed database servers have failed, the connection will time out when a time equal to double the connect\_timeout value elapses.



#### Information

If using a password file (.pgpass), describe the entries matching each server.

- Example 1:

```
host1:port1:dbname:user:password
host2:port2:dbname:user:password
```

- Example 2:

```
*:port:dbname:user:password
```

## 10.2.6 Using Embedded SQL

It is recommended that you use a connection service file. Refer to "10.2.4 Using a Connection Service File" for details.



### Point

If using a connection service file, either of the following methods is available:

- Set the service name as a string literal or host variable, as follows: tcp:postgresql://?service=my\_service
- Set the service name in the environment variable PGSERVICE, and use CONNECT TO DEFAULT

If a connection service file will not be used, use a literal or variable to specify the connection destination server information for target in the SQL statement below:

```
EXEC SQL CONNECT TO target [AS connection-name] [USER user-name];
```

#### Method used

```
dbname@host1,host2[:[port1][,port2]]
tcp:postgresql://host1,host2[:[port1][,port2]] [/dbname] [?target_session_attrs={read-
write | prefer-read | any}][&sslmode=verify-
full&sslservercertcn=targetServerCertificateCN]
```

- The above format cannot be specified directly without using a literal or variable.

Table 10.6 Information to be set

Argument	Explanation			
host1 host2	Specify the IP address or host name. IPv6 format addresses cannot be specified.			
port1 port2	Specify the connection destination port numbers, using a comma as the delimiter.  The port number can be omitted. If omitted, the default is 27500.			
dbname	Specify the database name.			
target_session_attrs	Specify the selection sequence of the servers to which the application will connect. Refer to "Target server" for details.			
sslmode	Specify this to encrypt communications. By default, this is disabled. The setting values for sslmode are as follows:			
	disable: Connect without SSL			
	allow: Connect without SSL, and if it fails, connect with SSL			
	prefer: Connect with SSL, and if it fails, connect without SSL			
	require: Connect always with SSL			
	verify-ca: Connect with SSL, using a certificate issued by a trusted CA (*1)			
	verify-full: Connect with SSL, using a certificate issued by a trusted CA to verify if the server host name matches the certificate (*1)			
sslservercertcn	This parameter is enabled only to perform SSL authentication (sslmode=verify-full).			
	Specify the server certificate CN. If this is omitted, the value will be null, and the server certificate CN will be authenticated using the host name specified in host.			

<sup>\*1:</sup> If specifying either "verify-ca" or "verify-full", use the system environment variable PGSSLROOTCERT (connection parameter sslrootcert) of your operating system to specify the CA certificate file as shown below.

Example)

Variable name: PGSSLROOTCERT Variable value: *cACertificateFile* 



Environment variables can also be used. Refer to "10.2.5 Using the C Library (libpq)" for information on environment variables.



If using the connection parameter connect\_timeout, this value is applied for connections to each of the specified hosts. If both multiplexed database servers have failed, the connection will time out when a time equal to double the connect\_timeout value elapses.

## 10.2.7 Using the psql Command

It is recommended that you use a connection service file. Refer to "10.2.4 Using a Connection Service File" for details.

If a connection service file will not be used, specify the following information in the psql command option/environment variable.

Table 10.7 Information to be set

Option (environment variable)	Explanation		
-h/host(PGHOST/ PGHOSTADDR)	Specify IP address 1 and IP address 2, or the host name, using a comma as the delimiter.  This can also be specified for the environment variable PGHOST or PGHOSTADDR.		
-p/port(PGPORT)	Specify the connection destination port numbers, using a comma as the delimiter. This can also be specified for the environment variable PGPORT.		
	Specify the port number corresponding to the IP address specified for the nth -h option as the nth -p option.		
	The port number can be omitted. If omitted, the default is 27500.		
	If <i>n</i> -h options are specified, and m -p options are specified then there will be error reported. The only exception is where m=n or m=1. In case only one port is specified, then the same is applied for all the hosts.		
(PGTARGETSESSIONA TTRS)	Specify the selection sequence of the servers to which the application will connect.  Refer to "Target server" for details.		
(PGSSLMODE)	Specify this to encrypt communications. By default, this is disabled. The setting values for PGSSLMODE are as follows:		
	disable: Connect without SSL		
	allow: Connect without SSL, and if it fails, connect with SSL		
	prefer: Connect with SSL, and if it fails, connect without SSL		
	require: Connect always with SSL		
	verify-ca: Connect with SSL, using a certificate issued by a trusted CA (*1)		
	verify-full: Connect with SSL, using a certificate issued by a trusted CA to verify if the server host name matches the certificate (*1)		
(PGXSSLSERVERCERT CN)	This environment variable is enabled only to perform SSL authentication (PGSSLMODE=verify-full).		
	Specify the server certificate CN. If this is omitted, the value will be null, and the server certificate CN will be authenticated using the host name specified in host.		

<sup>\*1:</sup> If specifying either "verify-ca" or "verify-full", use the system environment variable PGSSLROOTCERT (connection parameter sslrootcert) of your operating system to specify the CA certificate file as shown below.

Example)

Variable name: PGSSLROOTCERT Variable value: *cACertificateFile* 



If using the connection parameter connect\_timeout, this value is applied for connections to each of the specified hosts. If both multiplexed database servers have failed, the connection will time out when a time equal to double the connect\_timeout value elapses.



Use the same method as for psql commands to specify connection destination server information for other client commands used to specify connection destinations.

## **Chapter 11 Performance Tuning**

This chapter explains how to tune application performance.

## 11.1 Enhanced Query Plan Stability

FUJITSU Enterprise Postgres estimates the cost of query plans based on SQL statements and database statistical information, and selects the least expensive query plan. However, like other databases, FUJITSU Enterprise Postgres does not necessarily select the most suitable query plan. For example, it may suddenly select unsuitable query plan due to changes in the data conditions.

In mission-critical systems, stable performance is more important than improved performance, and changes in query plans case to be avoided. In this situation, by stabilizing the SQL statement query plan so that it does not change, deterioration of the application performance is suppressed.

### 11.1.1 Optimizer Hints

This section explains the basic feature content of the optimizer hint (pg\_hint\_plan).

Refer to the open-source software webpage for information on pg\_hint\_plan.

In FUJITSU Enterprise Postgres, the optimizer hints can be specified in all application interfaces.

#### **Description**

You can specify a query plan in each SQL statement.

#### **List of Features**

The main query plans that can be specified using this feature are as follows:

- Query methods
- Join methods
- Join sequences

#### Query methods

Specify which method to use to query the specified table.

The main features are as follows:

- SegScan (tableName)
- BitMapScan (tableName [indexName ... ])
- IndexScan (tableName [indexName ... ])
- IndexOnlyScan (tableName [indexName ... ])



- If the specified index does not exist, or is not related to the search condition column specified in the WHERE clause, for example, SeqScan will be used.

- Even if IndexOnlyScan is specified, IndexScan may be used if it is necessary to access the table because a row was updated, for example.
- If multiple query methods were specified for the same table, the method specified last will be used.

#### Join methods

Specify the join method.

The main features are as follows:

- NestLoop (tableName tableName [tableName ... ])
- MergeJoin (tableName tableName [tableName ... ])
- HashJoin (tableName tableName [tableName ... ])



- These cannot be specified for view tables and subqueries.
- If multiple methods were specified for the same table combination, the method specified last will be used.

#### Join sequences

The tables will be joined in the specified table sequence.

Specify the information using the following method:

- Leading ((table table))

The method used to specify [table] is as follows:

table = tableName or ( table table )



If multiple sequences were specified for the same table combination, the sequence specified last will be used.

#### **Usage method**

The use of this feature is explained below.

#### Method used to define this feature

Define this feature by specifying the format (block comment) " /\*+ ... \*/".

- To specify hint clauses in each SELECT statement, for example when there are multiple SELECT statements in the SQL statement, define all hint clauses in the first block comment.



#### Specifying hint clauses for the emp table and the dept table

```
WITH /*+ IndexScan(emp emp_age_index) IndexScan(dept dept_deptno_index) */ age30
AS (SELECT * FROM emp WHERE age BETWEEN 30 AND 39)
SELECT * FROM age30, dept WHERE age30.deptno = dept.deptno;
```

- To specify separate hint clauses for the same object in the SQL statement, define aliases in each object, and then specify hint clauses for those aliases.

.....



Specifying separate hint clauses for the emp table

```
WITH /*+ SeqScan(ta) IndexScan(tb) */ over100
AS (SELECT empno FROM emp ta WHERE salary > 1000000)
SELECT * FROM emp tb, over100 WHERE tb.empno = over100.empno AND tb.age < 30
```

- When using embedded SQL in C, the locations in which the hint clause block comment is specified are restricted. Refer to "6.4.2 Compiling Applications" for details.

#### **Usage notes**

- If a hint clause was specified in multiple block comments in the SQL statement, the hint clause specified in the second block comment and thereafter will be ignored.
- If characters other than those listed below appear before the hint clause in the SQL statement, they will be invalid even for hint clause block comments.
  - Space, tab, line feed
  - Letter (uppercase and lowercase), number
  - Underscore, comma
  - Brackets ()

#### 11.1.2 Locked Statistics

This section explains the basic feature content for locked statistics (pg\_dbms\_stats).

Refer to the open-source software webpage for information on pg\_dbms\_stats.

#### **Description**

Locks the statistics.

By using this feature to lock the statistics for performance obtained in job load testing in an environment that simulates a production environment, performance degradation caused by changes to the query plan after go-live can be suppressed.

Additionally, by using the export and import features, statistics that were checked in the test environment can also be reproduced in the production environment.

#### **List of Features**

The main features that can be specified using this feature are as follows.

#### [Features]

Feature	Details	Description	
Lock/unlock of the statistics	Lock Locks the statistics so that the currently selected query p remains selected.		
	Unlock	Unlocks the statistics.	
Backup/restore of the	Backup	Backs up the current statistics.	
statistics	Restore	Restores the statistics to the point when they were backed up, and then locks them.	
	Purge	Deletes backups that are no longer necessary.	
Backup/restore using	Export	Outputs the current statistics to an external file (binary format).	
external files	Import	Reads the statistics from an external file created by the export feature, and then locks them.	

#### [Target object]

Target resource	Range of feature	
Database	In the database	
Schema	In the schema	
Table	In the table	
Column	ID column	

#### **Usage method**

The use of this feature is explained below.

#### Method used to specify this feature

Specify this feature as an SQL function.

The methods used to specify the main features are shown in the table below.

Feature	Object	Function specified	
Lock	Database	dbms_stats.lock_database_stats()	
	Schema	dbms_stats.lock_schema_stats('schemaName')	
	Table	dbms_stats.lock_table_stats('schemaName.tableName')	
Unlock	Database	dbms_stats.unlock_database_stats()	
	Schema	dbms_stats.unlock_schema_stats('schemaName')	
	Table	dbms_stats.unlock_table_stats('schemaName.tableName')	
Import	Database	dbms_stats.import_database_stats('fullPathOfExportedFile')	
Backup	Database	dbms_stats.backup_database_stats('commentUsedForIdentification' )	
Restore	Database	[Format 1] dbms_stats.restore_database_stats('timestamp')  [Timestamp] Specify in the same format as the time column of the backup_history table. Backups earlier than the specified time will be restored.  [Format 2] dbms_stats.restore_stats(backupId)	
		[Backup ID]  Specify a value in the id column of the backup_history table. The specified backup will be restored.	
Purge	Backup	dbms_stats.purge_stats(backupId,flagUsedForDeletion)	
		[Backup ID] Specify a value in the id column of the backup_history table.	
		[Flag used for deletion] true: The target backup is forcibly deleted. false: The target backup is deleted only when there are also backups for the entire database.	

Remark 1: The export feature is executed using the COPY statement, not the SQL function.



#### Example 1: Locking the statistics of the entire database

```
userdb=# SELECT dbms_stats.lock_database_stats();
  lock_database_stats
-----tbl1
  tbl1_pkey
```

Note that the locked information can be referenced as follows:

```
userdb=# select relname from dbms_stats.relation_stats_locked;
  relname
-----tbl1
tbl1_pkey
```

#### Example 2: Unlocking the statistics of the entire database

```
userdb=# SELECT dbms_stats.unlock_database_stats();
  unlock_database_stats
------
tbl1
tbl1_pkey
```

#### Example 3: Backing up the statistics of the entire database

Note that the backed up statistics can be referenced as follows:

The ID:1 backup "backup1" is obtained for each database at "2014-03-04 11:08:40.315948+09". [Meaning of unit] d: database s: schema t: table c: column

#### Example 4: Exporting the statistics of the entire database

```
$ psql -d userdb -f export.sql
BEGIN
COMMIT
```

export.sql is the file in which the COPY statement is defined.

Refer to "export\_effective\_stats-<x>.sql\_sample" for information on the content of the COPY statement. "<x>" indicates the product version.

"export\_effective\_stats-<*x*>.sql\_sample" is stored as follows: *fujitsuEnterprisePostgresInstallDir*/share/doc/extension

#### Example 5: Importing the statistics of the entire database

#### **Usage notes**

- You must run the ANALYZE command once for the target tables of this feature. If the ANALYZE command is not run, the statistics cannot be locked.
  - Refer to "SQL Commands" in "Reference" in the PostgreSQL Documentation for information on the ANALYZE command.
- To use this feature to delete an object that has locked the statistics, use the unlock feature to delete the object lock information first.
- This feature does not specify the statistics value directly. It reproduces the status that has actually occurred. For this reason, if the text format is specified in the COPY statement when the export occurs, restore will not be possible. Always use the binary format when performing the export.

# Chapter 12 Scan Using a Vertical Clustered Index (VCI)

This chapter describes scanning using a VCI.

12		
	No	ote

This feature can only be used in Advanced Edition.

## 12.1 Operating Conditions

Faster aggregation can be achieved by using a VCI defined for all columns to be referenced.

This section describes the conditions under which a scan can use a VCI.

Whether to use VCI is determined based on cost estimation in the same way as normal indexes. Therefore, another execution plan will be selected if it is cheaper than a VCI even if a VCI is available.

#### SQL statements that can use VCIs

In addition to general SELECT statements, VCIs can be used for the SQL statements below (as long as they do not specify any of the elements listed in "SQL statements that cannot use VCIs" below):

- SELECT INTO
- CREATE TABLE AS SELECT
- CREATE MATERIALIZED VIEW ... AS SELECT
- CREATE VIEW ... AS SELECT
- COPY (SELECT ...) TO

#### SQL statements that cannot use VCIs

VCIs cannot be used for SQL statements that specify any of the following:

- Subquery to reference the column in which the parent query is referencing is specified
- Lock clause (such as FOR UPDATE)
- Cursor declared with WITH HOLD or scrollable
- SERIALIZABLE transaction isolation level
- Function or operator listed in "Functions and operators that do not use a VCI"
- User-defined function

Table 12.1 Functions and operators that cannot use VCIs

Classification		Function/operator	
Mathematical functions and operators	Random functions	random and setseed	
String functions and operators	String functions	format (if the <i>format</i> argument is specified), regexp_matches, regexp_split_to_array and regexp_split_to_table	
Date/time functions and operators	Date/time functions	age(timestamp), current_date, current_time, current_timestamp, localtime, localtimestamp, statement_timestamp and transaction_timestamp	
	Delaying execution functions	pg_sleep, pg_sleep_for, and pg_sleep_until	

Classification		Function/operator	
Enum support functions		All functions and operators	
Geometric functions and	operators	All functions and operators	
Network address function	ns and operators	All functions and operators	
Text search functions and	l operators	All functions and operators	
XML functions		All functions	
JSON functions and oper	ators	All functions and operators	
Sequence manipulation for	unctions	All functions	
Array functions and oper	ators	All functions and operators	
Range functions and open	rators	All functions and operators	
Aggregate functions	General-purpose aggregate functions	array_agg, json_agg, json_object_agg, string_agg and xmlagg	
	Aggregate functions for statistics	corr, covar_pop, covar_samp, regr_avgx, regr_avgy, regr_count, regr_intercept, regr_r2, regr_slope, regr_sxx, regr_sxy and regr_syy	
	Ordered-set aggregate functions	All functions	
	Hypothetical-set aggregate functions	All functions	
Window functions		All functions	
Subquery expressions		Subquery expressions with its row constructor specified on the left side	
Row and array comparisons		Row constructor and composite type comparisons	
Set returning functions		All functions	
System information functions		All functions	
System administration fu	nctions	All functions	
Trigger functions		All functions	
Session information func	tions	current_role and current_user	

## **12.2 Usage**

This section describes how to use a VCI in line with the following steps:



### 12.2.1 Designing

Design as follows before using a VCI.

- Execution multiplicity and number of parallel processes
- Parameters

Execution multiplicity and number of parallel processes

Determine the maximum number of SQL statements that can be executed simultaneously and the number of parallel processes based on the number of CPU cores that can be allocated for scans that use VCI to perform aggregate processing. Design in advance the multiplicity of SQL statements for executing scans that use VCI and the number of parallel processes for scans that use VCI.

For example, if the number of CPUs that can be allocated is 32 cores, then the maximum number of SQL statements that can be executed simultaneously is 8 and the number of parallel processes is 4.



A temporary file is created in /dev/shm or in a directory specified for the vci.smc\_directory parameter as the dynamic shared memory for each SQL statement during a scan using a VCI.

A temporary file is created in a directory under the data storage directory or in a directory specified for the vci.smc\_directory parameter as the dynamic shared memory for each SQL statement during a scan using a VCI.

Ensure that this directory has sufficient space to meet the memory requirements estimated for the execution multiplicity and number of parallel processes of SQL statements (refer to "Memory used per scanning" in "VCI Memory Requirements" in the Installation and Setup Guide for Server for details). If it does not have sufficient space when a scan is performed, SQL statements will return errors due to the insufficient memory.

#### **Parameters**

The VCI parallel scan feature cannot be used for setting parameters immediately after creating an instance.

Therefore, set the parameters below based on the values determined in "Execution multiplicity and number of parallel processes of SQL statements" above.

Parameter name	Description	Default	Value index
vci.max_parallel_degre e	Maximum number of VCI parallel	0	Specify the number of parallel processes.



Parameter name	Description	Default	Value index
	(background processes) to be used per SQL statement.		
vci.smc_directory	Directory name in which a temporary file is created as the dynamic shared memory during a scan using a VCI.	L /dev/shm W A directory (dataStorageDir\base\\pgsql_tmp) under the data storage directory	Specify a directory that has enough free space for the memory used for each query during the scan.
max_worker_processes	Maximum number of background processes that the system supports.	8	Add the value of the maximum number of SQL statements that can be executed simultaneously for scans that use VCI multiplied by vci.max_parallel_degree.



Refer to "Parameters" in the Operation Guide for information on the details of and how to set the parameters.

.....

### 12.2.2 Checking

Execute the SQL statement with "EXPLAIN ANALYZE" to check the following:

- If a VCI was used "Custom Scan (VCI...)" is displayed in the plan if a VCI was used.
- Number of parallel processes
   The number of parallel processes when the SQL statement is executed is displayed in "Allocated Workers". Check that
  it is running the designed number of parallel processes.
- Response
   Check if the execution time displayed in "Execution time" is as estimated.

The following shows an example of the output result of EXPLAIN ANALYZE:



A cost output by the execution plan that uses a VCI may be inaccurate. A VCI works if all or part of the best execution plan when the SQL statement was executed is replaced with an execution plan that uses a VCI. If the cost of the execution plan to be replaced is lower than a certain value (vci.cost\_threshold parameter), it will not be replaced or recalculated. Therefore, the cost of the original execution plan is output as is.

### 12.2.3 Evaluating

If the results in "12.2.2 Checking" is any of the following, tune accordingly:

If a VCI is not used

- Check if the "12.1 Operating Conditions" are met.
- Check if vci.enable is set to "on".
- A VCI may not be appropriately used when statistics are outdated, such as immediately after inserting a large amount of data. In such cases, execute the VACUUM ANALYZE statement or the ANALYZE statement.
- A VCI is not used if there is insufficient memory for VCI scan. This may occur during time-consuming transactions involving tables for which VCIs were defined. Set vci.log\_query to "on", and check if either "could not use VCI: local ROS size (%zu) exceeds limit (%zu)" or "out of memory during local ROS generation" is output. If it is, then increase the value of the vci.max\_local\_ros.

#### Response is not as expected

Tuning may improve response. Check the following:

- If vci.max\_parallel\_degree is not set or is set to 0, set an appropriate value according to "12.2.1 Designing".
- If there is a margin in the CPU usage, increase the value of vci.max\_parallel\_degree and check again. In addition, if the value that of max\_worker\_processes is lower than the maximum number of SQL statements that can be executed simultaneously for parallel scan multiplied by vci.max\_parallel\_degree, increase it and check again.

## 12.3 Usage Notes

This section provides notes on using VCI.

- Regardless of whether VCI is used, the content of the result does not change. However, records may be returned in a different order if the ORDER BY clause is not specified.
- To reduce resource consumption, edit postgresql.conf or use the SET statement to enable/disable vci.enable when you use this feature only for specific times or jobs (SQL applications).
- The optimizer hint (pg hint plan) cannot be specified for a VCI. The hint clause is ignored if it is specified.
- If a plan other than VCI is specified for the optimizer hint (pg\_hint\_plan), a VCI may be used. Therefore, if you specify a query plan with the hint clause, use the SET statement to set vci.enable to "off".
- The message below may be output when a scan that uses VCI is performed on the streaming replication standby server:

```
"LOG: recovery has paused"
"HINT: Execute pg_wal_replay_resume() to continue."
```

This message is output because application of the WAL to the VCI temporarily pauses due to the scan being performed.

- Even if a scan is performed using a VCI, information in the idx\_scan, idx\_tup\_read, and idx\_tup\_fetch columns of the collected statistics views, pg\_stat\_all\_indexes and pg\_stat\_user\_indexes, will not be updated.

- Currently, it is not possible to replace the query plan for parallel aggregation with the query plan using VCI. Therefore, if you create a VCI on a column of a partition table and aggregate (sum () etc.) on that column, one of the following plans will be selected. Use different setting parameters according to the situation of the target table.
  - Plan of the parallel aggregations using scan methods other than VCI scan

It is selected when max\_parallel\_workers\_per\_gather is 1 or more.

```
explain select sum(value) from test;

QUERY PLAN

Finalize Aggregate (cost=99906.30..99906.31 rows=1 width=8)

-> Gather (cost=99906.08..99906.29 rows=2 width=8)

Workers Planned: 2

-> Partial Aggregate (cost=98906.08..98906.09 rows=1 width=8)

-> Parallel Append (cost=0.00..94739.83 rows=1666500 width=4)

-> Parallel Seq Scan on test_1 (cost=0.00..43203.67 rows=833250 width=4)

-> Parallel Seq Scan on test_2 (cost=0.00..43203.67 rows=833250 width=4)
```

This plan is fast when the number of records to be aggregated (number of records that hit the search conditions) is very large. This is because the benefit of parallelizing aggregation is important, not the performance of scanning. For example, each parallel worker will perform a sequential scan and aggregate most of the scanned records.

- Plan that aggregates VCI scan results by a single multiplex

It is selected by setting max\_parallel\_workers\_per\_gather to 0 and not creating a query plan of parallel aggregate.

```
explain select sum(value) from test;

QUERY PLAN

Aggregate (cost=145571.00..145571.01 rows=1 width=8)

-> Append (cost=0.00..135572.00 rows=3999600 width=4)

-> Custom Scan (VCI Scan) using test_1_id_value_idx on test_1

(cost=0.00..57787.00 rows=1999800 width=4)

Allocated Workers: 2

-> Custom Scan (VCI Scan) using test_2_id_value_idx on test_2

(cost=0.00..57787.00 rows=1999800 width=4)

Allocated Workers: 2
```

This plan is fast when the number of aggregated items is not large or when the size of the aggregated column is smaller than the record size. This is because the scan performance is more important, so it is faster to aggregate the results of VCI scans of each partition.

- Originally, if there is only one partition to be accessed, the following VCI aggregation plan can be used. Below is an example of scanning only one partition with partition pruning.

```
explain select sum(value) from test where id < 1000001;

QUERY PLAN

Custom Scan (VCI Aggregate) (cost=62786.50..62786.51 rows=1 width=8)

Allocated Workers: 2

-> Custom Scan (VCI Scan) using test_l_id_value_idx on test_1 (cost=0.00..57787.00 rows=1999800 width=4)

Filter: (id < 1000001)
```

However, the current planner does not try to choose VCI aggregation because it creates a plan for parallel aggregation if the table is partitioned. So in this case, set max\_parallel\_workers\_per\_gather to 0 to force the planner to choose VCI aggregation.

## Appendix A Precautions when Developing Applications

This appendix describes precautions when developing applications with FUJITSU Enterprise Postgres.

## A.1 Precautions when Using Functions and Operators

This section describes notes for using functions and operators.

### A.1.1 General rules of Functions and Operators

This section describes general rules for using functions and operators. Ensure the general rules are followed when using functions and operators to develop applications.

#### General rules

- Specify the stated numbers for arguments when specifying numbers for arguments in functions.
- Specify the stated data types when specifying data types for functions. If you use a data type other than the stated data types, use CAST to explicitly convert the data type.
- Specify data types that can be compared when specifying data types for operators. If you use a data type that cannot be compared, use CAST to explicitly convert the data type.



Refer to "Functions and Operators" under "The SQL Language" in the PostgreSQL Documentation for information on the functions and operators available with FUJITSU Enterprise Postgres.

## A.1.2 Errors when Developing Applications that Use Functions and/ or Operators

This section provides examples of problems that may occur when developing applications that use functions and/or operators, and describes how to deal with them.

The error "Function \*\*\*\*\* does not exist" occurs when executing SQL

The following error will occur when executing an SQL statement that does not abide by the general rules for functions:

```
ERROR: Function ***** does not exist
```

Note: "\*\*\*\*\*" denotes the function for which the error occurred, and the data type of its arguments.

The cause of the error will be one of the following:

- The specified function does not exist.
- The wrong number of arguments or wrong argument data type was specified

#### Corrective action

Check the following points and correct any errors:

- Check if there are any errors in the specified function name, number of arguments, or argument data type, and revise accordingly.
- Check the argument data type of the function displayed in the message. If an unintended data type is displayed, use a function such as CAST to convert it.

The error "Operator does not exist" occurs when executing SQL

The following error will occur when executing an SQL statement that specifies a data type in the operator that cannot be compared:

```
ERROR: Operator does not exist: *****
```

Note: "\*\*\*\*\*" denotes the operator for which the error occurred, and the data type of the specified value.

#### Corrective action

Ensure the data type of the expressions specified on the left and right sides of the operator can be compared. If required, revise to ensure these data types can be compared by using a function such as CAST to explicitly convert them.

## A.2 Notes when Using Temporary Tables

In standard SQL, a temporary table can be defined in advance to enable an empty temporary table to be created automatically when the application connects to the database. However, in FUJITSU Enterprise Postgres, a temporary table must be created when the application connects to the database by explicitly using the CREATE TABLE statement.

If the same temporary table is repeatedly created and deleted during the same session, the system table might expand, and memory usage might increase. To prevent this, specify the CREATE TABLE statement to ensure the temporary table is reused.

For example, in cases where a temporary table would be created and deleted for repeatedly executed transactions, specify the CREATE TABLE statement as shown below:

- Specify "IF NOT EXISTS" to create a temporary table only if none exists when the transaction starts.
- Specify "ON COMMIT DELETE ROWS" to ensure all rows are deleted when the transaction ends.



Refer to "SQL Commands" under "Reference" in the PostgreSQL Documentation for information on the CREATE TABLE statement.

Examples of SQL using a temporary table are shown below:

Example of bad use (creating and deleting a temporary table)

```
BEGIN;
CREATE TEMPORARY TABLE mytable(coll CHAR(4), col2 INTEGER) ON COMMIT DROP;
(mytable processes)

COMMIT;
```

#### Example of good use (reusing a temporary table)

```
BEGIN;
CREATE TEMPORARY TABLE IF NOT EXISTS mytable(coll CHAR(4), col2 INTEGER) ON COMMIT
DELETE ROWS;
(mytable processes)

COMMIT;
```

## A.3 Implicit Data Type Conversions

An implicit data type conversion refers to a data type conversion performed automatically by FUJITSU Enterprise Postgres, without the need to explicitly specify the data type to convert to.

The combination of possible data type conversions differs, depending on whether the expression in the conversion source is a literal.

For non-literals, data types can only be converted to other types within the same range.

For literals, character string literal types can be converted to the target data type. Numeric literals are implicitly converted to specific numeric types. These implicitly converted numeric literals can then have their types converted to match the conversion target data type within the numeric type range. For bit character string literals, only the bit column data type can be specified. The following shows the range of type conversions for literals.

Table A.1 Data type combinations that contain literals and can be converted implicitly

Conversion target		Conversion source		
		Characte r literal (*1)	Numeric literal(*2)	Bit character string literal
Numeric type	SMALLINT	Y	N	N
	INTEGER	Y	Y (*3)	N
	BIGINT	Y	Y (*4)	N
	DECIMAL	Y	Y (*5)	N
	NUMERIC	Y	Y (*5)	N
	REAL	Y	N	N
	DOUBLE PRECISION	Y	N	N
	SMALLSERIAL	Y	N	N
	SERIAL	Y	Y (*3)	N
	BIGSERIAL	Y	Y (*4)	N
Currency type	MONEY	Y	N	N
Character type	CHAR	Y	N	N
	VARCHAR	Y	N	N
	NCHAR	Y	N	N
	NCHAR VARYING	Y	N	N
	TEXT	Y	N	N
Binary data type	BYTEA	Y	N	N
Date/time type	TIMESTAMP WITHOUT TIME ZONE	Y	N	N
	TIMESTAMP WITH TIME ZONE	Y	N	N
	DATE	Y	N	N
	TIME WITHOUT TIME ZONE	Y	N	N
	TIME WITH TIME ZONE	Y	N	N
	INTERVAL	Y	N	N
Boolean type	BOOLEAN	Y	N	N
Geometric type	POINT	Y	N	N
	LSEG	Y	N	N
	BOX	Y	N	N

Conversion target		С	Conversion source		
		Characte r literal (*1)	Numeric literal(*2)	Bit character string literal	
	PATH	Y	N	N	
	POLYGON	Y	N	N	
	CIRCLE	Y	N	N	
Network address type	CIDR	Y	N	N	
	INET	Y	N	N	
	MACADDR	Y	N	N	
	MACADDR8	Y	N	N	
Bit string type	BIT	Y	N	Y	
	BIT VARYING	Y	N	Y	
Text search type	TSVECTOR	Y	N	N	
	TSQUERY	Y	N	N	
UUID type	UUID	Y	N	N	
XML type	XML	Y	N	N	
JSON type	JSON	Y	N	N	

Y: Can be converted

Implicit data type conversions can be used when comparing or storing data.

The conversion rules differ, depending on the reason for converting. Purpose-specific explanations are provided below.

## A.3.1 Function Argument

Value expressions specified in a function argument will be converted to the data type of that function argument.



Refer to "Functions and Operators" under "The SQL Language" in the PostgreSQL Documentation for information on data types that can be specified in function arguments.

### A.3.2 Operators

Comparison operators, BETWEEN, IN

Combinations of data types that can be compared using comparison operators, BETWEEN, or IN are shown below.

N: Cannot be converted

<sup>\*1:</sup> Only strings that can be converted to the data type of the conversion target can be specified (such as "1" if the conversion target is a numeric type)

<sup>\*2: &</sup>quot;Y" indicates specific numeric types that are converted first.

<sup>\*3:</sup> Integers that can be expressed as INTEGER types can be specified

<sup>\*4:</sup> Integers that cannot be expressed as INTEGER types, but can be expressed as BIGINT types, can be specified

<sup>\*5:</sup> Integers that cannot be expressed as INTEGER or BIGINT types, but that can be expressed as NUMERIC types, or numeric literals that contain a decimal point or the exponent symbol (e), can be specified

Table A.2 Combinations of comparable data type

Left side	Right side		
	Numeric type	Charact er string type	Date/time type
Numeric type	Y	N	N
Character type	N	Y	N
Date/time type	N	N	Y

Y: Can be compared

N: Cannot be compared

When strings with different lengths are compared, the shorter one is padded with spaces to make the lengths match.

When numeric values with different precisions are compared, data will be converted to the type with the higher precision.

Set operation and CASE also follow the same rules.

#### Other operators

Value expressions specified in operators will be converted to data types that are valid for that operator.



Refer to "Functions and Operators" under "The SQL Language" in the PostgreSQL Documentation for information on data types that can be specified in operators.

### A.3.3 Storing Values

Value expressions specified in the VALUES clause of the INSERT statement or the SET clause of the UPDATE statement will be converted to the data type of the column in which they will be stored.

## A.4 Notes on Using Index

This section explains the notes on using the following indexes:

- SP-GiST index

#### A.4.1 SP-GiST Index

If more than 2 concurrent updates are performed on a table in which the SP-GiST index is defined, applications may stop responding. When this occur, all system processes including the Check Pointer process will also be in the state of no response. For these reasons, use of the SP-GiST index is not recommended.

## A.5 Notes on Using Multibyte Characters in Definition Names

Do not use multibyte characters in database names or user names if using a Windows database server.

Multibyte characters must not be used in database names or user names on non-Windows database servers, because certain conditions may apply or it may not be possible to connect to some clients.

Related notes and constraints are described below.

#### 1) Configuring the client encoding system

The client encoding system must be configured when the names are created.



Refer to "Character Set Support" in "Server Administration" in the PostgreSQL Documentation for information on how to configure the client encoding system.

#### 2) Encoding system of names used for connection

Ensure that the encoding system of names used for connection is the same as that of the database that was connected when these names were created.

The reasons for this are as follows:

- Storage system for names in FUJITSU Enterprise Postgres

The system catalog saves encoded names by using the encoding system of the database at the time the names were created.

- Encoding conversion policy when connected

When connected, names sent from the client are matched with names in the system catalog without performing encoding conversion.

Accordingly, if the database that was connected when the names were defined uses the EUC\_JP encoding system, but the database name is specified using UTF-8 encoding, then the database will be considered to be non-existent.

#### 3) Connection constraints

The table below shows the connection constraints for each client type, based on the following assumptions:

- The conditions described in 1) and 2) above are satisfied.
- The database name and user names use the same encoding system.

Client type	Client operating system		
Client type	Windows(R)	Linux	
JDBC driver	Cannot be connected	Cannot be connected	
ODBC driver	Cannot be connected	No connection constraints	
.NET Data Provider	Can only connect when the encoding system used for definitions is UTF-8	-	
SQLEmbedded SQL in C	Can only connect when the connection service file (pg_service.conf) is used	No connection constraints	
psql command	Can only connect when the connection service file (pg_service.conf) is used	No connection constraints	

# Appendix B Conversion Procedures Required due to Differences from Oracle Database

This appendix explains how to convert from an Oracle database to FUJITSU Enterprise Postgres, within the scope noted in "Chapter 9 Compatibility with Oracle Databases" from the following perspectives:

- Feature differences
- Specification differences

This document assumes that the version of the Oracle database to be converted is 7-10.2g.

## **B.1 Outer Join Operator (Perform Outer Join)**

#### **Features**

In the WHERE clause conditional expression, by adding the plus sign (+), which is the outer join operator, to the column of the table you want to add as a table join, it is possible to achieve an outer join that is the same as a joined table (OUTER JOIN).

### **B.1.1 Comparing with the ^= Comparison Operator**

#### Oracle database

```
SELECT *
FROM t1, t2
WHERE t1.col1(+) ^= t2.col1;
```

#### **FUJITSU Enterprise Postgres**

```
SELECT *
FROM t1, t2
WHERE t1.col1(+) != t2.col1;
```

#### Feature differences

Oracle database

The ^= comparison operator can be specified.

#### **FUJITSU Enterprise Postgres**

The ^= comparison operator cannot be specified.

#### Conversion procedure

Convert using the following procedure:

- 1. Locate the places where the keyword "^=" is used.
- 2. Ensure that the keyword, "(+)", is either on the right or left-hand side.
- 3. Change "^=" to "!=".

<sup>\*</sup> col1 is assumed to be CHAR(4) type

<sup>\*</sup> col1 is assumed to be CHAR(4) type

# **B.2 DECODE (Compare Values and Return Corresponding Results)**

#### **Features**

DECODE compares values of the conversion target value expression and the search values one by one, and if the values of the conversion target value expression and the search values match, a corresponding result value is returned.

## **B.2.1 Comparing Numeric Data of Character String Types and Numeric Characters**

#### **Oracle database**

#### **FUJITSU Enterprise Postgres**

```
SELECT DECODE( CAST(col1 AS INTEGER),

1000, 'ITEM-A',

2000, 'ITEM-B',

'ITEM-C')

FROM t1;
```

#### Feature differences

#### Oracle database

When the value expression is a string and the search value is a numeric, the string value will be converted to the data type of the comparison target numeric, so that they can be compared.

#### **FUJITSU Enterprise Postgres**

If the conversion target value expression is a string value, then no search value can be specified with numbers.

#### Conversion procedure

Since the data type that can be specified for the conversion target value expression is unknown, use CAST to explicitly convert the conversion target value expression (col1 in the example) to a numeric (INTEGER type in the example).

## **B.2.2 Obtaining Comparison Result from more than 50 Conditional Expressions**

#### Oracle database

```
SELECT DECODE(col1,

1,'A',

2,'B',

...

78,'BZ',

NULL,'UNKNOWN',

'OTHER')

FROM t1;
```

<sup>\*</sup> col1 is assumed to be CHAR(4) type

<sup>\*</sup> col1 is assumed to be CHAR(4) type

#### **FUJITSU Enterprise Postgres**

```
SELECT CASE

WHEN col1 = 1 THEN 'A'
WHEN col1 = 2 THEN 'B'
...
WHEN col1 = 78 THEN 'BZ'
WHEN col1 IS NULL THEN 'UNKNOWN'
ELSE 'OTHER'
END
FROM t1;
```

#### Feature differences

#### Oracle database

Search value with a maximum of 127 items (up to 255 arguments in total) can be specified.

#### **FUJITSU Enterprise Postgres**

Search value with a maximum of 49 items (up to 100 arguments in total) only can be specified.

#### Conversion procedure

Convert to the CASE expression using the following procedure:

- 1. Specify the DECODE conversion target value expression (col1 in the first argument, in the example) and the search value (1 in the second argument, in the example) for the CASE expression search condition. Specify the DECODE result value ('A' in the third argument, in the example) for the CASE expression THEN (WHEN col1 = 1 THEN 'A', in the example). Note that if the search value is NULL, specify "IS NULL" for the search condition for the CASE expression.
- 2. If the DECODE default value ('OTHER' in the last argument, in the example) is specified, specify the default value for the CASE expression ELSE (ELSE 'OTHER', in the example).

## **B.2.3 Obtaining Comparison Result from Values with Different Data**Types

#### Oracle database

#### **FUJITSU Enterprise Postgres**

<sup>\*</sup> col1 is assumed to be INTEGER type

<sup>\*</sup> col1 is assumed to be CHAR(4) type

<sup>\*</sup> col1 is assumed to be CHAR(4) type

#### Feature differences

#### Oracle database

The data types of all result values are converted to the data type of the first result value.

#### **FUJITSU Enterprise Postgres**

Results in an error.

#### Conversion procedure

Convert using the following procedure:

- 1. Check the literal data type for the first result value specified.
- 2. Change the literals specified for each result value to the literal data type checked in the step 1.

# B.3 SUBSTR (Extract a String of the Specified Length from Another String)

#### Features

SUBSTR returns the number of characters specified in the third argument (starting from the position specified in the second argument) from the string specified in the first argument.

Refer to "9.2.1 Notes on SUBSTR" for details on precautions when using SUBSTR.

## B.3.1 Specifying a Value Expression with a Data Type Different from the One that can be Specified for Function Arguments

#### **Oracle database**

#### **FUJITSU Enterprise Postgres**

#### Feature differences

#### Oracle database

If the type can be converted to a data type that can be specified for function arguments, conversion is performed implicitly.

#### **FUJITSU Enterprise Postgres**

If the data types are different from each other, or if loss of significance occurs, implicit conversion is not performed.

<sup>\*</sup> col1 and col2 are assumed to be CHAR type

<sup>\*</sup> col1 and col2 are assumed to be CHAR type

#### Conversion procedure

Since the data type of the string length is clear, first execute the following CREATE CAST only once so that the CHAR type value (col2 in the example) specified for the string length is implicitly converted to INTEGER type.

```
CREATE CAST (CHAR AS INTEGER) WITH INOUT AS IMPLICIT;
```

## B.3.2 Extracting a String with the Specified Format from a Datetime Type Value

#### Oracle database

```
SELECT SUBSTR( CURRENT_TIMESTAMP,

1,

8)

FROM DUAL;
```

#### **FUJITSU Enterprise Postgres**

#### Feature differences

#### Oracle database

A datetime value such as CURRENT\_TIMESTAMP can be specified for character value expressions.

#### **FUJITSU Enterprise Postgres**

A datetime value such as CURRENT\_TIMESTAMP cannot be specified for character value expressions.

#### Conversion procedure

First, specify TO\_CHAR for the SUBSTR character value expression.

Specify datetime type (CURRENT\_TIMESTAMP, in the example) in firstArg of TO\_CHAR, and specify the format template pattern ('DD-MON-YY HH.MI.SS.US PM', in the example) for secondArg to match with the result of SUBSTR before conversion.

 $TO\_CHAR\ specification\ format:\ TO\_CHAR(\textit{firstArg, secondArg})$ 



Refer to "Data Type Formatting Functions" in the PostgreSQL Documentation for information on format template patterns that can be specified for TO\_CHAR in FUJITSU Enterprise Postgres.

## **B.3.3 Concatenating a String Value with a NULL value**

#### **Oracle database**

<sup>\*</sup> col1 and col2 are assumed to be character string type, and col2 may contain NULL

#### **FUJITSU Enterprise Postgres**

```
SELECT SUBSTR( col1 || NVL(col2, '')
2,
5)
FROM t1;
```

#### Feature differences

#### Oracle database

NULL is handled as an empty string, and strings are joined.

#### **FUJITSU Enterprise Postgres**

NULL is not handled as an empty string, and the result of joining the strings becomes NULL.

#### Conversion procedure

Convert using the following procedure:

- 1. Locate the places where the keyword "||" is used.
- 2. Check if any of the value expressions can contain NULL if they can, then execute step 3.
- 3. Modify to NVL(valExpr,").

## **B.4 NVL (Replace NULL)**

#### **Features**

NVL converts NULL values.

## **B.4.1 Obtaining Result from Arguments with Different Data Types**

#### **Oracle database**

```
SELECT NVL( col1, col2)
FROM t1;
```

#### **FUJITSU Enterprise Postgres**

```
SELECT NVL( col1,

CAST(col2 AS VARCHAR(100)))

FROM t1;
```

#### Feature differences

#### Oracle database

Value expressions with different data types can be specified. If the first argument is a string value, then VARCHAR2 is returned, and if it is a numeric, then a numeric type with greater range is returned.

#### **FUJITSU Enterprise Postgres**

Value expressions with different data types cannot be specified.

<sup>\*</sup> col1 and col2 are assumed to be character string type, and col2 may contain NULL

 $<sup>\</sup>ast$  col1 is assumed to be VARCHAR(100) type, and col2 is assumed to be CHAR(100) type

<sup>\*</sup> col1 is assumed to be VARCHAR(100) type, and col2 is assumed to be CHAR(100) type

#### Conversion procedure

Since the data types that can be specified for the expressions in the two arguments are unknown, use the following steps to convert:

- 1. Check the data types specified for each of the two expressions.
- 2. Using the data type that is to be received as a result, explicitly convert the other argument with CAST.

## B.4.2 Operating on Datetime/Numeric, Including Adding Number of Days to a Particular Day

#### Oracle database

```
SELECT NVL( col1 + 10, CURRENT_DATE)
FROM t1;
```

#### **FUJITSU Enterprise Postgres**

```
SELECT NVL( CAST(coll AS DATE) + 10, CURRENT_DATE)
FROM t1;
```

#### Feature differences

#### Oracle database

Numbers can be operated (added to or subtracted from) with either TIMESTAMP WITHOUT TIME ZONE type or TIMESTAMP WITH TIME ZONE type. Operation result will be DATE type.

#### FUJITSU Enterprise Postgres

Numbers cannot be operated (added to or subtracted from) with neither TIMESTAMP WITHOUT TIME ZONE type nor TIMESTAMP WITH TIME ZONE type. However, numbers can be operated (added to or subtracted from) with DATE type.

#### Conversion procedure

Convert using the following procedure:

- 1. Search locations where the keyword "+" or "-" is used in addition or subtraction, and check if these operations are between numbers and TIMESTAMP WITHOUT TIME ZONE type or TIMESTAMP WITH TIME ZONE type.
- 2. If they are, use CAST to explicitly convert TIMESTAMP WITHOUT TIME ZONE type or TIMESTAMP WITH TIME ZONE type to DATE type.

## B.4.3 Calculating INTERVAL Values, Including Adding Periods to a Date

#### Oracle database

```
SELECT NVL( CURRENT_DATE + (col1 * 1.5), col2)
FROM tl;
```

#### **FUJITSU Enterprise Postgres**

```
SELECT NVL( CURRENT_DATE +
CAST(col1 * 1.5 AS
```

<sup>\*</sup> col1 is assumed to be TIMESTAMP WITHOUT TIME ZONE type or TIMESTAMP WITH TIME ZONE type

<sup>\*</sup> col1 is assumed to be TIMESTAMP WITHOUT TIME ZONE type or TIMESTAMP WITH TIME ZONE type

<sup>\*</sup> col1 and col2 are assumed to be INTERVAL YEAR TO MONTH types

```
INTERVAL YEAR TO MONTH), col2)
FROM t1;
```

\* col1 and col2 are assumed to be INTERVAL YEAR TO MONTH types

#### Feature differences

#### Oracle database

INTERVAL YEAR TO MONTH type multiplication and division result in INTERVAL YEAR TO MONTH type and any fraction (number of days) will be truncated.

#### **FUJITSU Enterprise Postgres**

INTERVAL YEAR TO MONTH type multiplication and division result in INTERVAL type and fractions (number of days) will not be truncated.

#### Conversion procedure

Convert using the following procedure:

- 1. Search locations where the keywords "\*" or "/" are used in multiplication or division, and check if the specified value is INTERVAL YEAR TO MONTH type.
- 2. If the value is INTERVAL YEAR TO MONTH type, use CAST to explicitly convert the operation result to INTERVAL YEAR TO MONTH type.

## B.5 DBMS\_OUTPUT (Output Messages)

#### **Features**

DBMS\_OUTPUT sends messages to clients such as psql from PL/pgSQL.

### **B.5.1 Outputting Messages Such As Process Progress Status**

#### **Oracle database**

```
set serveroutput on;...(1)
DECLARE
 v_col1
               CHAR(20);
               INTEGER;
 v_col2
 CURSOR c1 IS
   SELECT col1, col2 FROM t1;
BEGIN
 DBMS_OUTPUT.PUT_LINE('-- BATCH_001 Start --');
 OPEN c1;
 DBMS_OUTPUT.PUT_LINE('-- LOOP Start --');
   FETCH c1 INTO v_col1, v_col2;
   EXIT WHEN c1%NOTFOUND;
   DBMS OUTPUT.PUT('.');
 END LOOP;
 DBMS_OUTPUT.NEW_LINE; ...(2)
 DBMS_OUTPUT.PUT_LINE('-- LOOP End --');
 CLOSE c1;
 DBMS_OUTPUT.PUT_LINE('-- BATCH_001 End --');
EXCEPTION
 WHEN OTHERS THEN
    DBMS_OUTPUT.PUT_LINE('-- SQL Error --');
    DBMS_OUTPUT.PUT_LINE('ERROR : ' || SQLERRM );
```

```
END;
/
```

#### **FUJITSU Enterprise Postgres**

```
DO $$
DECLARE
   v_col1
               CHAR(20);
   v_col2
               INTEGER;
   c1 CURSOR FOR
       SELECT col1, col2 FROM t1;
BEGIN
   PERFORM DBMS_OUTPUT.SERVEROUTPUT(TRUE); ...(1)
   PERFORM DBMS_OUTPUT.ENABLE(NULL); ...(1)
   PERFORM DBMS_OUTPUT.PUT_LINE('-- BATCH_001 Start --');
   OPEN cl;
   PERFORM DBMS_OUTPUT.PUT_LINE('-- LOOP Start --');
   LOOP
       FETCH c1 INTO v_col1, v_col2;
       EXIT WHEN FOUND = false;
       PERFORM DBMS_OUTPUT.PUT('.');
   END LOOP;
   PERFORM DBMS_OUTPUT.NEW_LINE(); ...(2)
   PERFORM DBMS_OUTPUT.PUT_LINE('-- LOOP End --');
   CLOSE c1;
   PERFORM DBMS_OUTPUT.PUT_LINE('-- BATCH_001 End --');
EXCEPTION
   WHEN OTHERS THEN
       PERFORM DBMS_OUTPUT.PUT_LINE('-- SQL Error --');
       PERFORM DBMS_OUTPUT.PUT_LINE('ERROR : ' | SQLERRM );
END;
$$
```

#### (1) SERVEROUTPUT/ENABLE

#### Specification differences

Oracle database

Use SET statement and specify SERVEROUTPUT ON.

**FUJITSU Enterprise Postgres** 

Specify DBMS\_SQL.SERVEROUTPUT(TRUE).

#### Conversion procedure

Convert using the following procedure:

- $1. \ \ Check if a \ SET \ SERVEROUTPUT \ statement \ is \ specified \ before \ the \ PL/SQL \ block \ of \ a \ stored \ procedure.$
- If a SET SERVEROUTPUT straight after BEGIN
  of PL/pgSQL. If ON is specified to have messages output to a window, then specify TRUE. If OFF is specified, then
  specify FALSE.
- 3. Specify DBMS\_SQL.ENABLE only if SET SERVEROUTPUT is ON. The values to be specified for the argument are as follows:
  - If SIZE is specified for the SET SERVEROUTPUT statement, specify this size for the argument.

- If SIZE is not specified for the SET SERVEROUTPUT statement, then specify 2000 for Oracle10.1g or earlier, NULL for Oracle10.2g or later.

If DBMS\_SQL.ENABLE is specified for the PL/SQL block of the stored procedure, specify the same value as that argument.

#### (2) NEW\_LINE

#### Specification differences

#### Oracle database

If there is no argument for packageName.featureName, parenthesis can be omitted.

#### **FUJITSU Enterprise Postgres**

Even if there is no argument for packageName.featureName, parenthesis cannot be omitted.

#### Conversion procedure

Convert using the following procedure:

- 1. Locate the places where the keyword "DBMS\_OUTPUT.NEW\_LINE" is used in the stored procedure.
- 2. If there is no parenthesis after packageName.featureName, add the parenthesis.

## B.5.2 Receiving a Return Value from a Procedure (PL/SQL) Block (For GET\_LINES)

#### Oracle database

```
set serveroutput off;
DECLARE
                INTEGER;
   v_num
BEGIN
   DBMS_OUTPUT.DISABLE; ...(3)
   DBMS_OUTPUT.ENABLE(20000); ...(4)
   DBMS_OUTPUT.PUT_LINE('-- ITEM CHECK --');
   SELECT count(*) INTO v_num FROM t1;
   IF v_num = 0 THEN
       DBMS_OUTPUT.PUT_LINE('-- NO ITEM --');
       DBMS_OUTPUT.PUT_LINE('-- IN ITEM(' | v_num |  ') --');
    END IF;
END;
set serveroutput on;
DECLARE
   v_buffs
               DBMSOUTPUT_LINESARRAY; ...(5)
               INTEGER := 10;
   v_num
BEGIN
   DBMS_OUTPUT.GET_LINES(v_buffs, v_num); ...(5)
    FOR i IN 1..v_num LOOP
        DBMS_OUTPUT.PUT_LINE('LOG : ' || v_buffs(i)); ...(5)
```

```
END LOOP;
END;
/
```

#### **FUJITSU Enterprise Postgres**

```
DO $$
DECLARE
   v_num
                 INTEGER;
BEGIN
   PERFORM DBMS_OUTPUT.SERVEROUTPUT(FALSE);
   PERFORM DBMS_OUTPUT.DISABLE(); ...(3)
   PERFORM DBMS_OUTPUT.ENABLE(20000); ...(4)
PERFORM DBMS_OUTPUT.PUT_LINE('-- ITEM CHECK --');
   SELECT count(*) INTO v_num FROM t1;
   IF v_num = 0 THEN
       PERFORM DBMS_OUTPUT.PUT_LINE('-- NO ITEM --');
   ELSE
        PERFORM DBMS_OUTPUT.PUT_LINE('-- IN ITEM(' | v_num |  ') --');
   END IF;
END;
$$
DO $$
DECLARE
   v_buffs
                VARCHAR[]; ...(5)
                INTEGER := 10;
   v_num
BEGIN
   PERFORM DBMS_OUTPUT.SERVEROUTPUT(TRUE);
   SELECT lines, numlines INTO v_buffs, v_num FROM DBMS_OUTPUT.GET_LINES(v_num); ...(5)
   FOR i IN 1..v_num LOOP
       PERFORM DBMS_OUTPUT.PUT_LINE('LOG : ' | v_buffs[i]); ...(5)
   END LOOP;
END;
$$
```

#### (3) DISABLE

Same as the NEW\_LINE in the DBMS\_OUTPUT package. Refer to NEW\_LINE for information on specification differences and conversion procedures associated with specification differences.

#### (4) ENABLE

Same as NEW\_LINE in the DBMS\_OUTPUT package. Refer to NEW\_LINE for information on specification differences and conversion procedures associated with specification differences.

#### (5) GET\_LINES

Specification format for Oracle database

DBMS\_OUTPUT.GET\_LINES(firstArg, secondArg)

Specification differences

Oracle database

Obtained values are received with variables specified for arguments.

#### **FUJITSU Enterprise Postgres**

Since obtained values are the search results for DBMS\_OUTPUT.GET\_LINES, they are received with variables specified for the INTO clause of the SELECT statement.

#### Conversion procedure

Convert using the following procedure:

- 1. Locate the places where the keyword "DBMS\_OUTPUT.GET\_LINES" is used in the stored procedure.
- 2. Change the data type (DBMSOUTPUT\_LINESARRAY in the example) of the variable (v\_buffs in the example) specified as *firstArg* of DBMS\_OUTPUT.GET\_LINES into a VARCHAR type array (VARCHAR[] in the example).
- 3. Replace the DBMS\_OUTPUT.GET\_LINES location called with a SELECT INTO statement.
  - Use the literal "lines, numlines" in the select list.
  - Specify *firstArg* (v\_buffs in the example) and *secondArg* (v\_num in the example) configured in DBMS\_OUTPUT.GET\_LINES, in the INTO clause.
  - Use DBMS\_OUTPUT.GET\_LINES in the FROM clause. Specify only *secondArg* (v\_num in the example) before modification.
- 4. Identify the location that references *firstArg* (v\_buffs in the example), and change it to the PL/pgSQL array reference format (v\_buffs[i] in the example).

## B.5.3 Receiving a Return Value from a Procedure (PL/SQL) Block (For GET\_LINE)

#### Oracle database

#### **FUJITSU Enterprise Postgres**

<sup>\*</sup> Only the process to obtain a value is stated

```
SELECT line, status INTO v_buff1, v_num FROM DBMS_OUTPUT.GET_LINE(); ...(6)
EXIT WHEN v_num = 1;
    v_buff2 := v_buff2 || v_buff1;
END LOOP;

PERFORM DBMS_OUTPUT.PUT_LINE(v_buff2);
END;
$$
;
```

### (6) GET\_LINE

Specification format for Oracle database

DBMS\_OUTPUT.GET\_LINE(firstArg, secondArg)

#### Specification differences

### Oracle database

Obtained values are received with variables specified for arguments.

### **FUJITSU Enterprise Postgres**

Since obtained values are the search results for DBMS\_OUTPUT.GET\_LINES, they are received with variables specified for the INTO clause of the SELECT statement.

#### Conversion procedure

Convert using the following procedure:

- 1. Locate the places where the keyword "DBMS\_OUTPUT.GET\_LINE" is used in the stored procedure.
- 2. Replace the DBMS\_OUTPUT.GET\_LINE location called with a SELECT INTO statement.
  - Use the literal "line, status" in the select list.
  - Specify *firstArg* (v\_buff1 in the example) and *secondArg* (v\_num in the example) configured in DBMS\_OUTPUT.GET\_LINE, in the INTO clause.
  - Use DBMS\_OUTPUT.GET\_LINE in the FROM clause. Although arguments are not specified, parenthesis must be specified.

# **B.6 UTL\_FILE (Perform File Operation)**

# Features

UTL\_FILE reads and writes text files from PL/pgSQL.

# **B.6.1** Registering a Directory to Load and Write Text Files

### **Oracle database**

```
[Oracle9i or earlier]
Configure the following with initialization parameter

UTL_FILE_DIR='/home/fsep' ...(1)

[Oracle9.2i or later]
Configure the following with CREATE DIRECTORY statement

CREATE DIRECTORY DIR AS '/home/fsep'; ...(1)
```

<sup>\*</sup> Only the process to obtain a value is stated

### **FUJITSU Enterprise Postgres**

```
INSERT INTO UTL_FILE.UTL_FILE_DIR(dir)

VALUES('/home/fsep'); ...(1)
```

#### (1) UTL\_FILE\_DIR/CREATE DIRECTORY

#### Feature differences

### Oracle database

Configure the directory to be operated, using the CREATE DIRECTORY statement or the initialization parameter UTL\_FILE\_DIR.

# **FUJITSU Enterprise Postgres**

The directory to be operated cannot be configured using the CREATE DIRECTORY statement or the initialization parameter UTL\_FILE\_DIR.

### Conversion procedure

Configure the target directory information in the UTL\_FILE.UTL\_FILE\_DIR table using the INSERT statement. Note that this conversion procedure should be performed only once before executing the PL/pgSQL function.

- When using the initialization parameter UTL\_FILE\_DIR:
  - 1. Check the initialization parameter UTL\_FILE\_DIR value ('/home/fsep' in the example).
  - 2. Using the INSERT statement, specify and execute the directory name checked in step 1.
    - Specify UTL\_FILE.UTL\_FILE\_DIR(dir) for the INTO clause.
    - Using the character string literal ('/home/fsep' in the example), specify the target directory name for the VALUES clause.
    - If multiple directories are specified, execute the INSERT statement for each directory.
- When using the CREATE DIRECTORY statement:
  - 1. Check the directory name ('/home/fsep' in the example) registered with the CREATE DIRECTORY statement. To check, log in SQL\*Plus as a user with DBA privileges, and execute "show ALL\_DIRECTORIES;".
  - 2. Using the INSERT statement, specify and execute the directory name checked in step 1. Same steps are used to specify the INSERT statement as when using the initialization parameter UTL\_FILE\_DIR.

# **B.6.2 Checking File Information**

#### **Oracle database**

```
CREATE PROCEDURE read_file(fname VARCHAR2) AS

v_file    UTL_FILE.FILE_TYPE;
v_exists    BOOLEAN;
v_length    NUMBER;
v_bsize    INTEGER;
v_rbuff    VARCHAR2(1024);

BEGIN

UTL_FILE.FGETATTR('DIR', fname, v_exists, v_length, v_bsize); ...(2)

IF v_exists <> true THEN
    DBMS_OUTPUT.PUT_LINE('-- FILE NOT FOUND ---');
    RETURN;
END IF;
```

```
DBMS_OUTPUT.PUT_LINE('-- FILE DATA --');
   v_file := UTL_FILE.FOPEN('DIR', fname, 'r', 1024); ...(3)
   FOR i IN 1..3 LOOP
       UTL_FILE.GET_LINE(v_file, v_rbuff, 1024); ...(4)
       DBMS_OUTPUT.PUT_LINE(v_rbuff);
   END LOOP;
   DBMS_OUTPUT.PUT_LINE('... more');
   DBMS_OUTPUT.PUT_LINE('-- READ END --');
   UTL_FILE.FCLOSE(v_file); ...(5)
   RETURN;
EXCEPTION
   WHEN NO_DATA_FOUND THEN
       DBMS_OUTPUT.PUT_LINE('-- FILE END --');
       UTL_FILE.FCLOSE(v_file);
       RETURN;
   WHEN OTHERS THEN
       DBMS_OUTPUT.PUT_LINE('-- SQL Error --');
       DBMS_OUTPUT.PUT_LINE('ERROR : ' | SQLERRM );
       UTL_FILE.FCLOSE_ALL; ...(6)
       RETURN;
END;
set serveroutput on
call read_file('file01.txt');
```

### **FUJITSU Enterprise Postgres**

```
CREATE FUNCTION read_file(fname VARCHAR) RETURNS void AS $$
DECLARE
   v_file
              UTL_FILE.FILE_TYPE;
   v_exists BOOLEAN;
   v_length NUMERIC;
   v_bsize INTEGER;
   v_rbuff
              VARCHAR(1024);
BEGIN
   PERFORM DBMS OUTPUT.SERVEROUTPUT(TRUE);
   SELECT fexists, file_length, blocksize
      INTO v_exists, v_length, v_bsize
      FROM UTL_FILE.FGETATTR('/home/fsep', fname); ...(2)
   IF v_exists <> true THEN
       PERFORM DBMS_OUTPUT.PUT_LINE('-- FILE NOT FOUND --');
       RETURN;
   END IF;
   PERFORM DBMS_OUTPUT.PUT_LINE('-- FILE DATA --');
   v_file := UTL_FILE.FOPEN('/home/fsep', fname, 'w', 1024); ...(3)
   FOR i IN 1..3 LOOP
       v_rbuff := UTL_FILE.GET_LINE(v_file, 1024); ...(4)
       PERFORM DBMS_OUTPUT.PUT_LINE(v_rbuff);
   END LOOP;
   PERFORM DBMS_OUTPUT.PUT_LINE('... more');
   PERFORM DBMS_OUTPUT.PUT_LINE('-- READ END --');
   v_file := UTL_FILE.FCLOSE(v_file); ...(5)
```

### (2) FGETATTR

#### Specification format for Oracle database

UTL\_FILE.FGETATTR(firstArg, secondArg, thirdArg, fourthArg, fifthArg)

### Feature differences

#### Oracle database

If using a CREATE DIRECTORY statement (Oracle9.2i or later), specify a directory object name for the directory name.

#### **FUJITSU Enterprise Postgres**

A directory object name cannot be specified for the directory name.

#### Specification differences

#### Oracle database

Obtained values are received with variables specified for arguments.

# FUJITSU Enterprise Postgres

Since obtained values are the search results for UTL\_FILE.FGETATTR, they are received with variables specified for the INTO clause of the SELECT statement.

#### Conversion procedure

Convert using the following procedure. Refer to UTL\_FILE\_DIR/CREATE DIRECTORY for information on how to check if the directory object name corresponds to the actual directory name.

- $1. \ \ Locate \ the \ places \ where \ the \ keyword \ "UTL\_FILE.FOPEN" \ is \ used \ in \ the \ stored \ procedure.$
- 2. Check the actual directory name ('/home/fsep' in the example) that corresponds to the directory object name ('DIR' in the example).
- 3. Replace the directory object name ('DIR' in the example) in *firstArg* with the actual directory name ('/home/fsep' in the example) verified in step 2.
- 4. Replace the UTL\_FILE.FGETATTR location called with a SELECT INTO statement.
  - Use the literal "fexists, file\_length, blocksize" in the select list.
  - Specify *thirdArg*, *fourthArg*, and *fifthArg* (v\_exists, v\_length, v\_bsize, in the example) specified for UTL\_FILE.FGETATTR to the INTO clause in the same order as that of the arguments.
  - Use UTL\_FILE.FGETATTR in the FROM clause. Specify only the actual directory name for *firstArg* ('home/fsep' in the example) and *secondArg* (fname in the example) before modification for the arguments.

# (3) FOPEN

#### Specification format for Oracle

UTL\_FILE.FOPEN(firstArg, secondArg, thirdArg, fourthArg, fifthArg)

#### Feature differences

#### Oracle database

If using a CREATE DIRECTORY statement (Oracle9.2i or later), specify a directory object name for the directory name.

### **FUJITSU Enterprise Postgres**

A directory object name cannot be specified for the directory name.

#### Conversion procedure

Convert using the following procedure. Refer to UTL\_FILE\_DIR/CREATE DIRECTORY for information on how to check if the directory object name corresponds to the actual directory name.

- 1. Locate the places where the keyword "UTL\_FILE.FOPEN" is used in the stored procedure.
- 2. Check the actual directory name ('/home/fsep' in the example) that corresponds to the directory object name ('DIR' in the example).
- 3. Replace the directory object name ('DIR' in the example) in *firstArg* with the actual directory name ('/home/fsep' in the example) checked in step 1.

# (4) GET\_LINE

# Specification format for Oracle database

UTL\_FILE.GET\_LINE(firstArg, secondArg, thirdArg, fourthArg)

# Specification differences

# Oracle database

Obtained values are received with variables specified for arguments.

### **FUJITSU Enterprise Postgres**

Since obtained values are the returned value of UTL\_FILE.GET\_LINE, they are received with variables specified for substitution statement.

# Conversion procedure

Convert using the following procedure:

- 1. Locate the places where the keyword "UTL\_FILE.GET\_LINE" is used in the stored procedure.
- 2. Replace the UTL\_FILE.GET\_LINE location called with a value assignment (:=).
  - On the left-hand side, specify secondArg (v\_rbuff in the example) specified for UTL\_FILE.GET\_LINE.
  - Use UTL\_FILE.GET\_LINE in the right-hand side. Specify only *firstArg* (v\_file in the example) and *thirdArg* (1024 in the example) before modification.

# (5) FCLOSE

# Specification format for Oracle database

UTL\_FILE.FCLOSE(firstArg)

### Specification differences

### Oracle database

After closing, the file handler specified for the argument becomes NULL.

### **FUJITSU Enterprise Postgres**

After closing, set the file handler to NULL by assigning the return value of UTL\_FILE.FCLOSE to it.

### Conversion procedure

Convert using the following procedure:

- 1. Locate the places where the keyword "UTL\_FILE.FCLOSE" is used in the stored procedure.
- 2. Replace the UTL\_FILE.FCLOSE location called with a value assignment (:=) so that the file handler (v\_file in the example) becomes NULL.
  - On the left-hand side, specify the argument (v\_file in the example) specified for UTL\_FILE.FCLOSE.
  - Use UTL\_FILE.FCLOSE in the right-hand side. For the argument, specify the same value (v\_file in the example) as before modification.

### (6) FCLOSE\_ALL

Same as NEW\_LINE in the DBMS\_OUTPUT package. Refer to NEW\_LINE in the DBMS\_OUTPUT for information on specification differences and conversion procedures associated with specification differences.

# **B.6.3 Copying Files**

### Oracle database

```
CREATE PROCEDURE copy_file(fromname VARCHAR2, toname VARCHAR2) AS
BEGIN

UTL_FILE.FCOPY('DIR1', fromname, 'DIR2', toname, 1, NULL); ...(7)

RETURN;

EXCEPTION
WHEN OTHERS THEN
DBMS_OUTPUT_PUT_LINE('-- SQL Error --');

DBMS_OUTPUT.PUT_LINE('ERROR : ' || SQLERRM );
RETURN;

END;
/

set serveroutput on
call copy_file('file01.txt','file01_bk.txt');
```

# **FUJITSU Enterprise Postgres**

```
CREATE FUNCTION copy_file(fromname VARCHAR, toname VARCHAR) RETURNS void AS $$
BEGIN

PERFORM DBMS_OUTPUT.SERVEROUTPUT(TRUE);

PERFORM UTL_FILE.FCOPY('/home/fsep', fromname, '/home/backup', toname, 1, NULL); ...(7)

RETURN;

EXCEPTION

WHEN OTHERS THEN

PERFORM DBMS_OUTPUT.PUT_LINE('-- SQL Error --');

PERFORM DBMS_OUTPUT.PUT_LINE('ERROR : ' || SQLERRM );

RETURN;

END;
```

```
$$
LANGUAGE plpgsql;

SELECT copy_file('file01.txt','file01_bk.txt');
```

### (7) FCOPY

#### Specification format for Oracle database

UTL\_FILE.FCOPY(firstArg, secondArg, thirdArg, fourthArg, fifthArg, sixthArg)

### Feature differences

#### Oracle database

If using a CREATE DIRECTORY statement (Oracle9.2i or later), specify a directory object name for the directory name.

### **FUJITSU Enterprise Postgres**

A directory object name cannot be specified for the directory name.

#### Conversion procedure

Convert using the following procedure. Refer to UTL\_FILE\_DIR/CREATE DIRECTORY for information on how to check if the directory object name corresponds to the actual directory name.

- 1. Locate the places where the keyword "UTL\_FILE.FCOPY" is used in the stored procedure.
- 2. Check the actual directory names ('/home/fsep' and '/home/backup', in the example) that correspond to the directory object names ('DIR1' and 'DIR2', in the example) of *firstArg* and *thirdArg* argument.
- 3. Replace the directory object name ('DIR1' and 'DIR2', in the example) with the actual directory names ('/home/fsep' in the example) checked in step 1.

# **B.6.4 Moving/Renaming Files**

#### Oracle database

```
CREATE PROCEDURE move_file(fromname VARCHAR2, toname VARCHAR2) AS
BEGIN

UTL_FILE.FRENAME('DIR1', fromname, 'DIR2', toname, FALSE); ...(8)
RETURN;

EXCEPTION
WHEN OTHERS THEN
DBMS_OUTPUT.PUT_LINE('-- SQL Error --');
DBMS_OUTPUT.PUT_LINE('ERROR : ' || SQLERRM );
RETURN;

END;
/
set serveroutput on
call move_file('file01.txt','file02.txt');
```

#### **FUJITSU Enterprise Postgres**

```
CREATE FUNCTION move_file(fromname VARCHAR, toname VARCHAR) RETURNS void AS $$
BEGIN
PERFORM DBMS_OUTPUT.SERVEROUTPUT(TRUE);
```

```
PERFORM UTL_FILE.FRENAME('/home/fsep', fromname, '/home/backup', toname, FALSE); ...(8)
    RETURN;

EXCEPTION
    WHEN OTHERS THEN
        PERFORM DBMS_OUTPUT.PUT_LINE('-- SQL Error --');
        PERFORM DBMS_OUTPUT.PUT_LINE('ERROR : ' || SQLERRM );
        RETURN;

END;
$$
LANGUAGE plpgsql;

SELECT move_file('file01.txt','file02.txt');
```

# (8) FRENAME

Same as FCOPY for the UTL\_FILE package. Refer to FCOPY in the UTL\_FILE package for information on specification differences and conversion procedures associated with specification differences.

# B.7 DBMS\_SQL (Execute Dynamic SQL)

#### **Features**

For DBMS\_SQL, dynamic SQL can be executed from PL/pgSQL.

# **B.7.1 Searching Using a Cursor**

# Oracle database

```
CREATE PROCEDURE search_test(h_where CLOB) AS
   str_sql
              CLOB;
              INTEGER;
   v_cnt
              DBMS_SQL.VARCHAR2A;
   v_array
              INTEGER;
   v cur
   v_smpid
              INTEGER;
              VARCHAR2(20);
   v smpnm
   v_addbuff VARCHAR2(20);
   v_smpage INTEGER;
   errcd
              INTEGER;
              INTEGER;
   length
              INTEGER;
BEGIN
             := 'SELECT smpid, smpnm FROM smp_tbl WHERE ' || h_where || ' ORDER BY smpid';
   str_sql
             := 0;
   v_smpid
              := '';
   v_smpnm
   v_smpage
               := 0;
   v_cur := DBMS_SQL.OPEN_CURSOR; ...(1)
   v cnt :=
     CEIL(DBMS LOB.GETLENGTH(str sql)/1000);
   FOR idx IN 1 .. v_cnt LOOP
       v_array(idx) :=
           DBMS_LOB.SUBSTR(str_sql,
                           (idx-1)*1000+1);
```

```
END LOOP;
   {\tt DBMS\_SQL.PARSE}(v\_cur,\ v\_array,\ 1,\ v\_cnt,\ {\tt FALSE},\ {\tt DBMS\_SQL.NATIVE});\ \dots (2)
   DBMS_SQL.DEFINE_COLUMN(v_cur, 1, v_smpid);
   DBMS_SQL.DEFINE_COLUMN(v_cur, 2, v_smpnm, 10);
   ret := DBMS_SQL.EXECUTE(v_cur);
   LOOP
      v_addbuff := '';
       IF DBMS_SQL.FETCH_ROWS(v_cur) = 0 THEN
          EXIT;
       END IF;
       DBMS_OUTPUT.PUT_LINE('-----');
       DBMS_SQL.COLUMN_VALUE(v_cur, 1, v_smpid, errcd, length); ...(3)
       IF errcd = 1405 THEN ...(3)
        DBMS_OUTPUT.PUT_LINE('smpid
                                   = (NULL)');
       ELSE
        DBMS_OUTPUT.PUT_LINE('smpid = ' || v_smpid);
       END IF;
       DBMS_SQL.COLUMN_VALUE(v_cur, 2, v_smpnm, errcd, length);
       IF errcd = 1406 THEN
        v_addbuff := '... [len=' | length | | ']';
       IF errcd = 1405 THEN
        DBMS_OUTPUT.PUT_LINE('v_smpnm
                                       = (NULL)');
        DBMS_OUTPUT.PUT_LINE('v_smpnm = ' || v_smpnm || v_addbuff );
       END IF;
DBMS_OUTPUT.PUT_LINE('-----');
      DBMS_OUTPUT.NEW_LINE;
   END LOOP;
   DBMS_SQL.CLOSE_CURSOR(v_cur); ...(4)
   RETURN;
END;
Set serveroutput on
call search_test('smpid < 100');</pre>
```

# **FUJITSU Enterprise Postgres**

```
CREATE FUNCTION search_test(h_where text) RETURNS void AS $$

DECLARE

str_sql text;

v_cur INTEGER;
```

```
v_smpid INTEGER;
v_smpnm VARCHAR(20);
   v_addbuff VARCHAR(20);
   v_smpage INTEGER;
            INTEGER;
   errcd
   length
             INTEGER;
   ret
             INTEGER;
BEGIN
   PERFORM DBMS_OUTPUT.SERVEROUTPUT(TRUE);
   str_sql := 'SELECT smpid, smpnm FROM smp_tbl WHERE ' | h_where | | ' ORDER BY smpid';
   v_smpid
   v_smpnm
              := '';
   v_smpage
              := 0;
   v_cur := DBMS_SQL.OPEN_CURSOR(); ...(1)
   PERFORM DBMS_SQL.PARSE(v_cur, str_sql, 1); ...(2)
   PERFORM DBMS_SQL.DEFINE_COLUMN(v_cur, 1, v_smpid);
   PERFORM DBMS_SQL.DEFINE_COLUMN(v_cur, 2, v_smpnm, 10);
   ret := DBMS_SQL.EXECUTE(v_cur);
   LOOP
       v_addbuff := '';
       IF DBMS_SQL.FETCH_ROWS(v_{cur}) = 0 THEN
          EXIT;
       END IF;
       PERFORM
DBMS_OUTPUT.PUT_LINE('-----');
       SELECT value,column_error,actual_length
        INTO v_smpid, errcd, length
        FROM DBMS_SQL.COLUMN_VALUE(v_cur,
                                 v_smpid); ...(3)
       IF errcd = 22002 THEN ...(3)
        PERFORM DBMS_OUTPUT.PUT_LINE('smpid
                                           = (NULL)');
         PERFORM DBMS_OUTPUT.PUT_LINE('smpid = ' | v_smpid);
       END IF;
       SELECT value,column_error,actual_length INTO v_smpnm, errcd, length FROM
DBMS_SQL.COLUMN_VALUE(v_cur, 2, v_smpnm);
       IF errcd = 22001 THEN
        v_addbuff := '... [len=' || length || ']';
       END IF;
       IF errcd = 22002 THEN
        PERFORM DBMS OUTPUT.PUT LINE('v smpnm = (NULL)');
        PERFORM DBMS_OUTPUT.PUT_LINE('v_smpnm = ' | v_smpnm | v_addbuff );
       END IF;
       PERFORM
DBMS_OUTPUT.PUT_LINE('-----');
      PERFORM DBMS_OUTPUT.NEW_LINE();
   END LOOP;
   v_cur := DBMS_SQL.CLOSE_CURSOR(v_cur); ...(4)
   RETURN;
END:
$$
LANGUAGE plpgsql;
```

# (1) OPEN\_CURSOR

Same as NEW\_LINE in the DBMS\_OUTPUT package. Refer to NEW\_LINE in the DBMS\_OUTPUT package for information on specification differences and conversion procedures associated with specification differences.

### (2) PARSE

#### Specification format for Oracle database

DBMS\_SQL.PARSE(firstArg, secondArg, thirdArg, fourthArg, fifthArg)

#### Feature differences

### Oracle database

SQL statements can be specified with string table types (VARCHAR2A type, VARCHAR2S type). Specify this for *secondArg*.

DBMS\_SQL.NATIVE, DBMS\_SQL.V6, DBMS\_SQL.V7 can be specified for processing SQL statements.

### **FUJITSU Enterprise Postgres**

SQL statements cannot be specified with string table types.

DBMS\_SQL.NATIVE, DBMS\_SQL.V6, DBMS\_SQL.V7 cannot be specified for processing SQL statements.

#### Conversion procedure

Convert using the following procedure:

- 1. Locate the places where the keyword "DBMS\_SQL.PARSE" is used in the stored procedure.
- 2. Check the data type of the SQL statement specified for secondArg (v\_array in the example).
  - If the data type is either DBMS\_SQL.VARCHAR2A type or DBMS\_SQL.VARCHAR2S type, then it is a table type specification. Execute step 3 and continue the conversion process.
  - If the data type is neither DBMS\_SQL.VARCHAR2A type nor DBMS\_SQL.VARCHAR2S type, then it is a string specification. Execute step 7 and continue the conversion process.
- 3. Check the SQL statement (str\_sql in the example) before it was divided into DBMS\_SQL.VARCHAR2A type and DBMS\_SQL.VARCHAR2S type.
- 4. Delete the sequence of the processes (processes near FOR idx in the example) where SQL is divided into DBMS\_SQL.VARCHAR2A type and DBMS\_SQL.VARCHAR2S type.
- 5. Replace *secondArg* with the SQL statement (str\_sql in the example) before it is divided, that was checked in step
- 6. Delete thirdArg, fourthArg, and fifthArg (v\_cnt, FALSE, DBMS\_SQL.NATIVE, in the example).
- 7. If DBMS\_SQL.NATIVE, DBMS\_SQL.V6, and DBMS\_SQL.V7 are specified, then replace *thirdArg* with a numeric literal 1.
  - If either DBMS\_SQL.VARCHAR2A type or DBMS\_SQL.VARCHAR2S type is used, then sixthArg becomes relevant.
  - If neither DBMS\_SQL.VARCHAR2A type nor DBMS\_SQL.VARCHAR2S type is used, then *thirdArg* becomes relevant.

### (3) COLUMN\_VALUE

# Specification format for Oracle database

DBMS\_SQL.COLUMN\_VALUE(firstArg, secondArg, thirdArg, fourthArg, fifthArg)

#### Feature differences

#### Oracle database

The following error codes are returned for column\_error.

- 1406: fetched column value was truncated
- 1405: fetched column value is NULL

### **FUJITSU Enterprise Postgres**

The following error codes are returned for column\_error.

- 22001: string\_data\_right\_truncation
- 22002: null\_value\_no\_indicator\_parameter

#### Specification differences

### Oracle database

Obtained values are received with variables specified for arguments.

# **FUJITSU Enterprise Postgres**

Since obtained values are the search results for DBMS\_SQL.COLUMN\_VALUE, they are received with variables specified for the INTO clause of the SELECT statement.

### Conversion procedure

Convert using the following procedure:

- 1. Locate the places where the keyword "DBMS\_SQL.COLUMN\_VALUE" is used in the stored procedure.
- 2. Replace the DBMS\_SQL.COLUMN\_VALUE location called with a SELECT INTO statement.
  - Check the number of arguments (v\_smpid, errcd, and length in the example) specified after *secondArg* (1 in the example) of DBMS\_SQL.COLUMN\_VALUE.
  - Specify "value", "column\_error", and "actual\_length" in the select list, according to the number of arguments checked in the previous step (for example, if only *thirdArg* is specified, then specify "value" only.)
  - Specify *thirdArg*, *fourthArg*, and *fifthArg* (v\_smpid, errcd, length in the example) configured for DBMS\_SQL.COLUMN\_VALUE, for the INTO clause.
  - Use DBMS\_SQL.COLUMN\_VALUE in the FROM clause. Specify *firstArg*, *secondArg*, and *thirdArg* (v\_cur, 1, v\_smpid, in the example) before modification.
- 3. If the *fourthArg* (column\_error value in the example) is used, then check the location of the target variable (errcd in the example).
- 4. If a decision process is performed in the location checked, then modify the values used in the decision process as below:
  - 1406 to 22001
  - 1405 to 22002

### (4) CLOSE\_CURSOR

Specification format for Oracle database

 $DBMS\_SQL.CLOSE\_CURSOR(\textit{firstArg})$ 

# Specification differences

### Oracle database

After closing, the cursor specified in firstArg becomes NULL.

# FUJITSU Enterprise Postgres

After closing, set the cursor to NULL by assigning the return value of DBMS\_SQL.CLOSE\_CURSOR to it.

### Conversion procedure

Convert using the following procedure:

- $1. \ Locate the places where the keyword "DBMS\_SQL.CLOSE\_CURSOR" is used in the stored procedure.$
- 2. Set the cursor to NULL by assigning (:=) the return value of DBMS\_SQL.CLOSE\_CURSOR to it.
  - On the left-hand side, specify the argument (v\_cur in the example) specified for DBMS\_SQL.CLOSE\_CURSOR.
  - Use DBMS\_SQL.CLOSE\_CURSOR in the right-hand side. For the argument, specify the same value (v\_cur in the example) as before modification.

# Appendix C Tables Used by the Features Compatible with Oracle Databases

This chapter describes the tables used by the features compatible with Oracle databases.

# C.1 UTL\_FILE.UTL\_FILE\_DIR

Register the directory handled by the UTL\_FILE package in the UTL\_FILE.UTL\_FILE\_DIR table.

Name	Туре	Description
dir	text	Name of the directory handled by the UTL_FILE package

# Appendix D ECOBPG - Embedded SQL in COBOL

This appendix describes application development using embedded SQL in COBOL.

# D.1 Precautions when Using Functions and Operators

An embedded SQL program consists of code written in an ordinary programming language, in this case COBOL, mixed with SQL commands in specially marked sections. To build the program, the source code (\*.pco) is first passed through the embedded SQL preprocessor, which converts it to an ordinary COBOL program (\*.cob), and afterwards it can be processed by a COBOL compiler. (For details about the compiling and linking see "D.9 Processing Embedded SQL Programs".) Converted ECOBPG applications call functions in the libpq library through the embedded SQL library (ecpglib), and communicate with the PostgreSQL server using the normal frontend-backend protocol.

Embedded SQL has advantages over other methods for handling SQL commands from COBOL code. First, it takes care of the tedious passing of information to and from variables in your C program. Second, the SQL code in the program is checked at build time for syntactical correctness. Third, embedded SQL in COBOL is specified in the SQL standard and supported by many other SQL database systems. The PostgreSQL implementation is designed to match this standard as much as possible, and it is usually possible to port embedded SQL programs written for other SQL databases to PostgreSQL with relative ease.

As already stated, programs written for the embedded SQL interface are normal COBOL programs with special code inserted to perform database-related actions. This special code always has the form:

EXEC SQL ... END-EXEC

These statements syntactically take the place of a COBOL statement. Depending on the particular statement, they can appear at the data division or at the procedure division. Actual executable SQLs need to be placed at the procedure division, and host variable declarations need to be placed at data division. However, the precompiler does not validate their placements. Embedded SQL statements follow the case-sensitivity rules of normal SQL code, and not those of COBOL.

For COBOL code notation, "fixed' or "variable" can be used. In each line, columns 1 to 6 constitute the line number area, and column 7 is the indicator area. Embedded SQL programs also should be placed in area B (column 12 and beyond).

Note that sample code in this document omits indents for each area.

ECOBPG processes or outputs programs according to the COBOL code notation. COBOL code notation is specified using the ecobpg command. Note, however, that the following restrictions apply:

- For "fixed" notation, area B is from columns 12 to 72. Characters in column 73 and beyond are deleted in the precompiled source.
- For "variable" notation, area B is from column 12 to the last column of that record (up to column 251). Characters in column 252 and beyond are deleted in the precompiled source.

ECOBPG accepts as many COBOL statements as possible. Note, however, that the following restrictions apply:

- In declaring host variable section, you can't use debug line.
- Outside of declaring host variable section, you can use debug line, but you can't contain any SQL in debug lines.
- In declaring host variable section, you can't use commas or semicolons as separator. Use space instead.
- EXEC SQL VAR command, it can be used in ECPG, is not available in ECOBPG. Use REDEFINE clause of COBOL instead.

The following sections explain all the embedded SQL statements.

# **D.2 Managing Database Connections**

This section describes how to open, close, and switch database connections.

# D.2.1 Connecting to the Database Server

One connects to a database using the following statement:

```
EXEC SQL CONNECT TO target [AS connection-name] [USER user-name] END-EXEC.
```

The target can be specified in the following ways:

- dbname[@hostname][:port]
- tcp:postgresql://hostname[:port][/dbname][?options]
- unix:postgresql://hostname[:port][/dbname][?options]
- an SQL string literal containing one of the above forms
- a reference to a character variable containing one of the above forms (see examples)
- DEFAULT

If you specify the connection target literally (that is, not through a variable reference) and you don't quote the value, then the case-insensitivity rules of normal SQL are applied. In that case you can also double-quote the individual parameters separately as needed. In practice, it is probably less error-prone to use a (single-quoted) string literal or a variable reference. The connection target DEFAULT initiates a connection to the default database under the default user name. No separate user name or connection name can be specified in that case.

There are also different ways to specify the user name:

- username
- username/password
- username IDENTIFIED BY password
- username USING password

As above, the parameters username and password can be an SQL identifier, an SQL string literal, or a reference to a character variable

The connection-name is used to handle multiple connections in one program. It can be omitted if a program uses only one connection. The most recently opened connection becomes the current connection, which is used by default when an SQL statement is to be executed (see later in this chapter).

Here are some examples of CONNECT statements:

```
EXEC SQL CONNECT TO mydb@sql.mydomain.com END-EXEC.
```

EXEC SQL CONNECT TO tcp:postgresql://sql.mydomain.com/mydb AS myconnection USER john END-EXEC.

```
EXEC SQL BEGIN DECLARE SECTION END-EXEC.

01 TARGET PIC X(25).

01 USER PIC X(5).

EXEC SQL END DECLARE SECTION END-EXEC.

...

MOVE "mydb@sql.mydomain.com" TO TARGET.

MOVE "john" TO USER.

EXEC SQL CONNECT TO :TARGET USER :USER END-EXEC.
```

The last form makes use of the variant referred to above as character variable reference. For this purpose, only fixed-length string(no VARYING) variable can be used. Trailing spaces are ignored. You will see in later sections how COBOL variables can be used in SQL statements when you prefix them with a colon.

Be advised that the format of the connection target is not specified in the SQL standard. So if you want to develop portable applications, you might want to use something based on the last example above to encapsulate the connection target string somewhere.

# D.2.2 Choosing a Connection

SQL statements in embedded SQL programs are by default executed on the current connection, that is, the most recently opened one. If an application needs to manage multiple connections, then there are two ways to handle this.

The first option is to explicitly choose a connection for each SQL statement, for example:

```
EXEC SQL AT connection-name SELECT ... END-EXEC.
```

This option is particularly suitable if the application needs to use several connections in mixed order.

If your application uses multiple threads of execution, they cannot share a connection concurrently. You must either explicitly control access to the connection (using mutexes) or use a connection for each thread. If each thread uses its own connection, you will need to use the AT clause to specify which connection the thread will use.

The second option is to execute a statement to switch the current connection. That statement is:

```
EXEC SQL SET CONNECTION connection-name END-EXEC.
```

This option is particularly convenient if many statements are to be executed on the same connection. It is not thread-aware.

Here is an example program managing multiple database connections:

```
EXEC SQL BEGIN DECLARE SECTION END-EXEC.
   01 DBNAME PIC X(7).
EXEC SOL END DECLARE SECTION END-EXEC.
   EXEC SOL CONNECT TO testdb1 AS con1 USER testuser END-EXEC.
   EXEC SOL CONNECT TO testdb2 AS con2 USER testuser END-EXEC.
   EXEC SQL CONNECT TO testdb3 AS con3 USER testuser END-EXEC.
   This query would be executed in the last opened database "testdb3".
   EXEC SQL SELECT current_database() INTO :DBNAME END-EXEC.
   DISPLAY "current=" DBNAME " (should be testdb3)".
   Using "AT" to run a query in "testdb2"
   EXEC SQL AT con2 SELECT current_database() INTO :DBNAME END-EXEC.
   DISPLAY "current=" DBNAME " (should be testdb2)".
   Switch the current connection to "testdb1".
   EXEC SOL SET CONNECTION con1 END-EXEC.
   EXEC SQL SELECT current_database() INTO :DBNAME END-EXEC.
   DISPLAY "current=" DBNAME " (should be testdb1)".
   EXEC SQL DISCONNECT ALL END-EXEC.
```

This example would produce this output:

```
current=testdb3 (should be testdb3)
current=testdb2 (should be testdb2)
current=testdb1 (should be testdb1)
```

# D.2.3 Closing a Connection

To close a connection, use the following statement:

```
EXEC SQL DISCONNECT [connection] END-EXEC.
```

The connection can be specified in the following ways:

- connection-name
- DEFAULT
- CURRENT
- ALL

If no connection name is specified, the current connection is closed.

It is good style that an application always explicitly disconnect from every connection it opened.

# D.3 Running SQL Commands

Any SQL command can be run from within an embedded SQL application. Below are some examples of how to do that.

# D.3.1 Executing SQL Statements

### Creating a table:

```
EXEC SQL CREATE TABLE foo (number integer, ascii char(16)) END-EXEC.

EXEC SQL CREATE UNIQUE INDEX numl ON foo(number) END-EXEC.

EXEC SQL COMMIT END-EXEC.
```

### Inserting rows:

```
EXEC SQL INSERT INTO foo (number, ascii) VALUES (9999, 'doodad') END-EXEC.
EXEC SQL COMMIT END-EXEC.
```

#### Deleting rows:

```
EXEC SQL DELETE FROM foo WHERE number = 9999 END-EXEC.

EXEC SQL COMMIT END-EXEC.
```

# Updates:

```
EXEC SQL UPDATE foo

SET ascii = 'foobar'

WHERE number = 9999 END-EXEC.

EXEC SQL COMMIT END-EXEC.
```

SELECT statements that return a single result row can also be executed using EXEC SQL directly. To handle result sets with multiple rows, an application has to use a cursor; see "D.3.2 Using Cursors" below. (As a special case, an application can fetch multiple rows at once into an array host variable; see "Arrays".)

### Single-row select:

```
EXEC SQL SELECT foo INTO :FooBar FROM table1 WHERE ascii = 'doodad' END-EXEC.
```

Also, a configuration parameter can be retrieved with the SHOW command:

```
EXEC SQL SHOW search_path INTO :var END-EXEC.
```

The tokens of the form :something are *host variables*, that is, they refer to variables in the COBOL program. They are explained in "D.4 Using Host Variables".

# **D.3.2 Using Cursors**

To retrieve a result set holding multiple rows, an application has to declare a cursor and fetch each row from the cursor. The steps to use a cursor are the following: declare a cursor, open it, fetch a row from the cursor, repeat, and finally close it.

### Select using cursors:

```
EXEC SQL DECLARE foo_bar CURSOR FOR

SELECT number, ascii FROM foo

ORDER BY ascii END-EXEC.

EXEC SQL OPEN foo_bar END-EXEC.

EXEC SQL FETCH foo_bar INTO :FooBar, :DooDad END-EXEC.

...

EXEC SQL CLOSE foo_bar END-EXEC.

EXEC SQL CLOSE foo_bar END-EXEC.
```

For more details about declaration of the cursor, see "D.11.4 DECLARE", and refer to "SQL Commands" in "Reference" in the PostgreSQL Documentation for information on FETCH command.

Note: The ECOBPG DECLARE command does not actually cause a statement to be sent to the PostgreSQL backend. The cursor is opened in the backend (using the backend's DECLARE command) at the point when the OPEN command is executed.

# **D.3.3 Managing Transactions**

In the default mode, statements are committed only when EXEC SQL COMMIT is issued. The embedded SQL interface also supports autocommit of transactions (similar to libpq behavior) via the -t command-line option to ecobpg or via the EXEC SQL SET AUTOCOMMIT TO ON statement. In autocommit mode, each command is automatically committed unless it is inside an explicit transaction block. This mode can be explicitly turned off using EXEC SQL SET AUTOCOMMIT TO OFF.



Refer to "ecpg" in "PostgreSQL Client Applications" in the PostgreSQL Documentation for information on -t command-line option to ecobpg.

The following transaction management commands are available:

EXEC SQL COMMIT END-EXEC

Commit an in-progress transaction.

EXEC SQL ROLLBACK END-EXEC

Roll back an in-progress transaction.

EXEC SQL SET AUTOCOMMIT TO ON END-EXEC

Enable autocommit mode.

EXEC SQL SET AUTOCOMMIT TO OFF END-EXEC

Disable autocommit mode. This is the default.

# **D.3.4 Prepared Statements**

When the values to be passed to an SQL statement are not known at compile time, or the same statement is going to be used many times, then prepared statements can be useful.

The statement is prepared using the command PREPARE. For the values that are not known yet, use the placeholder "?":

```
EXEC SQL PREPARE stmt1 FROM "SELECT oid, datname FROM pg_database WHERE oid = ?" END-EXEC.
```

If a statement returns a single row, the application can call EXECUTE after PREPARE to execute the statement, supplying the actual values for the placeholders with a USING clause:

```
EXEC SQL EXECUTE stmt1 INTO :dboid, :dbname USING 1 END-EXEC.
```

If a statement returns multiple rows, the application can use a cursor declared based on the prepared statement. To bind input parameters, the cursor must be opened with a USING clause:

```
EXEC SQL PREPARE stmt1 FROM "SELECT oid,datname FROM pg_database WHERE oid > ?" END-EXEC.

EXEC SQL DECLARE foo_bar CURSOR FOR stmt1 END-EXEC.

* when end of result set reached, break out of while loop

EXEC SQL WHENEVER NOT FOUND GOTO FETCH-END END-EXEC.

EXEC SQL OPEN foo_bar USING 100 END-EXEC.

...

PERFORM NO LIMIT

EXEC SQL FETCH NEXT FROM foo_bar INTO :dboid, :dbname END-EXEC

END-PERFORM.

FETCH-END.

EXEC SQL CLOSE foo_bar END-EXEC.
```

When you don't need the prepared statement anymore, you should deallocate it:

```
EXEC SQL DEALLOCATE PREPARE name END-EXEC.
```

For more details about PREPARE, see "D.11.10 PREPARE". Also see "D.5 Dynamic SQL" for more details about using placeholders and input parameters.

# D.4 Using Host Variables

In "D.3 Running SQL Commands" you saw how you can execute SQL statements from an embedded SQL program. Some of those statements only used fixed values and did not provide a way to insert user-supplied values into statements or have the program process the values returned by the query. Those kinds of statements are not really useful in real applications. This section explains in detail how you can pass data between your COBOL program and the embedded SQL statements using a simple mechanism called host variables. In an embedded SQL program we consider the SQL statements to be guests in the COBOL program code which is the host language. Therefore the variables of the COBOL program are called host variables.

Another way to exchange values between PostgreSQL backends and ECOBPG applications is the use of SQL descriptors, described in "D.6 Using Descriptor Areas".

# D.4.1 Overview

Passing data between the COBOL program and the SQL statements is particularly simple in embedded SQL. Instead of having the program paste the data into the statement, which entails various complications, such as properly quoting the value, you can simply write the name of a COBOL variable into the SQL statement, prefixed by a colon. For example:

```
EXEC SQL INSERT INTO sometable VALUES (:v1, 'foo', :v2) END-EXEC.
```

This statements refers to two COBOL variables named v1 and v2 and also uses a regular SQL string literal, to illustrate that you are not restricted to use one kind of data or the other.

This style of inserting COBOL variables in SQL statements works anywhere a value expression is expected in an SQL statement.

# D.4.2 Declare Sections

To pass data from the program to the database, for example as parameters in a query, or to pass data from the database back to the program, the COBOL variables that are intended to contain this data need to be declared in specially marked sections, so the embedded SQL preprocessor is made aware of them.

This section starts with:

```
EXEC SQL BEGIN DECLARE SECTION END-EXEC.
```

and ends with:

```
EXEC SQL END DECLARE SECTION END-EXEC.
```

Between those lines, there must be normal COBOL variable declarations, such as:

```
01 INTX PIC S9(9) COMP VALUE 4.
01 FOO PIC X(15).
01 BAR PIC X(15).
```

As you can see, you can optionally assign an initial value to the variable. The variable's scope is determined by the location of its declaring section within the program.

You can have as many declare sections in a program as you like.

The declarations are also echoed to the output file as normal COBOL variables, so there's no need to declare them again. Variables that are not intended to be used in SQL commands can be declared normally outside these special sections.

The definition of a group item also must be listed inside a DECLARE section. Otherwise the preprocessor cannot handle these types since it does not know the definition.

# **D.4.3 Retrieving Query Results**

Now you should be able to pass data generated by your program into an SQL command. But how do you retrieve the results of a query? For that purpose, embedded SQL provides special variants of the usual commands SELECT and FETCH. These commands have a special INTO clause that specifies which host variables the retrieved values are to be stored in. SELECT is used for a query that returns only single row, and FETCH is used for a query that returns multiple rows, using a cursor.

Here is an example:

```
* assume this table:
* CREATE TABLE test (a int, b varchar(50));

*

EXEC SQL BEGIN DECLARE SECTION END-EXEC.
01 V1 PIC S9(9).
01 V2 PIC X(50) VARYING.
EXEC SQL END DECLARE SECTION END-EXEC.
...

EXEC SQL SELECT a, b INTO :V1, :V2 FROM test END-EXEC.
```

So the INTO clause appears between the select list and the FROM clause. The number of elements in the select list and the list after INTO (also called the target list) must be equal.

Here is an example using the command FETCH:

```
EXEC SQL BEGIN DECLARE SECTION END-EXEC.

01 V1 PIC S9(9).

01 V2 PIC X(50) VARYING.

EXEC SQL END DECLARE SECTION END-EXEC.

...

EXEC SQL DECLARE foo CURSOR FOR SELECT a, b FROM test END-EXEC.

...

PERFORM WITH

...

EXEC SQL FETCH NEXT FROM foo INTO :V1, :V2 END-EXEC
```

END-PERFORM.

Here the INTO clause appears after all the normal clauses.

# D.4.4 Type Mapping

When ECOBPG applications exchange values between the PostgreSQL server and the COBOL application, such as when retrieving query results from the server or executing SQL statements with input parameters, the values need to be converted between PostgreSQL data types and host language variable types (specifically COBOL language data types). One of the main points of ECOBPG is that it takes care of this automatically in most cases.

In this respect, there are two kinds of data types: Some simple PostgreSQL data types, such as integer and text, can be read and written by the application directly. Other PostgreSQL data types, such as timestamp and date can only be accessed through character strings. special library functions does not exist in ecobpg. (pgtypes, exists in ECPG, for COBOL is not implemented yet)

"Table D.1 Mapping Between PostgreSQL Data Types and COBOL Variable Types" shows which PostgreSQL data types correspond to which COBOL data types. When you wish to send or receive a value of a given PostgreSQL data type, you should declare a COBOL variable of the corresponding COBOL data type in the declare section.

Table D.1 Mapping Between PostgreSQL Data Types and COBOL Variable Types

PostgreSQL data type	COBOL Host variable type
smallint	PIC S9([1-4]) {BINARY COMP COMP-5}
integer	PIC S9([5-9]) {BINARY COMP COMP-5}
bigint	PIC S9([10-18]) {BINARY COMP COMP-5}
decimal	PIC S9(m)V9(n) PACKED-DECIMAL PIC 9(m)V9(n) DISPLAY (*1) PIC S9(m)V9(n) DISPLAY PIC S9(m)V9(n) DISPLAY SIGN TRAILING [SEPARATE] PIC S9(m)V9(n) DISPLAY SIGN LEADING [SEPARATE]
numeric	(same with decimal)
real	COMP-1
double precision	COMP-2
small serial	PIC S9([1-4]) {BINARY COMP COMP-5}
serial	PIC S9([1-9]) {BINARY COMP COMP-5}
bigserial	PIC S9([10-18]) {BINARY COMP COMP-5}
oid	PIC 9(9) {BINARY COMP COMP-5}
character(n), varchar(n), text	PIC X(n), PIC X(n) VARYING
name	PIC X(NAMEDATALEN)
boolean	BOOL(*2)
other types(e.g. timestamp)	PIC X(n), PIC X(n) VARYING

<sup>\*1:</sup> If no USAGE is specified, host variable is regarded as DISPLAY.

Body of BOOL is PIC X(1). '1' for true and '0' for false.

You can use some pattern of digits for integer(see table), but if database sends big number with more digits than specified, behavior is undefined.

<sup>\*2:</sup> Type definition is added automatically on pre-compiling.

### PostgreSQL data type

#### **COBOL** Host variable type

VALUE clause can't be used with VARYING. (Can be used with other types)

REDEFINE clause can be used, but it won't be validated on pre-compilation (The COBOL compiler will do this).

# **Handling Character Strings**

To handle SQL character string data types, such as varchar and text, there is a possible way to declare the host variables.

The way is using the PIC X(n) VARYING type (we call it VARCHAR type from now on), which is a special type provided by ECOBPG. The definition on type VARCHAR is converted into a group item consists of named variables. A declaration like:

```
01 VAR PIC X(180) VARYING.
```

is converted into:

```
01 VAR.
49 LEN PIC S9(4) COMP-5.
49 ARR PIC X(180).
```

if --varchar-with-named-member option is used, it is converted into:

```
01 VAR.
49 VAR-LEN PIC S9(4) COMP-5.
49 VAR-ARR PIC X(180).
```

You can use level 1 to 48 for VARCHAR. Don't use level 49 variable right after VARCHAR variable. To use a VARCHAR host variable as an input for SQL statement, LEN must be set the length of the string included in ARR.

To use a VARCHAR host variable as an output of SQL statement, the variable must be declared in a sufficient length. If the length is insufficient, it can cause a buffer overrun.

PIC X(n) and VARCHAR host variables can also hold values of other SQL types, which will be stored in their string forms.

# **Accessing Special Data Types**

ECOBPG doesn't have special support for date, timestamp, and interval types.

(ECPG has pgtypes, but ECOBPG doesn't.)

You can use PIC X(n) or VARCHAR for DB I/O with these types. See "Data Types" section in PostgreSQL's document.

### **Host Variables with Nonprimitive Types**

As a host variable you can also use arrays, typedefs, and group items.

# Arrays

To create and use array variables, OCCURENCE syntax is provided by COBOL.

The typical use case is to retrieve multiple rows from a query result without using a cursor. Without an array, to process a query result consisting of multiple rows, it is required to use a cursor and the FETCH command. But with array host variables, multiple rows can be received at once. The length of the array has to be defined to be able to accommodate all rows, otherwise a buffer overrun will likely occur.

Following example scans the pg\_database system table and shows all OIDs and names of the available databases:

```
EXEC SQL BEGIN DECLARE SECTION END-EXEC.

01 GROUP-ITEM.

05 DBID PIC S9(9) COMP OCCURS 8.

05 DBNAME PIC X(16) OCCURS 8.

01 I PIC S9(9) COMP.

EXEC SQL END DECLARE SECTION END-EXEC.

EXEC SQL CONNECT TO testdb END-EXEC.
```

```
* Retrieve multiple rows into arrays at once.

EXEC SQL SELECT oid, datname INTO :DBID, :DBNAME FROM pg_database END-EXEC.

PERFORM VARYING I FROM 1 BY 1 UNTIL I > 8

DISPLAY "oid=" DBID(I) ", dbname=" DBNAME(I)

END-PERFORM.

EXEC SQL COMMIT END-EXEC.

EXEC SQL DISCONNECT ALL END-EXEC.
```

You can use member of array as simple host variable by specifying subscript of array. For specifying subscript, use C-style "[1]", not COBOL-style "(1)". But subscript starts with 1, according to COBOL syntax.

```
EXEC SQL BEGIN DECLARE SECTION END-EXEC.

01 GROUP-ITEM.

05 DBID PIC S9(9) COMP OCCURS 8.

EXEC SQL END DECLARE SECTION END-EXEC.

EXEC SQL CONNECT TO testdb END-EXEC.

EXEC SQL SELECT oid INTO :DBID[1] FROM pg_database WHERE oid=1 END-EXEC.

DISPLAY "oid=" DBID(1)

EXEC SQL COMMIT END-EXEC.

EXEC SQL DISCONNECT ALL END-EXEC.
```

### Group Item

A group item whose subordinate item names match the column names of a query result, can be used to retrieve multiple columns at once. The group item enables handling multiple column values in a single host variable.

The following example retrieves OIDs, names, and sizes of the available databases from the pg\_database system table by using the pg\_database\_size() function. In this example, a group item variable dbinfo\_t with members whose names match each column in the SELECT result is used to retrieve one result row without putting multiple host variables in the FETCH statement.

```
EXEC SOL BEGIN DECLARE SECTION END-EXEC.
    01 DBINFO-T TYPEDEF.
       02 OID PIC S9(9) COMP.
       02 DATNAME PIC X(65).
       02 DBSIZE PIC S9(18) COMP.
    01 DBVAL TYPE DBINFO-T.
EXEC SQL END DECLARE SECTION END-EXEC.
   EXEC SQL DECLARE curl CURSOR FOR SELECT oid, datname, pg_database_size(oid) AS size
FROM pg_database END-EXEC.
   EXEC SQL OPEN curl END-EXEC.
    when end of result set reached, break out of loop
    EXEC SQL WHENEVER NOT FOUND GOTO END-FETCH END-EXEC.
    PERFORM NO LIMIT
        Fetch multiple columns into one structure.
        EXEC SQL FETCH FROM curl INTO :DBVAL END-EXEC
         Print members of the structure.
        DISPLAY "oid=" OID ", datname=" DATNAME ", size=" DBSIZE
    END-PERFORM.
```

```
END-FETCH.
EXEC SQL CLOSE curl END-EXEC.
```

group item host variables "absorb" as many columns as the group item as subordinate items. Additional columns can be assigned to other host variables. For example, the above program could also be restructured like this, with the size variable outside the group item:

```
EXEC SQL BEGIN DECLARE SECTION END-EXEC.
   01 DBINFO-T TYPEDEF.
       02 OID PIC S9(9) COMP.
       02 DATNAME PIC X(65).
   01 DBVAL TYPE DBINFO-T.
   01 DBSIZE PIC S9(18) COMP.
EXEC SQL END DECLARE SECTION END-EXEC.
   EXEC SQL DECLARE curl CURSOR FOR SELECT oid, datname, pg_database_size(oid) AS size
FROM pg_database END-EXEC.
   EXEC SQL OPEN curl END-EXEC.
   when end of result set reached, break out of loop
   EXEC SQL WHENEVER NOT FOUND GOTO END-FETCH END-EXEC.
   PERFORM NO LIMIT
       Fetch multiple columns into one structure.
       EXEC SQL FETCH FROM curl INTO :DBVAL, :DBSIZE END-EXEC
       Print members of the structure.
       DISPLAY "oid=" OID ", datname=" DATNAME ", size=" DBSIZE
   END-PERFORM
   FETCH-END.
   EXEC SQL CLOSE curl END-EXEC.
```

You can use only non-nested group items for host variable of SQL statement. Declaration of nested group items are OK, but you must specify non-nested part of group items for SQL. (VARCHAR, is translated to group item on pre-compilation, is not considered as offense of this rule.) When using inner item of group item in SQL, use C-struct like period separated syntax(not COBOL's A OF B). Here is example.

```
EXEC SQL BEGIN DECLARE SECTION END-EXEC.

01 NESTED-GROUP.

02 CHILD1.

03 A PIC X(10).

03 B PIC S9(9) COMP.

02 CHILD2.

03 A PIC X(10).

03 B PIC S9(9) COMP.

EXEC SQL END DECLARE SECTION END-EXEC.

* This SQL is valid. CHILD1 has no nested group items.

EXEC SQL SELECT * INTO :NESTED-GROUP.CHILD1 FROM TABLE1 END-EXEC.
```

For specifying basic item of group items, full specification is not needed if the specification is enough for identifying the item. This is from COBOL syntax. For more detail, see resources of COBOL syntax.

#### **TYPEDEF**

Use the typedef keyword to map new types to already existing types.

```
EXEC SQL BEGIN DECLARE SECTION END-EXEC.

01 MYCHARTYPE TYPEDEF PIC X(40).

01 SERIAL-T TYPEDEF PIC S9(9) COMP.

EXEC SQL END DECLARE SECTION END-EXEC.
```

Note that you could also use:

```
EXEC SQL TYPE SERIAL-T IS PIC S9(9) COMP-5. END-EXEC.
```

This declaration does not need to be part of a declare section.

# D.4.5 Handling Nonprimitive SQL Data Types

This section contains information on how to handle nonscalar and user-defined SQL-level data types in ECOBPG applications. Note that this is distinct from the handling of host variables of nonprimitive types, described in the previous section.

## **Arrays**

SQL-level arrays are not directly supported in ECOBPG. It is not possible to simply map an SQL array into a COBOL array host variable. This will result in undefined behavior. Some workarounds exist, however.

If a query accesses elements of an array separately, then this avoids the use of arrays in ECOBPG. Then, a host variable with a type that can be mapped to the element type should be used. For example, if a column type is array of integer, a host variable of type PIC S9(9) COMP can be used. Also if the element type is varchar or text, a host variable of type VARCHAR can be used.

Here is an example. Assume the following table:

The following example program retrieves the 4th element of the array and stores it into a host variable of type PIC S9(9) COMP-5:

```
EXEC SQL BEGIN DECLARE SECTION END-EXEC.

01 II PIC S9(9) COMP.

EXEC SQL END DECLARE SECTION END-EXEC.

EXEC SQL DECLARE curl CURSOR FOR SELECT ii[4] FROM t3 END-EXEC.

EXEC SQL OPEN curl END-EXEC.

EXEC SQL WHENEVER NOT FOUND GOTO END-FETCH END-EXEC.

PERFORM NO LIMIT

EXEC SQL FETCH FROM curl INTO :II END-EXEC

DISPLAY "ii=" II
END-PERFORM.
```

```
END-FETCH.
EXEC SQL CLOSE curl END-EXEC.
```

To map multiple array elements to the multiple elements in an array type host variables each element of array column and each element of the host variable array have to be managed separately, for example:

```
EXEC SQL BEGIN DECLARE SECTION END-EXEC.

01 GROUP-ITEM.

05 II_A PIC S9(9) COMP OCCURS 8.

EXEC SQL END DECLARE SECTION END-EXEC.

EXEC SQL DECLARE curl CURSOR FOR SELECT ii[1], ii[2], ii[3], ii[4] FROM t3 END-EXEC.

EXEC SQL OPEN curl END-EXEC.

EXEC SQL WHENEVER NOT FOUND GOTO END-FETCH END-EXEC.

PERFORM NO LIMIT

EXEC SQL FETCH FROM curl INTO :II_A[1], :II_A[2], :II_A[3], :II_A[4] END-EXEC

...

END-PERFORM.
```

#### Note again that.

```
EXEC SQL BEGIN DECLARE SECTION END-EXEC.

01 GROUP-ITEM.

05 II_A PIC S9(9) COMP OCCURS 8.

EXEC SQL END DECLARE SECTION END-EXEC.

EXEC SQL DECLARE curl CURSOR FOR SELECT ii FROM t3 END-EXEC.

EXEC SQL OPEN curl END-EXEC.

EXEC SQL WHENEVER NOT FOUND GOTO END-FETCH END-EXEC.

PERFORM NO LIMIT

* WRONG

EXEC SQL FETCH FROM curl INTO :II_A END-EXEC

...

END-PERFORM.
```

would not work correctly in this case, because you cannot map an array type column to an array host variable directly.

Another workaround is to store arrays in their external string representation in host variables of type VARCHAR. For more details about this representation.



#### See

Refer to "Arrays" in "Tutorial" in the PostgreSQL Documentation for information more details about this representation.

Note that this means that the array cannot be accessed naturally as an array in the host program (without further processing that parses the text representation).

# **Composite Types**

Composite types are not directly supported in ECOBPG, but an easy workaround is possible. The available workarounds are similar to the ones described for arrays above: Either access each attribute separately or use the external string representation.

For the following examples, assume the following type and table:

```
CREATE TYPE comp_t AS (intval integer, textval varchar(32));
CREATE TABLE t4 (compval comp_t);
INSERT INTO t4 VALUES ( (256, 'PostgreSQL') );
```

The most obvious solution is to access each attribute separately. The following program retrieves data from the example table by selecting each attribute of the type comp\_t separately:

```
EXEC SQL BEGIN DECLARE SECTION END-EXEC.
01 INTVAL PIC S9(9) COMP.
01 TEXTVAL PIC X(33) VARYING.
EXEC SQL END DECLARE SECTION END-EXEC.
* Put each element of the composite type column in the SELECT list.
EXEC SQL DECLARE curl CURSOR FOR SELECT (compval).intval, (compval).textval FROM t4 END-
EXEC.
EXEC SQL OPEN curl END-EXEC.
EXEC SQL WHENEVER NOT FOUND GOTO END-FETCH END-EXEC.
PERFORM NO LIMIT
   Fetch each element of the composite type column into host variables.
   EXEC SQL FETCH FROM curl INTO :INTVAL, :TEXTVAL END-EXEC
   DISPLAY "intval=" INTVAL ", textval=" ARR OF TEXTVAL
END-PERFORM.
END-FETCH.
EXEC SQL CLOSE curl END-EXEC.
```

To enhance this example, the host variables to store values in the FETCH command can be gathered into one group item. For more details about the host variable in the group item form, see "Group Item". To switch to the group item, the example can be modified as below. The two host variables, intval and textval, become subordinate items of the comp\_t group item, and the group item is specified on the FETCH command.

```
EXEC SQL BEGIN DECLARE SECTION END-EXEC.
01 COMP-T TYPEDEF.
   02 INTVAL PIC S9(9) COMP.
   02 TEXTVAL PIC X(33) VARYING.
01 COMPVAL TYPE COMP-T.
EXEC SQL END DECLARE SECTION END-EXEC.
* Put each element of the composite type column in the SELECT list.
EXEC SQL DECLARE curl CURSOR FOR SELECT (compval).intval, (compval).textval FROM t4 END-
EXEC.
EXEC SQL OPEN curl END-EXEC.
EXEC SQL WHENEVER NOT FOUND GOTO END-FETCH END-EXEC.
PERFORM NO LIMIT
   Put all values in the SELECT list into one structure.
   EXEC SQL FETCH FROM curl INTO : COMPVAL END-EXEC
   DISPLAY "intval=" INTVAL ", textval=" ARR OF TEXTVAL
END-PERFORM.
END-FETCH.
EXEC SQL CLOSE curl END-EXEC.
```

Although a group item is used in the FETCH command, the attribute names in the SELECT clause are specified one by one. This can be enhanced by using a \* to ask for all attributes of the composite type value.

```
EXEC SQL DECLARE curl CURSOR FOR SELECT (compval).* FROM t4 END-EXEC.

EXEC SQL OPEN curl END-EXEC.

EXEC SQL WHENEVER NOT FOUND GOTO END-FETCH END-EXEC.

PERFORM NO LIMIT
```

```
* Put all values in the SELECT list into one structure.

EXEC SQL FETCH FROM curl INTO :COMPVAL END-EXEC

DISPLAY "intval=" INTVAL ", textval=" ARR OF TEXTVAL

END-PERFORM.
```

This way, composite types can be mapped into structures almost seamlessly, even though ECOBPG does not understand the composite type itself.

Finally, it is also possible to store composite type values in their external string representation in host variables of type VARCHAR. But that way, it is not easily possible to access the fields of the value from the host program.

# **User-defined Base Types**

New user-defined base types are not directly supported by ECOBPG. You can use the external string representation and host variables of type VARCHAR, and this solution is indeed appropriate and sufficient for many types.

Here is an example using the data type complex.



Refer to "User-defined Types" in "Server Programming" in the PostgreSQL Documentation for information on the data type complex.

The external string representation of that type is (%lf,%lf), which is defined in the functions complex\_in() and complex\_out() functions. The following example inserts the complex type values (1,1) and (3,3) into the columns a and b, and select them from the table after that.

```
EXEC SQL BEGIN DECLARE SECTION END-EXEC.

01 A PIC X(64) VARYING.

01 B PIC X(64) VARYING.

EXEC SQL END DECLARE SECTION END-EXEC.

EXEC SQL INSERT INTO test_complex VALUES ('(1,1)', '(3,3)') END-EXEC.

EXEC SQL DECLARE curl CURSOR FOR SELECT a, b FROM test_complex END-EXEC.

EXEC SQL OPEN curl END-EXEC.

EXEC SQL WHENEVER NOT FOUND GOTO END-FETCH END-EXEC.

PERFORM NO LIMIT

EXEC SQL FETCH FROM curl INTO :A, :B END-EXEC

DISPLAY "a=" ARR OF A ", b=" ARR OF B

END-PERFORM.

END-FETCH.

EXEC SQL CLOSE curl END-EXEC.
```

Another workaround is avoiding the direct use of the user-defined types in ECOBPG and instead create a function or cast that converts between the user-defined type and a primitive type that ECOBPG can handle. Note, however, that type casts, especially implicit ones, should be introduced into the type system very carefully.

For example:

```
CREATE FUNCTION create_complex(r double precision, i double precision) RETURNS complex LANGUAGE SQL IMMUTABLE
AS $$ SELECT $1 * complex '(1,0)' + $2 * complex '(0,1)' $$;
```

After this definition, the following:

```
EXEC SQL BEGIN DECLARE SECTION END-EXEC.
01 A COMP-2.
```

```
01 B COMP-2.
01 C COMP-2.
01 D COMP-2.
EXEC SQL END DECLARE SECTION END-EXEC.

MOVE 1 TO A.
MOVE 2 TO B.
MOVE 3 TO C.
MOVE 4 TO D.

EXEC SQL INSERT INTO test_complex VALUES (create_complex(:A, :B), create_complex(:C, :D))
END-EXEC.
```

has the same effect as

```
EXEC SQL INSERT INTO test_complex VALUES ('(1,2)', '(3,4)') END-EXEC.
```

# **D.4.6 Indicators**

The examples above do not handle null values. In fact, the retrieval examples will raise an error if they fetch a null value from the database. To be able to pass null values to the database or retrieve null values from the database, you need to append a second host variable specification to each host variable that contains data. This second host variable is called the *indicator* and contains a flag that tells whether the datum is null, in which case the value of the real host variable is ignored. Here is an example that handles the retrieval of null values correctly:

```
EXEC SQL BEGIN DECLARE SECTION END-EXEC.

01 VAL PIC X(50) VARYING.

01 VAL_IND PIC S9(9) COMP-5.

EXEC SQL END DECLARE SECTION END-EXEC.

...

EXEC SQL SELECT b INTO :VAL :VAL_IND FROM test1 END-EXEC.
```

The indicator variable val\_ind will be zero if the value was not null, and it will be negative if the value was null.

The indicator has another function: if the indicator value is positive, it means that the value is not null, but it was truncated when it was stored in the host variable.

# D.5 Dynamic SQL

In many cases, the particular SQL statements that an application has to execute are known at the time the application is written. In some cases, however, the SQL statements are composed at run time or provided by an external source. In these cases you cannot embed the SQL statements directly into the COBOL source code, but there is a facility that allows you to call arbitrary SQL statements that you provide in a string variable.

# D.5.1 Executing Statements without a Result Set

The simplest way to execute an arbitrary SQL statement is to use the command EXECUTE IMMEDIATE. For example:

```
EXEC SQL BEGIN DECLARE SECTION END-EXEC.

01 STMT PIC X(30) VARYING.

EXEC SQL END DECLARE SECTION END-EXEC.

MOVE "CREATE TABLE test1 (...);" TO ARR OF STMT.

COMPUTE LEN OF STMT = FUNCTION STORED-CHAR-LENGTH (ARR OF STMT).

EXEC SQL EXECUTE IMMEDIATE :STMT END-EXEC.
```

EXECUTE IMMEDIATE can be used for SQL statements that do not return a result set (e.g., DDL, INSERT, UPDATE, DELETE). You cannot execute statements that retrieve data (e.g., SELECT) this way. The next section describes how to do that.

# D.5.2 Executing a Statement with Input Parameters

A more powerful way to execute arbitrary SQL statements is to prepare them once and execute the prepared statement as often as you like. It is also possible to prepare a generalized version of a statement and then execute specific versions of it by substituting parameters. When preparing the statement, write question marks where you want to substitute parameters later. For example:

```
EXEC SQL BEGIN DECLARE SECTION END-EXEC.

01 STMT PIC X(40) VARYING.

EXEC SQL END DECLARE SECTION END-EXEC.

MOVE "INSERT INTO test1 VALUES(?, ?);" TO ARR OF STMT.

COMPUTE LEN OF STMT = FUNCTION STORED-CHAR-LENGTH (ARR OF STMT).

EXEC SQL PREPARE MYSTMT FROM :STMT END-EXEC.

...

EXEC SQL EXECUTE MYSTMT USING 42, 'foobar' END-EXEC.
```

When you don't need the prepared statement anymore, you should deallocate it:

```
EXEC SQL DEALLOCATE PREPARE name END-EXEC.
```

# D.5.3 Executing a Statement with a Result Set

To execute an SQL statement with a single result row, EXECUTE can be used. To save the result, add an INTO clause.

```
EXEC SQL BEGIN DECLARE SECTION END-EXEC.

01 STMT PIC X(50) VARYING.

01 V1 PIC S9(9) COMP.

01 V2 PIC S9(9) COMP.

01 V3 PIC X(50) VARYING.

EXEC SQL END DECLARE SECTION END-EXEC.

MOVE "SELECT a, b, c FROM test1 WHERE a > ?" TO ARR OF STMT.

COMPUTE LEN OF STMT = FUNCTION STORED-CHAR-LENGTH (ARR OF STMT).

EXEC SQL PREPARE MYSTMT FROM :STMT END-EXEC.

...

EXEC SQL EXECUTE MYSTMT INTO :V1, :V2, :V3 USING 37 END-EXEC.
```

An EXECUTE command can have an INTO clause, a USING clause, both, or neither.

If a query is expected to return more than one result row, a cursor should be used, as in the following example. (See "D.3.2 Using Cursors" for more details about the cursor.)

```
PERFORM NO LIMIT

EXEC SQL FETCH cursor1 INTO :DBANAME,:DATNAME END-EXEC

DISPLAY "dbaname=" ARR OF DBANAME ", datname=" ARR OF DATNAME
END-PERFORM.

FETCH-END.

EXEC SQL CLOSE cursor1 END-EXEC.

EXEC SQL COMMIT END-EXEC.

EXEC SQL DISCONNECT ALL END-EXEC.
```

# **D.6 Using Descriptor Areas**

An SQL descriptor area is a more sophisticated method for processing the result of a SELECT, FETCH or a DESCRIBE statement. An SQL descriptor area groups the data of one row of data together with metadata items into one data group item. The metadata is particularly useful when executing dynamic SQL statements, where the nature of the result columns might not be known ahead of time. PostgreSQL provides a way to use Descriptor Areas: the named SQL Descriptor Areas.

# D.6.1 Named SQL Descriptor Areas

A named SQL descriptor area consists of a header, which contains information concerning the entire descriptor, and one or more item descriptor areas, which basically each describe one column in the result row.

Before you can use an SQL descriptor area, you need to allocate one:

```
EXEC SQL ALLOCATE DESCRIPTOR identifier END-EXEC.
```

The identifier serves as the "variable name" of the descriptor area. When you don't need the descriptor anymore, you should deallocate it:

```
EXEC SQL DEALLOCATE DESCRIPTOR identifier END-EXEC.
```

To use a descriptor area, specify it as the storage target in an INTO clause, instead of listing host variables:

```
EXEC SQL FETCH NEXT FROM mycursor INTO SQL DESCRIPTOR mydesc END-EXEC.
```

If the result set is empty, the Descriptor Area will still contain the metadata from the query, i.e. the field names.

For not yet executed prepared queries, the DESCRIBE statement can be used to get the metadata of the result set:

```
EXEC SQL BEGIN DECLARE SECTION END-EXEC.

01 SQL-STMT PIC X(30) VARYING.

EXEC SQL END DECLARE SECTION END-EXEC.

MOVE "SELECT * FROM table1" TO ARR OF SQL-STMT.

COMPUTE LEN OF SQL-STMT = FUNCTION STORED-CHAR-LENGTH ( ARR OF SQL-STMT ) .

EXEC SQL PREPARE STMT1 FROM :SQL-STMT END-EXEC.

EXEC SQL DESCRIBE STMT1 INTO SQL DESCRIPTOR MYDESC END-EXEC.
```

Before PostgreSQL 9.0, the SQL keyword was optional, so using DESCRIPTOR and SQL DESCRIPTOR produced named SQL Descriptor Areas. Now it is mandatory, omitting the SQL keyword is regarded as the syntax that produces SQLDA Descriptor Areas. However, ecobpg does not support SQLDA and it causes an error.

In DESCRIBE and FETCH statements, the INTO and USING keywords can be used to similarly: they produce the result set and the metadata in a Descriptor Area.

Now how do you get the data out of the descriptor area? You can think of the descriptor area as a group item with named fields. To retrieve the value of a field from the header and store it into a host variable, use the following command:

```
EXEC SQL GET DESCRIPTOR name :hostvar = field END-EXEC.
```

Currently, there is only one header field defined: COUNT, which tells how many item descriptor areas exist (that is, how many columns are contained in the result). The host variable needs to be of an integer type as PIC S9(9) COMP-5. To get a field from the item descriptor area, use the following command:

EXEC SQL GET DESCRIPTOR name VALUE num :hostvar = field END-EXEC.

num can be a host variable containing an integer as PIC S9(9) COMP-5.

hostvar must be PIC S9(9) COMP-5 if type of the field is integer. Possible fields are:

CARDINALITY (integer)

number of rows in the result set

DATA

actual data item (therefore, the data type of this field depends on the query)

DATETIME\_INTERVAL\_CODE (integer)

When TYPE is 9, DATETIME\_INTERVAL\_CODE will have a value of 1 for DATE, 2 for TIME, 3 for TIMESTAMP, 4 for TIME WITH TIME ZONE, or 5 for TIMESTAMP WITH TIME ZONE.

DATETIME\_INTERVAL\_PRECISION (integer)

not implemented

INDICATOR (integer)

the indicator (indicating a null value or a value truncation)

KEY\_MEMBER (integer)

not implemented

LENGTH (integer)

length of the datum in characters

NAME (string)

name of the column

NULLABLE (integer)

not implemented

OCTET\_LENGTH (integer)

length of the character representation of the datum in bytes

PRECISION (integer)

precision (for type numeric)

RETURNED\_LENGTH (integer)

length of the datum in characters

RETURNED\_OCTET\_LENGTH (integer)

length of the character representation of the datum in bytes

SCALE (integer)

scale (for type numeric)

TYPE (integer)

numeric code of the data type of the column

In EXECUTE, DECLARE and OPEN statements, the effect of the INTO and USING keywords are different. A Descriptor Area can also be manually built to provide the input parameters for a query or a cursor and USING SQL DESCRIPTOR name is the way to pass the input parameters into a parametrized query. The statement to build a named SQL Descriptor Area is below:

```
EXEC SQL SET DESCRIPTOR name VALUE num field = :hostvar END-EXEC.
```

PostgreSQL supports retrieving more than one record in one FETCH statement and storing the data in host variables in this case assumes that the variable is an array. E.g.:

```
EXEC SQL BEGIN DECLARE SECTION END-EXEC.

01 GROUP-ITEM.

05 IDNUM PIC S9(9) COMP OCCURS 5.

EXEC SQL END DECLARE SECTION END-EXEC.

EXEC SQL FETCH 5 FROM mycursor INTO SQL DESCRIPTOR mydesc END-EXEC.

EXEC SQL GET DESCRIPTOR mydesc VALUE 1 :IDNUM = DATA END-EXEC.
```

# **D.7 Error Handling**

This section describes how you can handle exceptional conditions and warnings in an embedded SQL program. There are two nonexclusive facilities for this.

- Callbacks can be configured to handle warning and error conditions using the WHENEVER command.
- Detailed information about the error or warning can be obtained from the sqlca variable.

# **D.7.1 Setting Callbacks**

One simple method to catch errors and warnings is to set a specific action to be executed whenever a particular condition occurs. In general:

```
EXEC SQL WHENEVER condition action END-EXEC.
```

condition can be one of the following:

# **SQLERROR**

The specified action is called whenever an error occurs during the execution of an SQL statement.

#### **SQLWARNING**

The specified action is called whenever a warning occurs during the execution of an SQL statement.

### NOT FOUND

The specified action is called whenever an SQL statement retrieves or affects zero rows. (This condition is not an error, but you might be interested in handling it specially.)

action can be one of the following:

# CONTINUE

This effectively means that the condition is ignored. This is the default.

#### **GOTO** label

### GO TO label

Jump to the specified label (using a COBOL goto statement).

#### **SQLPRINT**

Print a message to standard error. This is useful for simple programs or during prototyping. The details of the message cannot be configured.

# STOP

Call STOP, which will terminate the program.

# CALL name usingargs DO name usingargs

Call the specified functions with the following characters including arguments. Thus, syntaxes (including compiler depending) are able to be placed as well as the arguments. Though, there are some limitation as following:

- You can't use RETURNING, ON EXCEPTION or OVER FLOW clauses.
- In the called subprogram, you must specify CONTINUE for every action with WHENEVER statement.

The SQL standard only provides for the actions CONTINUE and GOTO (and GO TO).

Here is an example that you might want to use in a simple program. It prints a simple message when a warning occurs and aborts the program when an error happens:

```
EXEC SQL WHENEVER SQLWARNING SQLPRINT END-EXEC.
EXEC SQL WHENEVER SQLERROR STOP END-EXEC.
```

The statement EXEC SQL WHENEVER is a directive of the SQL preprocessor, not a COBOL statement. The error or warning actions that it sets apply to all embedded SQL statements that appear below the point where the handler is set, unless a different action was set for the same condition between the first EXEC SQL WHENEVER and the SQL statement causing the condition, regardless of the flow of control in the COBOL program. So neither of the two following COBOL program excerpts will have the desired effect:

```
* WRONG

* 
...

IF VERBOSE = 1 THEN

EXEC SQL WHENEVER SQLWARNING SQLPRINT END-EXEC

END-IF.

...

EXEC SQL SELECT ... END-EXEC.

...

* 
WRONG

* 
...

CALL SET-ERROR-HANDLER.

* 
(and execute "EXEC SQL WHENEVER SQLERROR STOP" in SET-ERROR-HANDLER)

...

EXEC SQL SELECT ... END-EXEC.

...
```

# D.7.2 sqlca

For more powerful error handling, the embedded SQL interface provides a global variable with the name sqlca (SQL communication area) that has the following group item:

(In a multithreaded program, every thread automatically gets its own copy of sqlca. This works similarly to the handling of the standard C global variable errno.)

sqlca covers both warnings and errors. If multiple warnings or errors occur during the execution of a statement, then sqlca will only contain information about the last one.

If no error occurred in the last SQL statement, SQLCODE will be 0 and SQLSTATE will be "00000". If a warning or error occurred, then SQLCODE will be negative and SQLSTATE will be different from "00000". A positive SQLCODE indicates a harmless condition, such as that the last query returned zero rows. SQLCODE and SQLSTATE are two different error code schemes; details appear below.

If the last SQL statement was successful, then SQLERRD(2) contains the OID of the processed row, if applicable, and SQLERRD(3) contains the number of processed or returned rows, if applicable to the command.

In case of an error or warning, SQLERRMC will contain a string that describes the error. The field SQLERRML contains the length of the error message that is stored in SQLERRMC (the result of FUNCTION STORED-CHAR-LENGTH. Note that some messages are too long to fit in the fixed-size sqlerrmc array; they will be truncated.

In case of a warning, the 3rd character of SQLWARN is set to W. (In all other cases, it is set to something different from W.) If the 2nd character of SQLWARN is set to W, then a value was truncated when it was stored in a host variable. The 1st character of SQLWARN is set to W if any of the other elements are set to indicate a warning.

The fields sqlcaid, sqlcabc, sqlerrp, and the remaining elements of sqlerrd and sqlwarn currently contain no useful information.

The structure sqlca is not defined in the SQL standard, but is implemented in several other SQL database systems. The definitions are similar at the core, but if you want to write portable applications, then you should investigate the different implementations carefully.

Here is one example that combines the use of WHENEVER and sqlca, printing out the contents of sqlca when an error occurs. This is perhaps useful for debugging or prototyping applications, before installing a more "user-friendly" error handler.

```
EXEC SQL WHENEVER SQLERROR GOTO PRINT_SQLCA END-EXEC.

PRINT_SQLCA.

DISPLAY "==== sqlca ====".

DISPLAY "SQLCODE: "SQLCODE.

DISPLAY "SQLERRML: "SQLERRML.

DISPLAY "SQLERRMC: "SQLERRMC.

DISPLAY "SQLERRD: "SQLERRMC.

DISPLAY "SQLERRD: "SQLERRD(1) " "SQLERRD(2) " "SQLERRD(3)" "SQLERRD(4) " "

SQLERRD(5) " "SQLERRD(6).

DISPLAY "SQLSTATE: "SQLSTATE.

DISPLAY "====================".
```

The result could look as follows (here an error due to a misspelled table name):

# D.7.3 SQLSTATE vs. SQLCODE

The fields SQLSTATE and SQLCODE are two different schemes that provide error codes. Both are derived from the SQL standard, but SQLCODE has been marked deprecated in the SQL-92 edition of the standard and has been dropped in later editions. Therefore, new applications are strongly encouraged to use SQLSTATE.

SQLSTATE is a five-character array. The five characters contain digits or upper-case letters that represent codes of various error and warning conditions. SQLSTATE has a hierarchical scheme: the first two characters indicate the general class of the condition, the last three characters indicate a subclass of the general condition. A successful state is indicated by the code 00000. The SQLSTATE codes are for the most part defined in the SQL standard. The PostgreSQL server natively supports

SQLSTATE error codes; therefore a high degree of consistency can be achieved by using this error code scheme throughout all applications.



#### See

Refer to "PostgreSQL Error Codes" in "Appendixes" in the PostgreSQL Documentation for further information.

SQLCODE, the deprecated error code scheme, is a simple integer. A value of 0 indicates success, a positive value indicates success with additional information, and a negative value indicates an error. The SQL standard only defines the positive value +100, which indicates that the last command returned or affected zero rows, and no specific negative values. Therefore, this scheme can only achieve poor portability and does not have a hierarchical code assignment. Historically, the embedded SQL processor for PostgreSQL has assigned some specific SQLCODE values for its use, which are listed below with their numeric value and their symbolic name. Remember that these are not portable to other SQL implementations. To simplify the porting of applications to the SQLSTATE scheme, the corresponding SQLSTATE is also listed. There is, however, no one-to-one or one-to-many mapping between the two schemes (indeed it is many-to-many), so you should consult the global SQLSTATE in each case.



#### Sec

Refer to "PostgreSQL Error Codes" in "Appendixes" in the PostgreSQL Documentation.

These are the assigned SQLCODE values:

0

Indicates no error. (SQLSTATE 00000)

100

This is a harmless condition indicating that the last command retrieved or processed zero rows, or that you are at the end of the cursor. (SQLSTATE 02000)

When processing a cursor in a loop, you could use this code as a way to detect when to abort the loop, like this:

```
PERFORM NO LIMIT

EXEC SQL FETCH ... END-EXEC

IF SQLCODE = 100 THEN

GO TO FETCH-END

END-IF

END-PERFORM.
```

But WHENEVER NOT FOUND GOTO ... effectively does this internally, so there is usually no advantage in writing this out explicitly.

-12

Indicates that your virtual memory is exhausted. The numeric value is defined as -ENOMEM. (SQLSTATE YE001)

-200

Indicates the preprocessor has generated something that the library does not know about. Perhaps you are running incompatible versions of the preprocessor and the library. (SQLSTATE YE002)

-201

This means that the command specified more host variables than the command expected. (SQLSTATE 07001 or 07002)

-202

This means that the command specified fewer host variables than the command expected. (SQLSTATE~07001~or~07002)

-203

This means a query has returned multiple rows but the statement was only prepared to store one result row (for example, because the specified variables are not arrays). (SQLSTATE 21000)

#### -204

The host variable is of type signed int and the datum in the database is of a different type and contains a value that cannot be interpreted as a signed int. The library uses strtol() for this conversion. (SQLSTATE 42804)

#### -205

The host variable is of type unsigned int and the datum in the database is of a different type and contains a value that cannot be interpreted as an unsigned int. The library uses strtoul() for this conversion. (SQLSTATE 42804)

#### -206

The host variable is of type float and the datum in the database is of another type and contains a value that cannot be interpreted as a float. The library uses strtod() for this conversion. (SQLSTATE 42804)

#### -207

The host variable is of type DECIMAL and the datum in the database is of another type and contains a value that cannot be interpreted as a DECIMAL or DISPLAY value. For the case of DISPLAY, this error happens if values in the database is too large for converting to DISPLAY value. (SQLSTATE 42804)

#### -208

The host variable is of type interval and the datum in the database is of another type and contains a value that cannot be interpreted as an interval value. (SQLSTATE 42804)

#### -209

The host variable is of type date and the datum in the database is of another type and contains a value that cannot be interpreted as a date value. (SQLSTATE 42804)

#### -210

The host variable is of type timestamp and the datum in the database is of another type and contains a value that cannot be interpreted as a timestamp value. (SQLSTATE 42804)

#### -211

This means the host variable is of type bool and the datum in the database is neither 't' nor 'f'. (SQLSTATE 42804)

#### -212

The statement sent to the PostgreSQL server was empty. (This cannot normally happen in an embedded SQL program, so it might point to an internal error.) (SQLSTATE YE002)

#### -213

A null value was returned and no null indicator variable was supplied. (SQLSTATE 22002)

#### -214

An ordinary variable was used in a place that requires an array. (SQLSTATE 42804)

#### -215

The database returned an ordinary variable in a place that requires array value. (SQLSTATE 42804)

#### -220

The program tried to access a connection that does not exist. (SQLSTATE 08003)

## -221

The program tried to access a connection that does exist but is not open. (This is an internal error.) (SQLSTATE YE002)

#### -230

The statement you are trying to use has not been prepared. (SQLSTATE 26000)

## -240

The descriptor specified was not found. The statement you are trying to use has not been prepared. (SQLSTATE 33000)

#### -241

The descriptor index specified was out of range. (SQLSTATE 07009)

-242

An invalid descriptor item was requested. (This is an internal error.) (SQLSTATE YE002)

-243

During the execution of a dynamic statement, the database returned a numeric value and the host variable was not numeric. (SQLSTATE 07006)

-244

During the execution of a dynamic statement, the database returned a non-numeric value and the host variable was numeric. (SQLSTATE 07006)

-400

Some error caused by the PostgreSQL server. The message contains the error message from the PostgreSQL server.

-401

The PostgreSQL server signaled that we cannot start, commit, or rollback the transaction. (SQLSTATE 08007)

-402

The connection attempt to the database did not succeed. (SQLSTATE 08001)

-403

Duplicate key error, violation of unique constraint. (SQLSTATE 23505)

-404

A result for the subquery is not single row. (SQLSTATE 21000)

-602

An invalid cursor name was specified. (SQLSTATE 34000)

-603

Transaction is in progress. (SQLSTATE 25001)

-604

There is no active (in-progress) transaction. (SQLSTATE 25P01)

-605

An existing cursor name was specified. (SQLSTATE 42P03)

# **D.8 Preprocessor Directives**

Several preprocessor directives are available that modify how the ecobpg preprocessor parses and processes a file.

# D.8.1 Including Files

To include an external file into your embedded SQL program, use:

```
EXEC SQL INCLUDE filename END-EXEC.

EXEC SQL INCLUDE <filename> END-EXEC.

EXEC SQL INCLUDE "filename" END-EXEC.
```

The embedded SQL preprocessor will look for a file named filename.pco, preprocess it, and include it in the resulting COBOL output. Thus, embedded SQL statements in the included file are handled correctly.

By default, the ecobpg preprocessor will search a file at the current directory. This behavior can be changed by the ecobpg commandline option.

First, the preprocessor tries to locate a file by specified file name at the current directory. If it fails and the file name does not end with .pco, the preprocessor also tries to locate a file with the suffix at the same directory.

The difference between EXEC SQL INCLUDE and COPY statement is whether precompiler processes embedded SQLs in the file, or not. If the file contains embedded SQLs, use EXEC SQL INCLUDE.



The include file name is case-sensitive, even though the rest of the EXEC SQL INCLUDE command follows the normal SQL case-sensitivity rules.

# D.8.2 The define and undef Directives

Similar to the directive #define that is known from C, embedded SQL has a similar concept:

```
EXEC SQL DEFINE name END-EXEC.
EXEC SQL DEFINE name value END-EXEC.
```

So you can define a name:

```
EXEC SQL DEFINE HAVE_FEATURE END-EXEC.
```

And you can also define constants:

```
EXEC SQL DEFINE MYNUMBER 12 END-EXEC.
EXEC SQL DEFINE MYSTRING 'abc' END-EXEC.
```

Use undef to remove a previous definition:

```
EXEC SQL UNDEF MYNUMBER END-EXEC.
```

Note that a constant in the SQL statement is only replaced by EXEC SQL DEFINE. The replacement may change the number of characters in a line, but ecobpg does not validate it after the replacement. Pay attention to the limitation of the number of characters in a line.

# D.8.3 ifdef, ifndef, else, elif, and endif Directives

You can use the following directives to compile code sections conditionally:

```
EXEC SQL ifdef name END-EXEC.
```

Checks a name and processes subsequent lines if name has been created with EXEC SQL define name.

```
EXEC SQL ifndef name END-EXEC.
```

Checks a name and processes subsequent lines if name has not been created with EXEC SQL define name.

```
EXEC SQL else END-EXEC.
```

Starts processing an alternative section to a section introduced by either EXEC SQL ifdef name or EXEC SQL ifndef name.

```
EXEC SQL elif name END-EXEC.
```

Checks name and starts an alternative section if name has been created with EXEC SQL define name.

```
EXEC SQL endif END-EXEC.
```

Ends an alternative section.

# Example:

```
EXEC SQL ifndef TZVAR END-EXEC.

EXEC SQL SET TIMEZONE TO 'GMT' END-EXEC.

EXEC SQL elif TZNAME END-EXEC.
```

```
EXEC SQL SET TIMEZONE TO TZNAME END-EXEC.

EXEC SQL else END-EXEC.

EXEC SQL SET TIMEZONE TO TZVAR END-EXEC.

EXEC SQL endif END-EXEC.
```

# D.9 Processing Embedded SQL Programs

Now that you have an idea how to form embedded SQL COBOL programs, you probably want to know how to compile them. Before compiling you run the file through the embedded SQL COBOL preprocessor, which converts the SQL statements you used to special function calls. After compiling, you must link with a special library that contains the needed functions. These functions fetch information from the arguments, perform the SQL command using the libpq interface, and put the result in the arguments specified for output.

The preprocessor program is called ecobpg and is included in a normal PostgreSQL installation. Embedded SQL programs are typically named with an extension .pco. If you have a program file called prog1.pco, you can preprocess it by simply calling:

```
ecobpg progl.pco
```

This will create a file called prog1.cob. If your input files do not follow the suggested naming pattern, you can specify the output file explicitly using the -o option.

The preprocessed file can be compiled normally, following the usage of the compiler.

The generated COBOL source files include library files from the PostgreSQL installation, so if you installed PostgreSQL in a location that is not searched by default, you have to add an option such as -I/usr/local/pgsql/include to the compilation command line.

To link an embedded SQL program, you need to include the libecpg library.

Again, you might have to add an option for library search like -L/usr/local/pgsql/lib to that command line.

If you manage the build process of a larger project using make, it might be convenient to include the following implicit rule to your makefiles:

The complete syntax of the ecobpg command is detailed in "D.12.1 ecobpg".

Currently, ecobpg does not support multi threading.

# **D.10 Large Objects**

Large objects are not supported by ECOBPG.

If you need to access large objects, use large objects interfaces of libpq instead.

# D.11 Embedded SQL Commands

This section describes all SQL commands that are specific to embedded SQL.

Command	Description
ALLOCATE DESCRIPTOR	Allocate an SQL descriptor area
CONNECT	Establish a database connection
DEALLOCATE DESCRIPTOR	Deallocate an SQL descriptor area
DECLARE	Define a cursor

Command	Description
DESCRIBE	Obtain information about a prepared statement or result set
DISCONNECT	Terminate a database connection
EXECUTE IMMEDIATE	Dynamically prepare and execute a statement
GET DESCRIPTOR	Get information from an SQL descriptor area
OPEN	Open a dynamic cursor
PREPARE	Prepare a statement for execution
SET AUTOCOMMIT	Set the autocommit behavior of the current session
SET CONNECTION	Select a database connection
SET DESCRIPTOR	Set information in an SQL descriptor area
ТҮРЕ	Define a new data type
VAR	Define a variable
WHENEVER	Specify the action to be taken when an SQL statement causes a specific class condition to be raised



Refer to the SQL commands listed in "SQL Commands" under "Reference" in the PostgreSQL Documentation, which can also be used in embedded SQL, unless stated otherwise.

# **D.11.1 ALLOCATE DESCRIPTOR**

# Name

ALLOCATE DESCRIPTOR -- allocate an SQL descriptor area

# **Synopsis**

ALLOCATE DESCRIPTOR name

# **Description**

ALLOCATE DESCRIPTOR allocates a new named SQL descriptor area, which can be used to exchange data between the PostgreSQL server and the host program.

Descriptor areas should be freed after use using the DEALLOCATE DESCRIPTOR command.

## **Parameters**

name

A name of SQL descriptor. This can be an SQL identifier or a host variable.

# **Examples**

EXEC SQL ALLOCATE DESCRIPTOR mydesc END-EXEC.

# Compatibility

ALLOCATE DESCRIPTOR is specified in the SQL standard.

# D.11.2 CONNECT

#### Name

CONNECT -- establish a database connection

# **Synopsis**

```
CONNECT TO connection_target [ AS connection_name ] [ USER connection_user_name ]
CONNECT TO DEFAULT
CONNECT connection user name
DATABASE connection_target
```

# **Description**

The CONNECT command establishes a connection between the client and the PostgreSQL server.

## **Parameters**

```
connection_target
  connection_target specifies the target server of the connection on one of several forms.
  [database_name][@host][:port]
      Connect over TCP/IP
   unix:postgresql://host [:port]/[database_name][?connection_option]
      Connect over Unix-domain sockets
  tcp:postgresql://host [:port]/[database_name][?connection_option]
      Connect over TCP/IP
   SQL string constant
```

containing a value in one of the above forms

host variable

host variable of fixed-length string (trailing spaces are ignored) containing a value in one of the above forms

connection\_name

An optional identifier for the connection, so that it can be referred to in other commands. This can be an SQL identifier or a host variable.

```
connection_user_name
```

The user name for the database connection.

This parameter can also specify user name and password, using one the forms user\_name/password, user\_name IDENTIFIED BY password, or user\_name USING password.

User name and password can be SQL identifiers, string constants, or host variables (fixed-length string, trailing spaces are ignored).

# **DEFAULT**

Use all default connection parameters, as defined by libpq.

### **Examples**

Here a several variants for specifying connection parameters:

```
EXEC SQL CONNECT TO "connectdb" AS main END-EXEC.
EXEC SQL CONNECT TO "connectdb" AS second END-EXEC.
EXEC SQL CONNECT TO "unix:postgresql://localhost/connectdb" AS main USER connectuser END-
EXEC.
EXEC SQL CONNECT TO 'connectdb' AS main END-EXEC.
EXEC SQL CONNECT TO 'unix:postgresql://localhost/connectdb' AS main USER :user END-EXEC.
EXEC SQL CONNECT TO :dbn AS :idt END-EXEC.
EXEC SQL CONNECT TO :dbn USER connectuser USING :pw END-EXEC.
EXEC SQL CONNECT TO @localhost AS main USER connectdb END-EXEC.
EXEC SQL CONNECT TO REGRESSDB1 as main END-EXEC.
EXEC SQL CONNECT TO connectdb AS :idt END-EXEC.
EXEC SQL CONNECT TO connectdb AS main USER connectuser/connectdb END-EXEC.
EXEC SQL CONNECT TO connectdb AS main END-EXEC.
EXEC SQL CONNECT TO connectdb@localhost AS main END-EXEC.
EXEC SQL CONNECT TO tcp:postgresql://localhost/ USER connectdb END-EXEC.
EXEC SQL CONNECT TO tcp:postgresql://localhost/connectdb USER connectuser IDENTIFIED BY
connectpw END-EXEC.
EXEC SQL CONNECT TO tcp:postgresql://localhost:20/connectdb USER connectuser IDENTIFIED BY
connectpw END-EXEC.
EXEC SQL CONNECT TO unix:postgresql://localhost/ AS main USER connectdb END-EXEC.
EXEC SQL CONNECT TO unix:postgresql://localhost/connectdb AS main USER connectuser END-
EXEC.
EXEC SQL CONNECT TO unix:postgresql://localhost/connectdb USER connectuser IDENTIFIED BY
"connectpw" END-EXEC.
EXEC SQL CONNECT TO unix:postgresql://localhost/connectdb USER connectuser USING
"connectpw" END-EXEC.
EXEC SQL CONNECT TO unix:postgresql://localhost/connectdb?connect_timeout=14 USER
connectuser END-EXEC.
```

Here is an example program that illustrates the use of host variables to specify connection parameters:

```
EXEC SQL BEGIN DECLARE SECTION END-EXEC.
   database name
   01 DBNAME PIC X(6).
   connection user name
   01 USER PIC X(8).
   connection string
   01 CONNECTION PIC X(38).
    01 VER PIC X(256).
EXEC SQL END DECLARE SECTION END-EXEC.
   MOVE "testdb" TO DBNAME.
   MOVE "testuser" TO USER.
   {\tt MOVE "tcp:postgresql://localhost:5432/testdb" TO CONNECTION.}
    EXEC SQL CONNECT TO :DBNAME USER :USER END-EXEC.
    EXEC SQL SELECT version() INTO :VER END-EXEC.
    EXEC SQL DISCONNECT END-EXEC.
    DISPLAY "version: " VER.
    EXEC SQL CONNECT TO :CONNECTION USER :USER END-EXEC.
    EXEC SQL SELECT version() INTO :VER END-EXEC.
    EXEC SQL DISCONNECT END-EXEC.
    DISPLAY "version: " VER.
```

# Compatibility

CONNECT is specified in the SQL standard, but the format of the connection parameters is implementation-specific.

#### See Also

DISCONNECT, SET CONNECTION

# **D.11.3 DEALLOCATE DESCRIPTOR**

#### Name

DEALLOCATE DESCRIPTOR -- deallocate an SQL descriptor area

# **Synopsis**

DEALLOCATE DESCRIPTOR name

# **Description**

DEALLOCATE DESCRIPTOR deallocates a named SQL descriptor area.

#### **Parameters**

name

The name of the descriptor which is going to be deallocated. This can be an SQL identifier or a host variable.

## **Examples**

EXEC SQL DEALLOCATE DESCRIPTOR mydesc END-EXEC.

## Compatibility

DEALLOCATE DESCRIPTOR is specified in the SQL standard.

## See Also

ALLOCATE DESCRIPTOR, GET DESCRIPTOR, SET DESCRIPTOR

# D.11.4 DECLARE

# Name

DECLARE -- define a cursor

# **Synopsis**

```
DECLARE cursor_name [ BINARY ] [ INSENSITIVE ] [ [ NO ] SCROLL ] CURSOR [ { WITH | WITHOUT } HOLD ] FOR prepared_name

DECLARE cursor_name [ BINARY ] [ INSENSITIVE ] [ [ NO ] SCROLL ] CURSOR [ { WITH | WITHOUT } HOLD ] FOR query
```

#### Description

DECLARE declares a cursor for iterating over the result set of a prepared statement. This command has slightly different semantics from the direct SQL command DECLARE: Whereas the latter executes a query and prepares the result set for retrieval, this embedded SQL command merely declares a name as a "loop variable" for iterating over the result set of a query; the actual execution happens when the cursor is opened with the OPEN command.

#### **Parameters**

cursor\_name

A cursor name. This can be an SQL identifier or a host variable.

prepared name

The name of a prepared query, either as an SQL identifier or a host variable.

query

A SELECT or VALUES command which will provide the rows to be returned by the cursor.

For the meaning of the cursor options, see DECLARE.



Refer to "SQL Commands" in "Reference" in the PostgreSQL Documentation for information on the SELECT, VALUES and DECLARE command.

# **Examples**

Examples declaring a cursor for a query:

```
EXEC SQL DECLARE C CURSOR FOR SELECT * FROM My_Table END-EXEC.

EXEC SQL DECLARE C CURSOR FOR SELECT Item1 FROM T END-EXEC.

EXEC SQL DECLARE cur1 CURSOR FOR SELECT version() END-EXEC.
```

An example declaring a cursor for a prepared statement:

```
EXEC SQL PREPARE stmt1 AS SELECT version() END-EXEC.
EXEC SQL DECLARE cur1 CURSOR FOR stmt1 END-EXEC.
```

# Compatibility

DECLARE is specified in the SQL standard.

#### See Also

OPEN, CLOSE, DECLARE



See

Refer to "SQL Commands" in "Reference" in the PostgreSQL Documentation for information on the CLOSE and DECLARE command.

# D.11.5 DESCRIBE

#### Name

DESCRIBE -- obtain information about a prepared statement or result set

# **Synopsis**

```
DESCRIBE [ OUTPUT ] prepared_name USING SQL DESCRIPTOR descriptor_name
DESCRIBE [ OUTPUT ] prepared_name INTO SQL DESCRIPTOR descriptor_name
```

# Description

DESCRIBE retrieves metadata information about the result columns contained in a prepared statement, without actually fetching a row.

#### **Parameters**

prepared\_name

The name of a prepared statement. This can be an SQL identifier or a host variable.

descriptor\_name

A descriptor name. It can be an SQL identifier or a host variable.

# **Examples**

```
EXEC SQL ALLOCATE DESCRIPTOR mydesc END-EXEC.

EXEC SQL PREPARE stmt1 FROM :sql_stmt END-EXEC.

EXEC SQL DESCRIBE stmt1 INTO SQL DESCRIPTOR mydesc END-EXEC.

EXEC SQL GET DESCRIPTOR mydesc VALUE 1 :charvar = NAME END-EXEC.

EXEC SQL DEALLOCATE DESCRIPTOR mydesc END-EXEC.
```

# Compatibility

DESCRIBE is specified in the SQL standard.

#### See Also

ALLOCATE DESCRIPTOR, GET DESCRIPTOR

# **D.11.6 DISCONNECT**

# Name

DISCONNECT -- terminate a database connection

# **Synopsis**

```
DISCONNECT connection_name

DISCONNECT [ CURRENT ]

DISCONNECT DEFAULT

DISCONNECT ALL
```

# **Description**

DISCONNECT closes a connection (or all connections) to the database.

## **Parameters**

connection\_name

A database connection name established by the CONNECT command.

#### **CURRENT**

Close the "current" connection, which is either the most recently opened connection, or the connection set by the SET CONNECTION command. This is also the default if no argument is given to the DISCONNECT command.

#### **DEFAULT**

Close the default connection.

#### ALL

Close all open connections.

# **Examples**

```
EXEC SQL CONNECT TO testdb AS DEFAULT USER testuser END-EXEC.

EXEC SQL CONNECT TO testdb AS con1 USER testuser END-EXEC.

EXEC SQL CONNECT TO testdb AS con2 USER testuser END-EXEC.

EXEC SQL CONNECT TO testdb AS con3 USER testuser END-EXEC.

* close con3

EXEC SQL DISCONNECT CURRENT END-EXEC.

* close DEFAULT

EXEC SQL DISCONNECT DEFAULT END-EXEC.

* close con2 and con1

EXEC SQL DISCONNECT ALL END-EXEC.
```

# Compatibility

DISCONNECT is specified in the SQL standard.

#### See Also

CONNECT, SET CONNECTION

# **D.11.7 EXECUTE IMMEDIATE**

#### Name

EXECUTE IMMEDIATE -- dynamically prepare and execute a statement

# **Synopsis**

EXECUTE IMMEDIATE string

# **Description**

EXECUTE IMMEDIATE immediately prepares and executes a dynamically specified SQL statement, without retrieving result rows.

## **Parameters**

string

A literal string or a host variable containing the SQL statement to be executed.

# **Examples**

Here is an example that executes an INSERT statement using EXECUTE IMMEDIATE and a host variable named command:

```
MOVE "INSERT INTO test (name, amount, letter) VALUES ('db: ''rl''', 1, 'f')" TO ARR OF cmd. COMPUTE LEN OF cmd = FUNCTION STORED-CHAR-LENGTH(ARR OF cmd). EXEC SQL EXECUTE IMMEDIATE :cmd END-EXEC.
```

#### Compatibility

EXECUTE IMMEDIATE is specified in the SQL standard.

# D.11.8 GET DESCRIPTOR

#### Name

GET DESCRIPTOR -- get information from an SQL descriptor area

# **Synopsis**

```
GET DESCRIPTOR descriptor_name :hostvariable = descriptor_header_item [, ...]

GET DESCRIPTOR descriptor_name VALUE column_number :hostvariable = descriptor_item [, ...]
```

#### Description

GET DESCRIPTOR retrieves information about a query result set from an SQL descriptor area and stores it into host variables. A descriptor area is typically populated using FETCH or SELECT before using this command to transfer the information into host language variables.

This command has two forms: The first form retrieves descriptor "header" items, which apply to the result set in its entirety. One example is the row count. The second form, which requires the column number as additional parameter, retrieves information about a particular column. Examples are the column name and the actual column value.

#### **Parameters**

descriptor\_name

A descriptor name.

descriptor\_header\_item

A token identifying which header information item to retrieve. Only COUNT, to get the number of columns in the result set, is currently supported.

column\_number

The number of the column about which information is to be retrieved. The count starts at 1.

descriptor\_item

A token identifying which item of information about a column to retrieve. See Section 33.7.1 for a list of supported items.

hostvariable

A host variable that will receive the data retrieved from the descriptor area.

#### **Examples**

An example to retrieve the number of columns in a result set:

```
EXEC SQL GET DESCRIPTOR d :d_count = COUNT END-EXEC.
```

An example to retrieve a data length in the first column:

```
EXEC SQL GET DESCRIPTOR d VALUE 1 :d_returned_octet_length = RETURNED_OCTET_LENGTH END-EXEC.
```

An example to retrieve the data body of the second column as a string:

```
EXEC SQL GET DESCRIPTOR d VALUE 2 :d_data = DATA END-EXEC.
```

Here is an example for a whole procedure of executing SELECT current\_database(); and showing the number of columns, the column data length, and the column data:

```
EXEC SQL BEGIN DECLARE SECTION END-EXEC.

01 D-COUNT PIC S9(9) COMP-5.

01 D-DATA PIC X(1024).
```

```
01 D-RETURNED-OCTET-LENGTH PIC S9(9) COMP.
EXEC SQL END DECLARE SECTION END-EXEC.
   EXEC SQL CONNECT TO testdb AS con1 USER testuser END-EXEC.
   EXEC SQL ALLOCATE DESCRIPTOR d END-EXEC.
   Declare, open a cursor, and assign a descriptor to the cursor
   EXEC SQL DECLARE cur CURSOR FOR SELECT current_database() END-EXEC.
   EXEC SQL OPEN cur END-EXEC.
   EXEC SQL FETCH NEXT FROM cur INTO SQL DESCRIPTOR d END-EXEC.
   Get a number of total columns
   EXEC SQL GET DESCRIPTOR d :D-COUNT = COUNT END-EXEC.
   DISPLAY "d_count
   Get length of a returned column
   EXEC SQL GET DESCRIPTOR d VALUE 1 :D-RETURNED-OCTET-LENGTH = RETURNED_OCTET_LENGTH END-
EXEC.
   DISPLAY "d_returned_octet_length = " D-RETURNED-OCTET-LENGTH.
  Fetch the returned column as a string
   EXEC SQL GET DESCRIPTOR d VALUE 1 :D-DATA = DATA END-EXEC.
   DISPLAY "d_data
                                    = " D-DATA.
  Closing
   EXEC SQL CLOSE cur END-EXEC.
   EXEC SQL COMMIT END-EXEC.
   EXEC SQL DEALLOCATE DESCRIPTOR d END-EXEC.
   EXEC SQL DISCONNECT ALL END-EXEC.
```

## When the example is executed, the result will look like this:

#### Compatibility

GET DESCRIPTOR is specified in the SQL standard.

#### See Also

ALLOCATE DESCRIPTOR, SET DESCRIPTOR

# **D.11.9 OPEN**

#### Name

OPEN -- open a dynamic cursor

#### **Synopsis**

```
OPEN cursor_name

OPEN cursor_name USING value [, ...]

OPEN cursor_name USING SQL DESCRIPTOR descriptor_name
```

# Description

OPEN opens a cursor and optionally binds actual values to the placeholders in the cursor's declaration. The cursor must previously have been declared with the DECLARE command. The execution of OPEN causes the query to start executing on the server.

#### **Parameters**

cursor\_name

The name of the cursor to be opened. This can be an SQL identifier or a host variable.

#### value

A value to be bound to a placeholder in the cursor. This can be an SQL constant, a host variable, or a host variable with indicator.

descriptor\_name

The name of a descriptor containing values to be bound to the placeholders in the cursor. This can be an SQL identifier or a host variable.

## **Examples**

```
EXEC SQL OPEN a END-EXEC.

EXEC SQL OPEN d USING 1, 'test' END-EXEC.

EXEC SQL OPEN c1 USING SQL DESCRIPTOR mydesc END-EXEC.

EXEC SQL OPEN :curname1 END-EXEC.
```

## Compatibility

OPEN is specified in the SQL standard.

# See Also

**DECLARE**, CLOSE



See

Refer to "SQL Commands" in "Reference" in the PostgreSQL Documentation for information on the CLOSE command.

# D.11.10 PREPARE

#### Name

PREPARE -- prepare a statement for execution

# **Synopsis**

PREPARE name FROM string

#### Description

PREPARE prepares a statement dynamically specified as a string for execution. This is different from the direct SQL statement PREPARE, which can also be used in embedded programs. The EXECUTE command is used to execute either kind of prepared statement.

# **Parameters**

prepared\_name

An identifier for the prepared query.

# string

A literal string or a host variable containing a preparable statement, one of the SELECT, INSERT, UPDATE, or DELETE.

## **Examples**

```
MOVE "SELECT * FROM test1 WHERE a = ? AND b = ?" TO ARR OF STMT.

COMPUTE LEN OF STMT = FUNCTION STORED-CHAR-LENGTH (ARR OF STMT).

EXEC SQL ALLOCATE DESCRIPTOR indesc END-EXEC.

EXEC SQL ALLOCATE DESCRIPTOR outdesc END-EXEC.

EXEC SQL PREPARE foo FROM :STMT END-EXEC.

EXEC SQL EXECUTE foo USING SQL DESCRIPTOR indesc INTO SQL DESCRIPTOR outdesc END-EXEC.
```

# Compatibility

PREPARE is specified in the SQL standard.

#### See Also

**EXECUTE** 



Refer to "SQL Commands" in "Reference" in the PostgreSQL Documentation for information on the EXECUTE command.

......

# **D.11.11 SET AUTOCOMMIT**

## Name

SET AUTOCOMMIT -- set the autocommit behavior of the current session

#### **Synopsis**

```
SET AUTOCOMMIT \{ = | TO \} \{ ON | OFF \}
```

#### **Description**

SET AUTOCOMMIT sets the autocommit behavior of the current database session. By default, embedded SQL programs are not in autocommit mode, so COMMIT needs to be issued explicitly when desired. This command can change the session to autocommit mode, where each individual statement is committed implicitly.

## Compatibility

SET AUTOCOMMIT is an extension of PostgreSQL ECOBPG.

# **D.11.12 SET CONNECTION**

#### Name

SET CONNECTION -- select a database connection

# **Synopsis**

```
SET CONNECTION [ TO | = ] connection_name
```

# **Description**

SET CONNECTION sets the "current" database connection, which is the one that all commands use unless overridden.

#### **Parameters**

connection\_name

A database connection name established by the CONNECT command.

**DEFAULT** 

Set the connection to the default connection.

# **Examples**

```
EXEC SQL SET CONNECTION TO con2 END-EXEC.

EXEC SQL SET CONNECTION = con1 END-EXEC.
```

# Compatibility

SET CONNECTION is specified in the SQL standard.

#### See Also

CONNECT, DISCONNECT

# **D.11.13 SET DESCRIPTOR**

## Name

SET DESCRIPTOR -- set information in an SQL descriptor area

# **Synopsis**

```
SET DESCRIPTOR descriptor_name descriptor_header_item = value [, ...]

SET DESCRIPTOR descriptor_name VALUE number descriptor_item = value [, ...]
```

# **Description**

SET DESCRIPTOR populates an SQL descriptor area with values. The descriptor area is then typically used to bind parameters in a prepared query execution.

This command has two forms: The first form applies to the descriptor "header", which is independent of a particular datum. The second form assigns values to particular datums, identified by number.

# **Parameters**

descriptor\_name

A descriptor name.

descriptor\_header\_item

A token identifying which header information item to set. Only COUNT, to set the number of descriptor items, is currently supported.

number

The number of the descriptor item to set. The count starts at 1.

descriptor\_item

A token identifying which item of information to set in the descriptor. See Section 33.7.1 for a list of supported items.

value

A value to store into the descriptor item. This can be an SQL constant or a host variable.

# **Examples**

```
EXEC SQL SET DESCRIPTOR indesc COUNT = 1 END-EXEC.

EXEC SQL SET DESCRIPTOR indesc VALUE 1 DATA = 2 END-EXEC.

EXEC SQL SET DESCRIPTOR indesc VALUE 1 DATA = :val1 END-EXEC.

EXEC SQL SET DESCRIPTOR indesc VALUE 2 DATA = 'some string', INDICATOR = :val1 END-EXEC.

EXEC SQL SET DESCRIPTOR indesc VALUE 2 INDICATOR = :val2null, DATA = :val2 END-EXEC.
```

# Compatibility

SET DESCRIPTOR is specified in the SQL standard.

#### See Also

ALLOCATE DESCRIPTOR, GET DESCRIPTOR

# **D.11.14 TYPE**

#### Name

TYPE -- define a new data type

# **Synopsis**

```
TYPE type_name IS ctype
```

# Description

The TYPE command defines a new COBOL type. It is equivalent to putting a typedef into a declare section.

This command is only recognized when ecobpgpg is run with the -c option.

A level number of 01 is automatically added to type\_name item. Thus, the level number must not to be specified externally. To define a group item, a level number needs to be specified to the each subordinate items.

For reasons of internal implementation, "TYPE" must be placed just after "EXEC SQL", without containing newline. For other place, you can use newline.

## **Parameters**

type\_name

The name for the new type. It must be a valid COBOL type name.

ctype

A COBOL type specification (including expression format specification).

# **Examples**

```
EXEC SQL TYPE CUSTOMER IS

02 NAME PIC X(50) VARYING.

02 PHONE PIC S9(9) COMP. END-EXEC.

EXEC SQL TYPE CUST-IND IS

02 NAME_IND PIC S9(4) COMP.

02 PHONE_IND PIC S9(4) COMP. END-EXEC.

EXEC SQL TYPE INTARRAY IS

02 INT PIC S9(9) OCCURS 20. END-EXEC.

EXEC SQL TYPE STR IS PIC X(50) VARYING. END-EXEC.

EXEC SQL TYPE STRING IS PIC X(10). END-EXEC.
```

Here is an example program that uses EXEC SQL TYPE:

```
EXEC SQL TYPE TT IS
    02 V PIC X(256) VARYING.
    02 I PIC S9(9) COMP. END-EXEC.
EXEC SOL TYPE TT-IND IS
    02 V-IND PIC S9(4) COMP.
    02 I-IND PIC S9(4) COMP. END-EXEC.
EXEC SQL BEGIN DECLARE SECTION END-EXEC.
   01 T TYPE TT.
   01 T-IND TYPE TT-IND.
EXEC SQL END DECLARE SECTION END-EXEC.
   EXEC SQL CONNECT TO testdb AS con1 END-EXEC.
   EXEC SQL SELECT current_database(), 256 INTO :T :T-IND LIMIT 1 END-EXEC.
   DISPLAY "t.v = " ARR OF V OF T.
   DISPLAY "t.i = " I OF T.
   DISPLAY "t_ind.v_ind = " V-IND OF T-IND.
   DISPLAY "t_ind.i_ind = " I-IND OF T-IND.
   EXEC SQL DISCONNECT con1 END-EXEC.
```

## Compatibility

The TYPE command is a PostgreSQL extension.

# **D.11.15 VAR**

#### Name

VAR-define a variable

# **Synopsis**

VAR varname IS ctype

# **Description**

The VAR command defines a host variable. It is equivalent to an ordinary COBOL variable definition inside a declare section.

When translating, a level number 01 is added. Thus, the level number must not to be specified externally.

To define a group item, a level number needs to be specified to the each subordinate items.

For reasons of internal implementation, "VAR" must be placed just after "EXEC SQL", without containing newline. For other place, you can use newline.

## **Parameters**

varname

A COBOL variable name.

ctype

A COBOL type specification.

# **Examples**

```
EXEC SQL VAR VC IS PIC X(10) VARYING. END-EXEC.
EXEC SQL VAR BOOL-VAR IS BOOL. END-EXEC.
```

# Compatibility

The VAR command is a PostgreSQL extension.

# D.11.16 WHENEVER

#### Name

WHENEVER -- specify the action to be taken when an SQL statement causes a specific class condition to be raised

# **Synopsis**

```
WHENEVER { NOT FOUND | SQLERROR | SQLWARNING } action
```

# Description

Define a behavior which is called on the special cases (Rows not found, SQL warnings or errors) in the result of SQL execution.

#### **Parameters**

See Section "D.7.1 Setting Callbacks" or a description of the parameters.

## **Examples**

```
EXEC SQL WHENEVER NOT FOUND CONTINUE END-EXEC.

EXEC SQL WHENEVER SQLWARNING SQLPRINT END-EXEC.

EXEC SQL WHENEVER SQLWARNING DO "warn" END-EXEC.

EXEC SQL WHENEVER SQLERROR sqlprint END-EXEC.

EXEC SQL WHENEVER SQLERROR CALL "print2" END-EXEC.

EXEC SQL WHENEVER SQLERROR DO handle_error USING "select" END-EXEC.

EXEC SQL WHENEVER SQLERROR DO sqlnotice USING 0 1 END-EXEC.

EXEC SQL WHENEVER SQLERROR DO "sqlprint" END-EXEC.

EXEC SQL WHENEVER SQLERROR GOTO error_label END-EXEC.

EXEC SQL WHENEVER SQLERROR STOP END-EXEC.
```

A typical application is the use of WHENEVER NOT FOUND GOTO to handle looping through result sets:

```
EXEC SQL CONNECT TO testdb AS conl END-EXEC.

EXEC SQL ALLOCATE DESCRIPTOR d END-EXEC.

EXEC SQL DECLARE cur CURSOR FOR SELECT current_database(), 'hoge', 256 END-EXEC.

EXEC SQL OPEN cur END-EXEC.

* when end of result set reached, break out of while loop

EXEC SQL WHENEVER NOT FOUND GOTO NOTFOUND END-EXEC.

PERFORM NO LIMIT

EXEC SQL FETCH NEXT FROM cur INTO SQL DESCRIPTOR d END-EXEC

...

END-PERFORM.

NOTFOUND.

EXEC SQL CLOSE cur END-EXEC.

EXEC SQL COMMIT END-EXEC.
```

```
EXEC SQL DEALLOCATE DESCRIPTOR d END-EXEC. EXEC SQL DISCONNECT ALL END-EXEC.
```

# Compatibility

WHENEVER is specified in the SQL standard, but most of the actions are PostgreSQL extensions.

# **D.12 PostgreSQL Client Applications**

This part contains reference information for PostgreSQL client applications and utilities. Not all of these commands are of general utility; some might require special privileges. The common feature of these applications is that they can be run on any host, independent of where the database server resides.

When specified on the command line, user and database names have their case preserved—the presence of spaces or special characters might require quoting. Table names and other identifiers do not have their case preserved, except where documented, and might require quoting.

# D.12.1 ecobpg

#### Name

ecobpg -- embedded SQL COBOL preprocessor

## **Synopsis**

ecobpg [option...] file...

# **Description**

ecobpg is the embedded SQL preprocessor for COBOL programs. It converts COBOL programs with embedded SQL statements to normal COBOL code by replacing the SQL invocations with special function calls. The output files can then be processed with any COBOL compiler tool chain.

ecobpg will convert each input file given on the command line to the corresponding COBOL output file. Input files preferably have the extension .pco, in which case the extension will be replaced by .cob to determine the output file name. If the extension of the input file is not .pco, then the output file name is computed by appending .cob to the full file name. The output file name can also be overridden using the -o option.

#### **Options**

ecobpg accepts the following command-line arguments:

-c

Automatically generate certain COBOL code from SQL code. Currently, this works for EXEC SQL TYPE.

#### -I directory

Specify an additional include path, used to find files included via EXEC SQL INCLUDE. Defaults are: (current directory), /usr/local/include, the PostgreSQL include directory which is defined at compile time (default: /usr/local/pgsql/include), and /usr/include, in that order.

## -o filename

Specifies that ecobpg should write all its output to the given filename.

#### -f format

Specifies the COBOL code notation. For "format", specify either of the following. If omitted, "fixed" is used.

#### fixed

Specifies fixed format notation. Up to 72 columns can be specified for area B. Characters in column 73 and beyond are deleted in the precompiled source.

#### variable

Specifies variable format notation. Up to 251 columns can be specified for area B. Characters in column 252 and beyond are deleted in the precompiled source.

#### -r option

Selects run-time behavior. Option can be one of the following:

#### prepare

Prepare all statements before using them. Libecpg will keep a cache of prepared statements and reuse a statement if it gets executed again. If the cache runs full, libecpg will free the least used statement.

#### questionmarks

Allow question mark as placeholder for compatibility reasons. This used to be the default long ago.

-t

Turn on autocommit of transactions. In this mode, each SQL command is automatically committed unless it is inside an explicit transaction block. In the default mode, commands are committed only when EXEC SQL COMMIT is issued.

#### --varchar-with-named-member

When converting VARCHAR host variable, adding name of the variable to members as prefix. Instead of LEN and ARR, (varname)-ARR and (varname)-LEN will be used.

#### -E encode

Specify the COBOL source encoding: "UTF8", "SJIS", or "EUC\_JP".

If this option is omitted, the encoding is processed based on the locale.

-V

Print additional information including the version and the "include" path.

#### --version

Print the ecobpg version and exit.

-?

--help

Show help about ecobpg command line arguments, and exit.

#### **Notes**

When compiling the preprocessed COBOL code files, the compiler needs to be able to find the library files in the PostgreSQL include directory.

Programs using COBOL code with embedded SQL have to be linked against the libecpg library, for example using the linker options.

The value of either of these directories that is appropriate for the installation can be found out using pg\_config.



See

Refer to "pg\_config" in "Reference" in the PostgreSQL Documentation.

# **Examples**

If you have an embedded SQL COBOL source file named prog1.pco, you can create an executable program using the following command:

ecobpg progl.pco

# Appendix E Quantitative Limits

This appendix lists the quantitative limits of FUJITSU Enterprise Postgres.

Table E.1 Length of identifier

Item	Limit
Database name	Up to 63 bytes (*1) (*2)
Schema name	Up to 63 bytes (*1) (*2)
Table name	Up to 63 bytes (*1) (*2)
View name	Up to 63 bytes (*1) (*2)
Index name	Up to 63 bytes (*1) (*2)
Table space name	Up to 63 bytes (*1) (*2)
Cursor name	Up to 63 bytes (*1) (*2)
Function name	Up to 63 bytes (*1) (*2)
Aggregate function name	Up to 63 bytes (*1) (*2)
Trigger name	Up to 63 bytes (*1) (*2)
Constraint name	Up to 63 bytes (*1) (*2)
Conversion name	Up to 63 bytes (*1) (*2)
Role name	Up to 63 bytes (*1) (*2)
Cast name	Up to 63 bytes (*1) (*2)
Collation sequence name	Up to 63 bytes (*1) (*2)
Encoding method conversion name	Up to 63 bytes (*1) (*2)
Domain name	Up to 63 bytes (*1) (*2)
Extension name	Up to 63 bytes (*1) (*2)
Operator name	Up to 63 bytes (*1) (*2)
Operator class name	Up to 63 bytes (*1) (*2)
Operator family name	Up to 63 bytes (*1) (*2)
Rewrite rule name	Up to 63 bytes (*1) (*2)
Sequence name	Up to 63 bytes (*1) (*2)
Text search settings name	Up to 63 bytes (*1) (*2)
Text search dictionary name	Up to 63 bytes (*1) (*2)
Text search parser name	Up to 63 bytes (*1) (*2)
Text search template name	Up to 63 bytes (*1) (*2)
	i
Data type name	Up to 63 bytes (*1) (*2)

<sup>\*1:</sup> This is the character string byte length when converted by the server character set character code.

Table E.2 Database object

	Item	Limit
Number of databases		Less than 4,294,967,296 (*1)

<sup>\*2:</sup> If an identifier that exceeds 63 bytes in length is specified, the excess characters are truncated and it is processed.

Item	Limit
Number of schemas	Less than 4,294,967,296 (*1)
Number of tables	Less than 4,294,967,296 (*1)
Number of views	Less than 4,294,967,296 (*1)
Number of indexes	Less than 4,294,967,296 (*1)
Number of table spaces	Less than 4,294,967,296 (*1)
Number of functions	Less than 4,294,967,296 (*1)
Number of aggregate functions	Less than 4,294,967,296 (*1)
Number of triggers	Less than 4,294,967,296 (*1)
Number of constraints	Less than 4,294,967,296 (*1)
Number of conversion	Less than 4,294,967,296 (*1)
Number of roles	Less than 4,294,967,296 (*1)
Number of casts	Less than 4,294,967,296 (*1)
Number of collation sequences	Less than 4,294,967,296 (*1)
Number of encoding method conversions	Less than 4,294,967,296 (*1)
Number of domains	Less than 4,294,967,296 (*1)
Number of extensions	Less than 4,294,967,296 (*1)
Number of operators	Less than 4,294,967,296 (*1)
Number of operator classes	Less than 4,294,967,296 (*1)
Number of operator families	Less than 4,294,967,296 (*1)
Number of rewrite rules	Less than 4,294,967,296 (*1)
Number of sequences	Less than 4,294,967,296 (*1)
Number of text search settings	Less than 4,294,967,296 (*1)
Number of text search dictionaries	Less than 4,294,967,296 (*1)
Number of text search parsers	Less than 4,294,967,296 (*1)
Number of text search templates	Less than 4,294,967,296 (*1)
Number of data types	Less than 4,294,967,296 (*1)
Number of enumerator type labels	Less than 4,294,967,296 (*1)
Number of default access privileges defined in the ALTER DEFAULT PRIVILEGES statement	Less than 4,294,967,296 (*1)
Number of large objects	Less than 4,294,967,296 (*1)
Number of index access methods	Less than 4,294,967,296 (*1)

<sup>\*1:</sup> The total number of all database objects must be less than 4,294,967,296.

# Table E.3 Schema element

Item	Limit
Number of columns that can be defined in one table	From 250 to 1600 (according to the data type)
Table row length	Up to 400 gigabytes
Number of columns comprising a unique constraint	Up to 32 columns
Data length comprising a unique constraint	Less than 2,000 bytes (*1) (*2)

Item	Limit
Table size	Up to one terabyte
Search condition character string length in a trigger definition statement	Up to 800 megabytes (*1) (*2)
Item size	Up to 1 gigabyte

<sup>\*1:</sup> Operation might proceed correctly even if operations are performed with a quantity outside the limits.

# Table E.4 Index

Item	Limit
Number of columns comprising a key (including VCI)	Up to 32 columns
Key length (other than VCI)	Less than 2,000 bytes (*1)

<sup>\*1:</sup> This is the character string byte length when converted by the server character set character code.

Table E.5 Data types and attributes that can be handled

Item		Limit	
Character	Data length		Data types and attributes that can be handled (*1)
	Specification length (n)		Up to 10,485,760 characters (*1)
Numeric	External decimal expression		Up to 131,072 digits before the decimal point, and up to 16,383 digits after the decimal point
	Internal binary	2 bytes	From -32,768 to 32,767
	expression	4 bytes	From -2,147,483,648 to 2,147,483,647
	8 bytes	From -9,223,372,036,854,775,808 to 9,223,372,036,854,775,807	
	Internal decimal expression		Up to 13,1072 digits before the decimal point, and up to 16,383 digits after the decimal point
Floating point expression	4 bytes	From -3.4E+38 to -7.1E-46, 0, or from 7.1E-46 to 3.4E+38	
	8 bytes	From -1.7E+308 to -2.5E-324, 0, or from 2.5E-324 to 1.7E+308	
bytea		Up to one gigabyte minus 53 bytes	
Large object		Up to two gigabytes	

<sup>\*1:</sup> This is the character string byte length when converted by the server character set character code.

# Table E.6 Function definition

Item	Limit
Number of arguments that can be specified	Up to 100
Number of variable names that can be specified in the declarations section	No limit

<sup>\*2:</sup> This is the character string byte length when converted by the server character set character code.

Item	Limit
Number of SQL statements or control statements that can be specified in a function processing implementation	No limit

Table E.7 Data operation statement			
Item	Limit		
Maximum number of connections for one process in an application (remote access)	4,000 connections		
Number of expressions that can be specified in a selection list	Up to 1,664		
Number of tables that can be specified in a FROM clause	No limit		
Number of unique expressions that can be specified in a selection list/DISTINCT clause/ORDER BY clause/GROUP BY clause within one SELECT statement	Up to 1,664		
Number of expressions that can be specified in a GROUP BY clause	No limit		
Number of expressions that can be specified in an ORDER BY clause	No limit		
Number of SELECT statements that can be specified in a UNION clause/INTERSECT clause/EXCEPT clause	Up to 4,000 (*1)		
Number of nestings in joined tables that can be specified in one view	Up to 4,000 (*1)		
Number of functions or operator expressions that can be specified in one expression	Up to 4,000 (*1)		
Number of expressions that can be specified in one row constructor	Up to 1,664		
Number of expressions that can be specified in an UPDATE statement SET clause	Up to 1,664		
Number of expressions that can be specified in one row of a VALUES list	Up to 1,664		
Number of expressions that can be specified in a RETURNING clause	Up to 1,664		
Total expression length that can be specified in the argument list of one function specification	Up to 800 megabytes (*2)		
Number of cursors that can be processed simultaneously by one session	No limit		
Character string length of one SQL statement	Up to 800 megabytes (*1) (*3)		
Number of input parameter specifications that can be specified in one dynamic SQL statement	No limit		
Number of tokens that can be specified in one SQL statement	Up to 10,000		
Number of values that can be specified as a list in a WHERE clause IN syntax	No limit		
Number of expressions that can be specified in a USING clause	No limit		
Number of JOINs that can be specified in a joined table	Up to 4,000 (*1)		
Number of expressions that can be specified in COALESCE	No limit		
Number of WHEN clauses that can be specified for CASE in a simple format or a searched format	No limit		

Item	Limit
Data size per record that can be updated or inserted by one SQL statement	Up to one gigabyte minus 53 bytes
Number of objects that can share a lock simultaneously	Up to 256,000 (*1)

<sup>\*1:</sup> Operation might proceed correctly even if operations are performed with a quantity outside the limits.

# Table E.8 Data sizes

Item	Limit
Data size per record for input data files (COPY statement, psql command \copy meta command)	Up to 800 megabytes (*1)
Data size per record for output data files (COPY statement, psql command \copy meta command)	Up to 800 megabytes (*1)

<sup>\*1:</sup> Operation might proceed correctly even if operations are performed with a quantity outside the limits.

<sup>\*2:</sup> The total number of all database objects must be less than 4,294,967,296.

<sup>\*3:</sup> This is the character string byte length when converted by the server character set character code.

# Appendix F Reference

# F.1 JDBC Driver

Refer to the Java API Reference for information on PostgreSQL JDBC driver.

# F.2 ODBC Driver

# F.2.1 List of Supported APIs

The following table shows the support status of APIs:

Function name	Support status
SQLAllocConnect	Y
SQLAllocEnv	Y
SQLAllocHandle	Y
SQLAllocStmt	Y
SQLBindCol	Y
SQLBindParameter	Y
SQLBindParam	Y
SQLBrowseConnect	Y
SQLBulkOperations	Y
SQLCancel	Y
SQLCancelHandle	N
SQLCloseCursor	Y
SQLColAttribute	Y
SQLColAttributeW	Y
SQLColAttributes	Y
SQLColAttributesW	Y
SQLColumnPrivileges	Y
SQLColumnPrivilegesW	Y
SQLColumns	Y
SQLColumnsW	Y
SQLCompleteAsync	N
SQLConnect	Y
SQLConnectW	Y
SQLCopyDesc	Y
SQLDataSources	Y
SQLDataSourcesW	Y
SQLDescribeCol	Y

Function name	Support status
SQLDescribeColW	Y
SQLDescribeParam	Y
SQLDisconnect	Y
SQLDriverConnect	Y
SQLDriverConnectW	Y
SQLDrivers	Y
SQLEndTran	Y
SQLError	Y
SQLErrorW	Y
SQLExecDirect	Y
SQLExecDirectW	Y
SQLExecute	Y
SQLExtendedFetch	Y
SQLFetch	Y
SQLFetchScroll	Y
SQLForeignKeys	Y
SQLForeignKeysW	Y
SQLFreeConnect	Y
SQLFreeEnv	Y
SQLFreeHandle	Y
SQLFreeStmt	Y
SQLGetConnectAttr	Y
SQLGetConnectAttrW	Y
SQLGetConnectOption	Y
SQLGetConnectOptionW	Y
SQLGetCursorName	Y
SQLGetCursorNameW	Y
SQLGetData	Y
SQLGetDescField	Y
SQLGetDescFieldW	Y
SQLGetDescRec	Y
SQLGetDescRecW	Y
SQLGetDiagField	Y
SQLGetDiagFieldW	Y
SQLGetDiagRec	Y
SQLGetDiagRecW	Y
SQLGetEnvAttr	Y
SQLGetFunctions	Y
SQLGetInfo	Y

SQLGetInfoW         Y           SQLGetStmtAttr         Y           SQLGetStmtOption         Y           SQLGetTypeInfo         Y           SQLGetTypeInfoW         Y           SQLMoreResults         Y           SQLMativeSql         Y           SQLNativeSqlW         Y           SQLNumParams         Y           SQLNumParams         Y           SQLParamData         Y           SQLParamOptions         Y           SQLPrepare         Y           SQLPrepare         Y           SQLPrepare         Y           SQLPrimaryKeys         Y           SQLProcedureColumns         Y           SQLProcedureColumns         Y           SQLProcedureS         Y           SQLProceduresW         Y           SQLProceduresW         Y           SQLPoceduresW         Y           SQLPoceduresW         Y           SQLPoceduresW         Y           SQLSetConnectAttr         Y           SQLSetConnectOption         Y           SQLSetConnectOption         Y           SQLSetDescRee         Y           SQLSetDescRee         Y	Function name	Support status
SQLGetStmtAttrW         Y           SQLGetStmtOption         Y           SQLGetTypeInfo         Y           SQLGetTypeInfoW         Y           SQLMoreResults         Y           SQLNativeSql         Y           SQLNativeSqlW         Y           SQLNativeSqlW         Y           SQLNumParams         Y           SQLNumParams         Y           SQLNumResultCols         Y           SQLParamData         Y           SQLParamOptions         Y           SQLParamOptions         Y           SQLPrepare         Y           SQLPrepare         Y           SQLPrepare         Y           SQLPrimary Keys         Y           SQLPrimary KeysW         Y           SQLProcedureColumns         Y           SQLProcedureColumnsW         Y           SQLProcedureS         Y           SQLProcedureS         Y           SQLProcedureS         Y           SQLProcedureSW         Y           SQLProcedureSW         Y           SQLPocedureSW         Y           SQLSetConnectAttr         Y           SQLSetConnectAttrW         Y	SQLGetInfoW	Y
SQLGetStmtOption         Y           SQLGetTypeInfo         Y           SQLMoreResults         Y           SQLNativeSql         Y           SQLNativeSqlW         Y           SQLNumParams         Y           SQLNumResultCols         Y           SQLParamData         Y           SQLPrepare         Y           SQLPrepare         Y           SQLPrepare         Y           SQLPrimaryKeys         Y           SQLPrimaryKeysW         Y           SQLProcedureColumns         Y           SQLProcedureColumnsW         Y           SQLProcedureS         Y           SQLSendoco	SQLGetStmtAttr	Y
SQL.GetTypeInfo         Y           SQL.MoreResults         Y           SQL.NativeSql         Y           SQL.NativeSqlW         Y           SQL.NumParams         Y           SQL.NumResultCols         Y           SQL.ParamOptions         Y           SQL.ParamOptions         Y           SQL.Prepare         Y           SQL.PrepareW         Y           SQL.PrimaryKeys         Y           SQL.PrimaryKeysW         Y           SQL.ProcedureColumns         Y           SQL.ProcedureColumnsW         Y           SQL.ProcedureS         Y </td <td>SQLGetStmtAttrW</td> <td>Y</td>	SQLGetStmtAttrW	Y
SQL.GetTypeInfoW         Y           SQL.MoreResults         Y           SQL.NativeSql         Y           SQL.NativeSqlW         Y           SQL.NumParams         Y           SQL.NumParams         Y           SQL.ParamData         Y           SQL.ParamOptions         Y           SQL.Prepare         Y           SQL.Prepare         Y           SQL.PrepareW         Y           SQL.PropareW         Y           SQL.ProcedureSw         Y           SQL.ProcedureColumns         Y           SQL.ProcedureColumns         Y           SQL.ProcedureS         Y	SQLGetStmtOption	Y
SQLMoreResults  SQLNativeSql  SQLNativeSqlW  Y  SQLNumParams  Y  SQLNumResultCols  Y  SQLParamData  Y  SQLPrepare  Y  SQLPrepare  Y  SQLPrimaryKeys  Y  SQLPrimaryKeys  Y  SQLProcedureColumns  Y  SQLProcedureColumns  Y  SQLProcedureS  Y  SQLProcedureS  Y  SQLProcedureS  Y  SQLProcedureS  Y  SQLProcedureS  Y  SQLPocedureS  Y  SQLPocedureS  Y  SQLPocedureS  Y  SQLPocedureS  Y  SQLProcedureS  Y  SQLSetConnectAttr  Y  SQLSetConnectOption  Y  SQLSetConnectOption  Y  SQLSetConnectOptionW  Y  SQLSetConnectOptionW  Y  SQLSetConsectOptionW  Y  SQLSetConsectOptionW  Y  SQLSetConsectOptionW  Y  SQLSetConsectOptionW  Y  SQLSetCursorName  Y  SQLSetCursorNameW  Y  SQLSetCursorNa	SQLGetTypeInfo	Y
SQLNativeSql         Y           SQLNumParams         Y           SQLNumResultCols         Y           SQLParamData         Y           SQLPramOptions         Y           SQLPrepare         Y           SQLPrepareW         Y           SQLPrimaryKeys         Y           SQLPrimaryKeysW         Y           SQLProcedureColumns         Y           SQLProcedureS         Y           SQLProcedures         Y           SQLProceduresW         Y           SQLProceduresW         Y           SQLPutData         Y           SQLRowCount         Y           SQLSetConnectAttr         Y           SQLSetConnectOption         Y           SQLSetConnectOption         Y           SQLSetConnectOptionW         Y           SQLSetConnectOptionW         Y           SQLSetDescField         Y           SQLSetDescField         Y           SQLSetDescField         Y           SQLSetEnvAttr         Y           SQLSetPos         Y           SQLSetPos         Y           SQLSetStmtAttr         Y           SQLSetStmtAttr         Y	SQLGetTypeInfoW	Y
SQLNumParams  SQLNumResultCols  SQLParamData  SQLParamOptions  Y  SQLPrepare  Y  SQLPrepare  Y  SQLPrimaryKeys  Y  SQLPrimaryKeys  Y  SQLProcedureColumns  Y  SQLProcedureColumns  Y  SQLProcedureS  Y  SQLSetConnectAttr  Y  SQLSetConnectAttr  Y  SQLSetConnectAttr  Y  SQLSetConnectOption  Y  SQLSetConnectOption  Y  SQLSetConnectOptionW  Y  SQLSetConnectOptionW  Y  SQLSetCursorName  Y  SQLSetCursorNameW  Y  SQLSetCursorNameW  Y  SQLSetDescField  Y  SQLSetEnvAttr  Y  SQLSetEnvAttr  Y  SQLSetEnvAttr  Y  SQLSetStmtAttr  Y  SQLSetStmtAttr  Y  SQLSetStmtAttrW  Y  SQLSetStmtAttrW  Y  SQLSetStmtAttrW  Y  SQLSetStmtAttrW  Y	SQLMoreResults	Y
SQLNumParams  SQLNumResultCols  Y  SQLParamData  Y  SQLPrepare  Y  SQLPrepare  Y  SQLPrepareW  Y  SQLPrimaryKeys  Y  SQLPrimaryKeysW  Y  SQLProcedureColumns  Y  SQLProcedureColumns  Y  SQLProcedureS  Y  SQLSetConnectAttr  Y  SQLSetConnectAttr  Y  SQLSetConnectAttr  Y  SQLSetConnectOption  Y  SQLSetConnectOption  Y  SQLSetConnectOptionW  Y  SQLSetConnectOptionW  Y  SQLSetCursorName  Y  SQLSetCursorNameW  Y  SQLSetCursorNameW  Y  SQLSetDescField  Y  SQLSetDescField  Y  SQLSetDescField  Y  SQLSetDescRec  Y  SQLSetDescRec  Y  SQLSetEnvAttr  Y  SQLSetEnvAttr  Y  SQLSetEnvAttr  Y  SQLSetStmtAttr  Y  SQLSetStmtAttr  Y  SQLSetStmtAttr  Y  SQLSetStmtAttrW  Y  SQLSetStmtAttrW  Y  SQLSetStmtAttrW  Y  SQLSetStmtAttrW  Y	SQLNativeSql	Y
SQLNumResultCols  SQLParamData  SQLParamOptions  Y  SQLPrepare  Y  SQLPrepare  Y  SQLPrimaryKeys  Y  SQLPrimaryKeys  Y  SQLProcedureColumns  Y  SQLProcedureS  Y  SQLProcedureS  Y  SQLProcedures  Y  SQLProcedures  Y  SQLPocedures  Y  SQLSetConnectAttr  Y  SQLSetConnectAttr  Y  SQLSetConnectOption  Y  SQLSetConnectOption  Y  SQLSetConnectOption  Y  SQLSetCursorName  Y  SQLSetCursorName  Y  SQLSetDescField  Y  SQLSetDescField  Y  SQLSetDescRec  Y  SQLSetDescRec  Y  SQLSetEnvAttr  Y  SQLSetEnvAttr  Y  SQLSetEnvAttr  Y  SQLSetScrollOptions  N  SQLSetStmtAttr  Y	SQLNativeSqlW	Y
SQLParamOptions SQLPrepare SQLPrepare Y SQLPrepare Y SQLPrimaryKeys Y SQLPrimaryKeys Y SQLProcedureColumns Y SQLProcedureS Y SQLProcedureS Y SQLProcedureS Y SQLProcedureS Y SQLProcedureS Y SQLPocedureS Y SQLSetConnectAttr Y SQLSetConnectAttr Y SQLSetConnectAttr Y SQLSetConnectOption Y SQLSetConnectOption Y SQLSetConnectOption Y SQLSetCursorName Y SQLSetCursorName Y SQLSetDescField Y SQLSetDescField Y SQLSetDescRec Y SQLSetDescRec Y SQLSetEnvAttr Y SQLSetEnvAttr Y SQLSetSorollOptions N SQLSetStmtAttr Y	SQLNumParams	Y
SQLParamOptions  SQLPrepare  SQLPrepare  Y  SQLPrepareW  Y  SQLPrimaryKeys  Y  SQLPrimaryKeysW  Y  SQLProcedureColumns  Y  SQLProcedureColumnsW  Y  SQLProcedureS  Y  SQLProcedureS  Y  SQLProcedureS  Y  SQLProcedureS  Y  SQLProcedureS  Y  SQLProcedureSW  Y  SQLPutData  Y  SQLPutData  Y  SQLSetConnectAttr  Y  SQLSetConnectAttr  Y  SQLSetConnectOption  Y  SQLSetConnectOption  Y  SQLSetConnectOptionW  Y  SQLSetCursorName  Y  SQLSetDescField  Y  SQLSetDescField  Y  SQLSetDescRec  Y  SQLSetEnvAttr  Y  SQLSetEnvAttr  Y  SQLSetEnvAttr  Y  SQLSetParam  Y  SQLSetPos  Y  SQLSetStmtAttr  Y	SQLNumResultCols	Y
SQLParamOptions  SQLPrepare  SQLPrepare  Y  SQLPrepareW  Y  SQLPrimaryKeys  Y  SQLPrimaryKeysW  Y  SQLProcedureColumns  Y  SQLProcedureColumnsW  Y  SQLProcedureS  Y  SQLProcedureS  Y  SQLProcedureS  Y  SQLProcedureS  Y  SQLProcedureS  Y  SQLProcedureSW  Y  SQLPutData  Y  SQLPutData  Y  SQLSetConnectAttr  Y  SQLSetConnectAttr  Y  SQLSetConnectOption  Y  SQLSetConnectOption  Y  SQLSetConnectOptionW  Y  SQLSetCursorName  Y  SQLSetDescField  Y  SQLSetDescField  Y  SQLSetDescRec  Y  SQLSetEnvAttr  Y  SQLSetEnvAttr  Y  SQLSetEnvAttr  Y  SQLSetParam  Y  SQLSetPos  Y  SQLSetStmtAttr  Y	SQLParamData	Y
SQLPrepare Y SQLPrepareW Y SQLPrimaryKeys Y SQLPrimaryKeysW Y SQLProcedureColumns Y SQLProcedureColumnsW Y SQLProcedureSW Y SQLProceduresW Y SQLPutData Y SQLRowCount Y SQLSetConnectAttr Y SQLSetConnectAttrW Y SQLSetConnectOption Y SQLSetConnectOptionW Y SQLSetConnectOptionW Y SQLSetConnectOptionW Y SQLSetConsectOptionW Y SQLSetSetSetConsectOptionW Y SQLSetStrutAttr Y SQLSetStrutAttr Y SQLSetStrutAttrW Y SQLSetStrutAttrW Y SQLSetStrutAttrW Y		Y
SQLPrepareW SQLPrimaryKeys SQLPrimaryKeysW SQLProcedureColumns Y SQLProcedureColumnsW SQLProcedureS Y SQLProcedureS Y SQLProceduresW Y SQLPutData Y SQLRowCount Y SQLSetConnectAttr Y SQLSetConnectAttrW Y SQLSetConnectOption Y SQLSetConnectOptionW Y SQLSetCursorName Y SQLSetCursorName Y SQLSetDescRec Y SQLSetDescRec Y SQLSetParam Y SQLSetParam Y SQLSetStonloptions N SQLSetStonloptions N SQLSetStonloptions N SQLSetStonloptions N SQLSetStonloptions N SQLSetStmtAttr Y		Y
SQLPrimaryKeys Y SQLProcedureColumns Y SQLProcedureColumnsW Y SQLProcedureS Y SQLProcedureS Y SQLProcedureSW Y SQLPutData Y SQLPutData Y SQLRowCount Y SQLSetConnectAttr Y SQLSetConnectAttrW Y SQLSetConnectOption Y SQLSetCursorName Y SQLSetCursorName Y SQLSetDescField Y SQLSetDescRec Y SQLSetParam Y SQLSetPos Y SQLSetPos Y SQLSetPos Y SQLSetStmtAttr Y SQLSetStmtAttr Y SQLSetStmtAttrW Y SQLSetStmtOption Y	· ·	Y
SQLPrimaryKeysW SQLProcedureColumns Y SQLProcedureColumnsW SQLProcedureS Y SQLProcedureS Y SQLProcedureSW Y SQLPutData Y SQLRowCount Y SQLSetConnectAttr Y SQLSetConnectAttrW Y SQLSetConnectOption Y SQLSetCursorName Y SQLSetCursorName Y SQLSetDescRec Y SQLSetDescRec Y SQLSetParam Y SQLSetParam Y SQLSetPos SQLSetStmtAttr Y		Y
SQLProcedureColumnsW SQLProcedureS Y SQLProcedureS Y SQLProcedureS Y SQLProcedureS Y SQLProceduresW Y SQLPutData Y SQLRowCount Y SQLSetConnectAttr Y SQLSetConnectAttrW Y SQLSetConnectOption Y SQLSetConnectOptionW Y SQLSetCursorName Y SQLSetCursorNameW Y SQLSetDescField Y SQLSetSetrivAttr Y SQLSetStrivAttr Y SQLSetStrivAttr Y SQLSetStrivAttr Y SQLSetStmtAttr Y		Y
SQLProcedureS         Y           SQLProcedures         Y           SQLProceduresW         Y           SQLPutData         Y           SQLRowCount         Y           SQLSetConnectAttr         Y           SQLSetConnectAttrW         Y           SQLSetConnectOption         Y           SQLSetConnectOptionW         Y           SQLSetCursorName         Y           SQLSetCursorNameW         Y           SQLSetDescField         Y           SQLSetDescField         Y           SQLSetEnvAttr         Y           SQLSetParam         Y           SQLSetPos         Y           SQLSetScrollOptions         N           SQLSetStmtAttr         Y           SQLSetStmtAttr         Y           SQLSetStmtAttrW         Y           SQLSetStmtOption         Y		Y
SQLProceduresW         Y           SQLPutData         Y           SQLRowCount         Y           SQLSetConnectAttr         Y           SQLSetConnectAttrW         Y           SQLSetConnectOption         Y           SQLSetConnectOptionW         Y           SQLSetCursorName         Y           SQLSetCursorNameW         Y           SQLSetDescField         Y           SQLSetDescField         Y           SQLSetEnvAttr         Y           SQLSetEnvAttr         Y           SQLSetPos         Y           SQLSetPos         Y           SQLSetScrollOptions         N           SQLSetStmtAttr         Y           SQLSetStmtAttrW         Y           SQLSetStmtOption         Y		Y
SQLProceduresW         Y           SQLPutData         Y           SQLRowCount         Y           SQLSetConnectAttr         Y           SQLSetConnectAttrW         Y           SQLSetConnectOption         Y           SQLSetConnectOptionW         Y           SQLSetCursorName         Y           SQLSetCursorNameW         Y           SQLSetDescField         Y           SQLSetDescRec         Y           SQLSetEnvAttr         Y           SQLSetParam         Y           SQLSetPos         Y           SQLSetScrollOptions         N           SQLSetStmtAttr         Y           SQLSetStmtAttrW         Y           SQLSetStmtOption         Y		Y
SQLPutData         Y           SQLRowCount         Y           SQLSetConnectAttr         Y           SQLSetConnectOption         Y           SQLSetConnectOptionW         Y           SQLSetCursorName         Y           SQLSetCursorNameW         Y           SQLSetDescField         Y           SQLSetDescRec         Y           SQLSetEnvAttr         Y           SQLSetParam         Y           SQLSetPos         Y           SQLSetScrollOptions         N           SQLSetStmtAttr         Y           SQLSetStmtAttr         Y           SQLSetStmtAttrW         Y           SQLSetStmtOption         Y		Y
SQLRowCount         Y           SQLSetConnectAttr         Y           SQLSetConnectOption         Y           SQLSetConnectOptionW         Y           SQLSetCursorName         Y           SQLSetCursorNameW         Y           SQLSetDescField         Y           SQLSetDescRec         Y           SQLSetEnvAttr         Y           SQLSetParam         Y           SQLSetPos         Y           SQLSetScrollOptions         N           SQLSetStmtAttr         Y           SQLSetStmtAttrW         Y           SQLSetStmtOption         Y		Y
SQLSetConnectAttr Y   SQLSetConnectOption Y   SQLSetConnectOptionW Y   SQLSetCursorName Y   SQLSetCursorNameW Y   SQLSetDescField Y   SQLSetDescRec Y   SQLSetEnvAttr Y   SQLSetParam Y   SQLSetPos Y   SQLSetScrollOptions N   SQLSetStmtAttr Y   SQLSetStmtAttrW Y   SQLSetStmtOption Y		Y
SQLSetConnectOption Y   SQLSetConnectOptionW Y   SQLSetCursorName Y   SQLSetCursorNameW Y   SQLSetDescField Y   SQLSetDescRec Y   SQLSetEnvAttr Y   SQLSetParam Y   SQLSetPos Y   SQLSetScrollOptions N   SQLSetStmtAttr Y   SQLSetStmtAttr Y   SQLSetStmtAttr Y   SQLSetStmtAttr Y   SQLSetStmtAttr Y   SQLSetStmtAttr Y   SQLSetStmtOption Y		Y
SQLSetConnectOption         Y           SQLSetConnectOptionW         Y           SQLSetCursorName         Y           SQLSetCursorNameW         Y           SQLSetDescField         Y           SQLSetDescRec         Y           SQLSetEnvAttr         Y           SQLSetParam         Y           SQLSetPos         Y           SQLSetScrollOptions         N           SQLSetStmtAttr         Y           SQLSetStmtAttrW         Y           SQLSetStmtOption         Y		Y
SQLSetConnectOptionW         Y           SQLSetCursorName         Y           SQLSetCursorNameW         Y           SQLSetDescField         Y           SQLSetDescRec         Y           SQLSetEnvAttr         Y           SQLSetParam         Y           SQLSetPos         Y           SQLSetScrollOptions         N           SQLSetStmtAttr         Y           SQLSetStmtAttrW         Y           SQLSetStmtOption         Y		
SQLSetCursorName         Y           SQLSetCursorNameW         Y           SQLSetDescField         Y           SQLSetDescRec         Y           SQLSetEnvAttr         Y           SQLSetParam         Y           SQLSetPos         Y           SQLSetScrollOptions         N           SQLSetStmtAttr         Y           SQLSetStmtAttrW         Y           SQLSetStmtOption         Y		Y
SQLSetCursorNameW         Y           SQLSetDescField         Y           SQLSetDescRec         Y           SQLSetEnvAttr         Y           SQLSetParam         Y           SQLSetPos         Y           SQLSetScrollOptions         N           SQLSetStmtAttr         Y           SQLSetStmtAttrW         Y           SQLSetStmtOption         Y	<u> </u>	Y
SQLSetDescField         Y           SQLSetDescRec         Y           SQLSetEnvAttr         Y           SQLSetParam         Y           SQLSetPos         Y           SQLSetScrollOptions         N           SQLSetStmtAttr         Y           SQLSetStmtAttrW         Y           SQLSetStmtOption         Y		Y
SQLSetDescRec         Y           SQLSetEnvAttr         Y           SQLSetParam         Y           SQLSetPos         Y           SQLSetScrollOptions         N           SQLSetStmtAttr         Y           SQLSetStmtAttrW         Y           SQLSetStmtOption         Y		
SQLSetEnvAttr         Y           SQLSetParam         Y           SQLSetPos         Y           SQLSetScrollOptions         N           SQLSetStmtAttr         Y           SQLSetStmtAttrW         Y           SQLSetStmtOption         Y		
SQLSetParam         Y           SQLSetPos         Y           SQLSetScrollOptions         N           SQLSetStmtAttr         Y           SQLSetStmtAttrW         Y           SQLSetStmtOption         Y		
SQLSetPos       Y         SQLSetScrollOptions       N         SQLSetStmtAttr       Y         SQLSetStmtAttrW       Y         SQLSetStmtOption       Y		
SQLSetScrollOptions       N         SQLSetStmtAttr       Y         SQLSetStmtAttrW       Y         SQLSetStmtOption       Y		
SQLSetStmtAttr Y SQLSetStmtAttrW Y SQLSetStmtOption Y		
SQLSetStmtAttrW Y SQLSetStmtOption Y		
SQLSetStmtOption Y		Y
		Y

Function name	Support status
SQLSpecialColumnsW	Y
SQLStatistics	Y
SQLStatisticsW	Y
SQLTablePrivileges	Y
SQLTablePrivilegesW	Y
SQLTables	Y
SQLTablesW	Y
SQLTransact	Y

Y: Supported N: Not supported

# F.3 .NET Data Provider

There are the following ways to develop applications using Fujitsu Npgsql .NET Data Provider:

- Use the Fujitsu Npgsql.NET Data Provider API (classes and methods) directly.

Fujitsu Npgsql .NET Data Provider is created based on the open source software Npgsql. Refer to the "Npgsql - .Net Data Provider for Postgresql" for information on the APIs:



#### See

Refer to the Installation and Setup Guide for Client for the version of Npgsql that Fujitsu Npgsql .NET Data Provider is based on.

- Use the API of the .NET System.Data.Common namespace

It is possible to create applications that do not rely on a provider when you use the System.Data.Common namespace. Refer to the "Writing Provider-Independent Code in ADO.NET" in MSDN Library for information.



# See

Refer to the Npgsql API Reference for information on Npgsql API classes and methods.

# F.4 C Library (libpq)



#### See

Refer to "libpq - C Library" in "Client Interfaces" in the PostgreSQL Documentation.

# F.5 Embedded SQL in C



# See

 $Refer \ to \ "ECPG - Embedded \ SQL \ in \ C" \ in \ "Client \ Interfaces" \ in \ the \ PostgreSQL \ Documentation.$ 

# <u>Index</u>

[A]	Notes on the Query Builder44
Additional Notes on each Type Plugin42	NVL86
71	
[B]	[O]
BIND_VARIABLE99	OPEN_CURSOR102
	Outer Join Operator (+)80
[C]	
CLOSE_CURSOR100	[P]
Code examples for applications55,67	PARSE102
COLUMN_VALUE100	Path of library70
Comparison operator	Path of the library file70
	Pattern matching3
[D]	Precompiling example
DBMS_OUTPUT87	PUT
DBMS_SQL97	PUTF
DECODE83	
DEFINE_COLUMN101	PUT_LINE89,96
DISABLE88	[S]
DUAL Table	Scan Using a Vertical Clustered Index (VCI)
DUAL Table	
[E]	Settings for encrypting communication data for connection to the
	server
ENABLE 88	String functions and operators
Encoding System Settings	SUBSTR85
Entry information of subprogram70	
Example of specifying the hint clause56,69	[T]
EXECUTE101	Type Plugins41
r=1	n.n
[F]	[U]
FCLOSE92	Use the API of the .NET System.Data.Common namespace.213
FCLOSE_ALL93	Use the Fujitsu Npgsql.NET Data Provider API (classes and
FCOPY	methods) directly213
FETCH_ROWS101	UTL_FILE91
FFLUSH93	
FGETATTR93	[W]
FOPEN94	When setting from outside with environment variables49
FRENAME94	When specifying in the connection URI49
	When using Add-OdbcDsn22
[G]	When using ODBCConf.exe21
GET_LINE90	Ç
021_21.	
[1]	
IS_OPEN95	
_	
[L]	
Language settings	
Libraries to use	
[M]	
Management of database resources with Server Explorer6	
Manipulation of database resources with TableAdapter6	
•	
[N]	
NEW_LINE89,95	
Notes on automatically generating update-type SQL statements	
45	
Notes on metadata	
Notes on Server Explorer	
Notes on TableAdapter41	



# FUJITSU Enterprise Postgres 12



# Operation Guide

Linux > Windows >



# FUJITSU Enterprise Postgres 12



Operation Guide (Linux)

# **Preface**

# Purpose of this document

The FUJITSU Enterprise Postgres database system extends the PostgreSQL features and runs on the Linux platform. This document is the FUJITSU Enterprise Postgres Operation Guide.

#### Intended readers

This document is intended for those who install and operate FUJITSU Enterprise Postgres.

Readers of this document are assumed to have general knowledge of:

- PostgreSQL
- SQL
- Linux

## Structure of this document

This document is structured as follows:

# Chapter 1 Operating FUJITSU Enterprise Postgres

Describes how to operate FUJITSU Enterprise Postgres.

#### Chapter 2 Starting an Instance and Creating a Database

Describes how to start a FUJITSU Enterprise Postgres instance, and how to create a database.

# Chapter 3 Backing Up the Database

Describes how to back up the database.

# Chapter 4 Configuring Secure Communication Using Secure Sockets Layer

Describes communication data encryption between the client and the server.

# Chapter 5 Protecting Storage Data Using Transparent Data Encryption

Describes how to encrypt the data to be stored in the database.

### Chapter 6 Data Masking

Describes the data masking feature.

# Chapter 7 Periodic Operations

Describes the periodic database operations that must be performed on FUJITSU Enterprise Postgres.

# Chapter 8 Streaming Replication Using WebAdmin

Describes how to create a streaming replication cluster using WebAdmin.

# Chapter 9 Installing and Operating the In-memory Feature

Describes how to install and operate the in-memory feature.

# Chapter 10 Parallel Query

Describes the factors taken into consideration by FUJITSU Enterprise Postgres when performing parallel queries.

# Chapter 11 High-Speed Data Load

Describes how to install and operate high-speed data load.

# Chapter 12 Global Meta Cache

Describes how to use Grobal Meta Cache feature.

# Chapter 13 Backup/Recovery Using the Copy Command

Describes backup and recovery using the copy command created by the user.

# Chapter 14 Actions when an Error Occurs

Describes how to perform recovery when disk failure or data corruption occurs.

#### Appendix A Parameters

Describes the FUJITSU Enterprise Postgres parameters.

#### Appendix B System Administration Functions

Describes the system administration functions of FUJITSU Enterprise Postgres.

#### Appendix C System Views

Describes how to use the system view in FUJITSU Enterprise Postgres.

#### Appendix D Tables Used by Data Masking

Describes the tables used by the data masking feature.

# Appendix E Tables Used by High-Speed Data Load

Describes the tables used by high-speed data load.

## Appendix F Starting and Stopping the Web Server Feature of WebAdmin

Describes how to start and stop WebAdmin (Web server feature).

#### Appendix G WebAdmin Wallet

Describes how to use the Wallet feature of WebAdmin.

# Appendix H WebAdmin Disallow User Inputs Containing Hazardous Characters

Describes characters not allowed in WebAdmin.

#### Appendix I Copy Command Samples that Use the Advanced Copy Feature of the ETERNUS Disk Array

Describes copy command samples that use the advanced copy feature of the ETERNUS disk array.

# Appendix J Collecting Failure Investigation Data

Describes how to collect information for initial investigation.

## **Export restrictions**

Exportation/release of this document may require necessary procedures in accordance with the regulations of your resident country and/or US export control laws.

# Issue date and version

Edition 1.0: July 2020

# Copyright

Copyright 2015-2020 FUJITSU LIMITED

# Contents

Chapter 1 Operating FUJITSU Enterprise Postgres	
1.1 Operating Methods	
1.2 Starting WebAdmin	2
1.2.1 Logging in to WebAdmin	2
1.3 Starting pgAdmin	
1.3.1 Starting pgAdmin	
1.3.2 Adding an Instance	4
1.3.3 Connecting/Disconnecting an Instance	<i>e</i>
1.4 Operations Using Commands	e
1.5 Operating Environment of FUJITSU Enterprise Postgres	
1.5.1 Operating Environment	
1.5.2 File Composition	
1.6 Notes on Compatibility of Applications Used for Operations	9
1.7 Notes on Upgrading Database Instances	10
1.7.1 Additional Steps for upgrading to FUJITSU Enterprise Postgres with Vertical Clustered Index (VCI) Enabled	11
Chapter 2 Starting an Instance and Creating a Database	12
2.1 Starting and Stopping an Instance	
2.1.1 Using WebAdmin.	
2.1.2 Using Server Commands.	
2.2 Creating a Database	
2.2.1 Using pgAdmin.	
2.2.2 Using Client Commands.	
2.2.2 Using Choir Communication	1
Chapter 3 Backing Up the Database	18
3.1 Periodic Backup	19
3.2 Backup Methods	19
3.2.1 Using WebAdmin	19
3.2.2 Using Server Commands	20
	0.0
Chapter 4 Configuring Secure Communication Using Secure Sockets Layer	
4.1 Configuring Communication Data Encryption.	
4.1.1 Issuing a Certificate.	
4.1.2 Deploying a Server Certificate File and a Server Private Key File	
4.1.3 Distributing a CA Certificate File to the Client.	
4.1.4 Configuring the Operating Environment for the Database Server.	
4.1.5 Configuring the Operating Environment for the Client	
4.1.6 Performing Database Multiplexing	25
Chapter 5 Protecting Storage Data Using Transparent Data Encryption	26
5.1 Protecting Data Using Encryption.	
5.2 Setting the Master Encryption Key	
5.3 Opening the Keystore	
5.4 Encrypting a Tablespace	
5.5 Checking an Encrypted Tablespace	
5.6 Managing the Keystore	
5.6.1 Changing the Master Encryption Key	
5.6.2 Changing the Keystore Passphrase	
5.6.3 Enabling Automatic Opening of the Keystore	
5.6.4 Backing Up and Recovering the Keystore	
5.7 Backing Up and Restoring/Recovering the Database.	
5.8 Importing and Exporting the Database	
5.10 Operations in Cluster Systems	
5.10 Operations in Cluster Systems.	7
5.10.1 HA Clusters that do not Use Database Multiplexing	36

5.11 Security-Related Notes	37
5.12 Tips for Installing Built Applications	38
Chapter 6 Data Masking	30
6.1 Masking Policy	
6.1.1 Masking Target	
6.1.2 Masking Type	
6.1.3 Masking Condition.	
6.1.4 Masking Format	
6.2 Usage Method.	
6.2.1 Creating a Masking Policy.	
6.2.2 Changing a Masking Policy.	
6.2.3 Confirming a Masking Policy	
6.2.4 Enabling and Disabling a Masking Policy	
6.2.5 Deleting a Masking Policy	
6.3 Data Types for Masking	
6.4 Security Notes.	
Chapter 7 Periodic Operations	49
7.1 Configuring and Monitoring the Log	
7.2 Monitoring Disk Usage and Securing Free Space	
7.2.1 Monitoring Disk Usage	
7.2.2 Securing Free Disk Space	
7.3 Automatically Closing Connections	50
7.4 Monitoring the Connection State of an Application	
7.4.1 Using the View (pg_stat_activity)	51
7.4.2 Using pgAdmin	51
7.5 Reorganizing Indexes	52
7.6 Monitoring Database Activity	53
7.6.1 Information that can be Collected	54
7.6.2 Collection Configuration	55
7.6.3 Information Reset	55
Chapter 8 Streaming Replication Using WebAdmin	57
8.1 Creating a Standby Instance	57
8.2 Promoting a Standby Instance	
8.3 Converting an Asynchronous Replication to Synchronous	
8.4 Converting a Synchronous Replication to Asynchronous	
8.5 Joining a Replication Cluster	
Chapter 9 Installing and Operating the In-memory Feature	62
9.1 Installing Vertical Clustered Index (VCI)	
9.1.1 Evaluating whether to Install VCI	
9.1.2 Estimating Resources.	63
9.1.3 Setting up	63
9.1.3.1 Setting Parameters	63
9.1.3.2 Installing the Extensions.	64
9.1.3.3 Creating a VCI	64
9.1.3.4 Confirming that the VCI has been Created	65
9.1.4 Data that can Use VCI	
9.1.4.1 Relation Types	65
9.1.4.2 Data Types	66
9.2 Operating VCI	67
9.2.1 Commands that cannot be Used for VCI	68
9.2.2 Data Preload Feature	69
Chapter 10 Parallel Query	70
10.1 CDLU and Calculation	70

10.2 Increase of Workers during Runtime.	70
Chapter 11 High-Speed Data Load	71
11.1 Installing High-Speed Data Load	
11.1.1 Deciding whether to Install	71
11.1.2 Estimating Resources	71
11.1.3 Setup	72
11.1.3.1 Setting Parameters.	72
11.1.3.2 Installing the Extension.	73
11.2 Using High-Speed Data Load	73
11.2.1 Loading Data	73
11.2.2 Recovering from a Data Load that Ended Abnormally	74
11.3 Removing High-Speed Data Load	75
11.3.1 Removing the Extension.	75
Chapter 12 Global Meta Cache	77
12.1 Usage	77
12.1.1 Deciding Whether to Enable the Global Meta Cache Feature	77
12.1.2 Estimating Memory for Global Meta Cache	77
12.1.3 How the GMC Memory Area Is Used	77
12.1.4 Enabling the Global Meta Cache Feature	77
12.1.5 Estimating Resources	78
12.2 Statistics.	78
12.2.1 System View	78
Chapter 13 Backup/Recovery Using the Copy Command	79
13.1 Configuration of the Copy Command	79
13.2 Backup Using the Copy Command	83
13.3 Recovery Using the Copy Command	84
13.4 Copy Command Interface	85
13.4.1 Copy Command for Backup	
13.4.2 Copy Command for Recovery	87
Chapter 14 Actions when an Error Occurs	89
14.1 Recovering from Disk Failure (Hardware)	90
14.1.1 Using WebAdmin	90
14.1.2 Using Server Command	91
14.2 Recovering from Data Corruption	95
14.2.1 Using WebAdmin	
14.2.2 Using the pgx_rcvall Command	
14.3 Recovering from an Incorrect User Operation	97
14.3.1 Using WebAdmin	98
14.3.2 Using the pgx_rcvall Command	
14.4 Actions in Response to an Application Error	
14.4.1 When using the view (pg_stat_activity)	
14.4.2 Using the ps Command	
14.4.3 Using pgAdmin	
14.5 Actions in Response to an Access Error.	
14.6 Actions in Response to Insufficient Space on the Data Storage Destination	
14.6.1 Using a Tablespace	
14.6.2 Replacing the Disk with a Larger Capacity Disk	
14.6.2.1 Using WebAdmin.	
14.6.2.2 Using Server Commands.	
14.7 Actions in Response to Insufficient Space on the Backup Data Storage Destination	
14.7.1.1 Using Wah Admin	
14.7.1.1 Using WebAdmin	
14.7.2 Replacing the Disk with a Larger Capacity Disk.	
THE REPRESENTATION OF THE PRICE WITH A DRIEGH CAUGUITY DISK	

14.7.2.1 Using WebAdmin	109
14.7.2.2 Using Server Commands	110
14.8 Actions in Response to Insufficient Space on the Transaction Log Storage Destination	114
14.8.1 Replacing the Disk with a Larger Capacity Disk	
14.8.1.1 Using WebAdmin	
14.8.1.2 Using Server Commands	
14.9 Errors in More Than One Storage Disk	
14.10 Actions in Response to Instance Startup Failure	
14.10.1 Errors in the Configuration File	
14.10.2 Errors Caused by Power Failure or Mounting Issues	
14.10.3 Other Errors	
14.10.3.1 Using WebAdmin	
14.10.3.2 Using Server Commands	
14.11 Actions in Response to Failure to Stop an Instance	
14.11.1 Using WebAdmin	
14.11.2 Using Server Commands	
14.11.2.1 Stopping the Instance Using the Fast Mode	
14.11.2.2 Stopping the Instance Using the Immediate Mode	
14.11.2.3 Forcibly Stopping the Server Process	
14.12 Actions in Response to Failure to Create a Streaming Replication Standby Instance	
14.13 Actions in Response to Error in a Distributed Transaction	
14.14 I/O Errors Other than Disk Failure	
14.14.1 Network Error with an External Disk	
14.14.2 Errors Caused by Power Failure or Mounting Issues	
14.15 Anomaly Detection and Resolution	
14.15.1 Port Number and Backup Storage Path Anomalies	
14.15.2 Mirroring Controller Anomalies	122
Appendix A Parameters	124
Appoint A A T didinoto in the same and a sam	
Appendix B System Administration Functions	128
B.1 WAL Mirroring Control Functions.	128
B.2 Transparent Data Encryption Control Functions.	
B.3 Data Masking Control Functions	129
B.3.1 pgx_alter_confidential_policy	129
B.3.2 pgx_create_confidential_policy	
B.3.3 pgx_drop_confidential_policy	
B.3.4 pgx_enable_confidential_policy	139
B.3.5 pgx_update_confidential_values	140
B.4 VCI Data Load Control Function.	141
B.5 High-Speed Data Load Control Functions	142
A	4.46
Appendix C System Views	
C.1 pgx_tablespaces	
C.2 pgx_stat_lwlock	
C.3 pgx_stat_latch	
C.4 pgx_stat_walwriter	
C.5 pgx_stat_sql	
C.6 pgx_stat_gmc	144
Appendix D Tables Used by Data Masking	1Δ6
D.1 pgx_confidential_columns	
D.2 pgx_confidential_policies	
D.3 pgx_confidential_values	
Appendix E Tables Used by High-Speed Data Load	148
E.1 pgx_loader_state	148
Appendix F Starting and Stopping the Web Server Feature of WebAdmin	149

F.1 Starting the Web Server Feature of WebAdmin	149
F.2 Stopping the Web Server Feature of WebAdmin	149
Appendix G WebAdmin Wallet	151
G.1 Creating a Credential.	151
G.2 Using a Credential.	152
Appendix H WebAdmin Disallow User Inputs Containing Hazardous Characters	153
Appendix I Copy Command Samples that Use the Advanced Copy Feature of the ETERNUS Disk Array	154
Appendix J Collecting Failure Investigation Data	156
Index	157

# Chapter 1 Operating FUJITSU Enterprise Postgres

This chapter describes how to operate FUJITSU Enterprise Postgres.

# 1.1 Operating Methods

There are two methods of managing FUJITSU Enterprise Postgres operations:

- Operation management using GUI tools
- Operation management using commands



See

Before performing database multiplexing using database multiplexing, refer to "Database Multiplexing Mode" in the Cluster Operation Guide (Database Multiplexing).

### Operation management using GUI tools

This involves managing operations using the WebAdmin and pgAdmin GUI tools.

- Management using WebAdmin

This removes the requirement for complex environment settings and operational design for backup and recovery that is usually required for running a database. It enables you to easily and reliably monitor the state of the database, create a streaming replication cluster, back up the database, and restore it even if you do not have expert knowledge of databases.

- Management using pgAdmin

When developing applications and maintaining the database, you can use pgAdmin to perform simple operations on database objects, such as:

- Rebuild indexes and update statistics
- Create, delete, and update database objects

In addition, from pgAdmin of FUJITSU Enterprise Postgres, you can use the expanded features provided by FUJITSU Enterprise Postgres on the PostgreSQL SQL commands.



See

Refer to pgAdmin Help for information on the expanded features of pgAdmin provided by FUJITSU Enterprise Postgres.

### Operation management using commands

You can use commands for configuring and operating the database and managing operations.



You cannot combine WebAdmin and server commands to perform the following operations:

- Use commands to operate an instance created using WebAdmin.
- Use WebAdmin to recover a database backed up using commands.

For instances created with WebAdmin, however, backup can be obtained with the pgx\_dmpall command. Also, WebAdmin can perform recovery by using the backup obtained with the pgx\_dmpall command.

- To operate an instance created using the initdb command in WebAdmin, the instance needs to be imported using WebAdmin.

- You can perform backup and restoration in pgAdmin, but the backup data obtained with WebAdmin and pgx\_dmpall is not compatible with the backup data obtained with pgAdmin.
- Refer to pgAdmin Help for other notes on pgAdmin.

### Features used in each phase

The following table lists the features used in each phase for GUI-based operations and command-based operations.

Operation		Operation with the GUI	Operation with commands
Setup	Creating an instance	WebAdmin is used. The server machine capacity, and the optimum parameter for operations using WebAdmin, are set automatically.	The configuration file is edited directly using the initdb command.
	Creating a standby instance	WebAdmin is used.  WebAdmin performs a base backup of the source instance and creates a standby instance.	A standby instance is created using the pg_basebackup command.
	Changing the configuration files	WebAdmin is used.	The configuration file is edited directly.
Starting and stopping	an instance	WebAdmin is used.	The pg_ctl command is used.
Creating a database		This is defined using pgAdmin of the GUI tool, or using the psql command or the application after specifying the DDL statement.	
Backing up the database		WebAdmin, or the pgx_dmpall command, is used.	It is recommended that the pgx_dmpall command be used. Recovery to the latest database can be performed.
Database recovery		WebAdmin is used.	To use the backup that was performed using the pgx_dmpall command, the pgx_rcvall command is used.
Monitoring	Database errors	The status in the WebAdmin window can be checked. (*1)	The messages that are output to the database server log are monitored (*1)
	Disk space	The status in the WebAdmin window can be checked. A warning will be displayed if the free space falls below 20%. (*1)	This is monitored using the df command of the operating system, for example. (*1)
	Connection status	This can be checked using pgAdmin of the GUI tool, or referencing pg_stat_activity of the standard statistics view from psql or the application.	

<sup>\*1:</sup> This can be used together with system log monitoring using operations management middleware (Systemwalker Centric Manager, for example).

# 1.2 Starting WebAdmin

This section describes how to start and log in to WebAdmin.

# 1.2.1 Logging in to WebAdmin

This section describes how to log in to WebAdmin.

### **User environment**

It is recommended to use the following browsers with WebAdmin:

- Internet Explorer 11
- Microsoft Edge (Build41 or later)

WebAdmin will work with other browsers, such as Firefox and Chrome, however, the look and feel may be slightly different.

### Startup URL for WebAdmin

In the browser address bar, type the startup URL of the WebAdmin window in the following format:

http://hostNameOrIpAddress:portNumber/

- hostNameOrIpAddress: The host name or IP address of the server where WebAdmin is installed.
- portNumber. The port number of WebAdmin. The default port number is 27515.



For a server with IP address "192.0.2.0" and port number "27515"

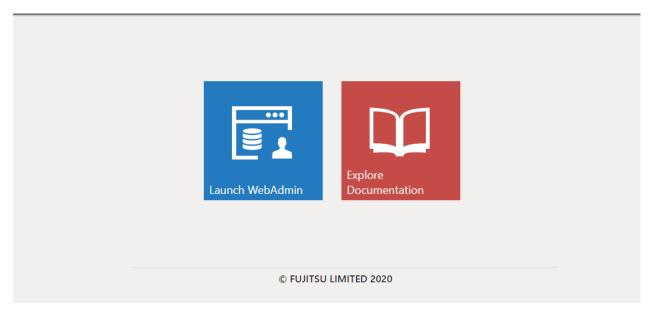
http://192.0.2.0:27515/

The startup URL window shown below is displayed. From this window you can log in to WebAdmin or access the product documentation.

......

# **FUJITSU Enterprise Postgres**







- You must start the Web server feature of WebAdmin before using WebAdmin.
- Refer to "Appendix F Starting and Stopping the Web Server Feature of WebAdmin" for information on how to start the Web server feature of WebAdmin.

### Log in to WebAdmin

Click [Launch WebAdmin] in the startup URL window to start WebAdmin and display the login window.

To log in, specify the following values:

- [User name]: User name (OS user account) of the instance administrator
- [Password]: Password corresponding to the user name



Use the OS user account as the user name of the instance administrator. Refer to "Creating an Instance Administrator" in the Installation and Setup Guide for Server for details.

# 1.3 Starting pgAdmin

This section describes how to start pgAdmin, how to add an instance required for managing a database, and how to connect to and disconnect from the instance.

You can use pgAdmin on the Windows client.

### 1.3.1 Starting pgAdmin

This section explains how to start pgAdmin if you are using it from the product "FUJITSU Enterprise Postgres Client (AAbit) x SPz" (where AA is "32" or "64", "x" is the product version, and "z" is the product level (x SPz)).

Windows Server(R) 2012

From the [Start] screen, start [pgAdmin 4 (AAbit) (x SPz)].

Windows(R) 8.1 or Windows Server(R) 2012 R2

From the [Apps] view, start [pgAdmin 4 (AAbit) (x SPz)].

Windows(R) 10, Windows Server(R) 2016, or Windows Server(R) 2019

Click [Start] >> [All apps] >> [FUJITSU Enterprise Postgres Client(AAbit) x SPz] and start [pgAdmin 4 (AAbit) (x SPz)].



- You must start the instance to be connected to before using pgAdmin.
- Refer to "2.1 Starting and Stopping an Instance" for information on how to start an instance.
- When using pgAdmin4 with Microsoft Edge, enable network access by loopback in Microsoft Edge. Also, add Microsoft Edge to the loopback exclusion list.

# 1.3.2 Adding an Instance

This section describes how to add an instance to be connected to.

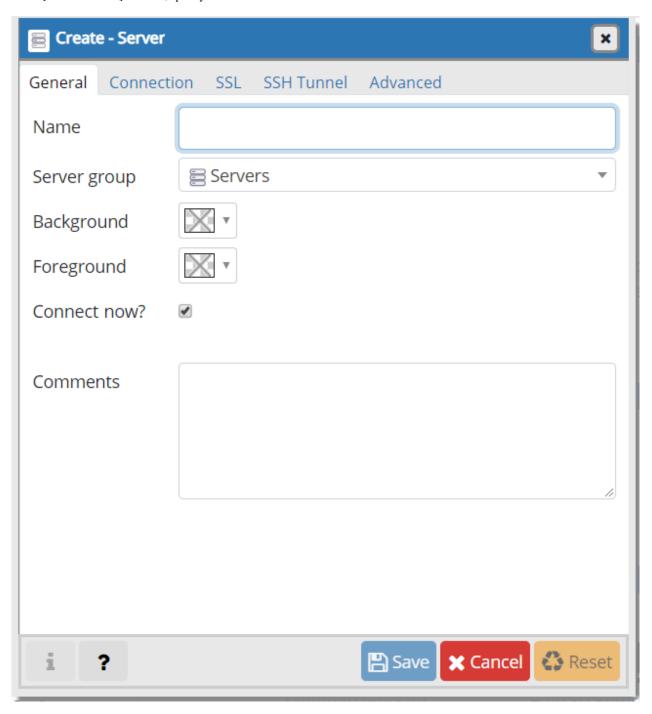


If you use a link-local address with version 6 of the TCP/IP protocol, you may encounter the following error when registering the server with pgAdmin. Therefore, do not use link-local addresses.

unsupported format character '"' (0x22) at index 96

 $1. \ \ In the \ [Browser] \ pane, \ right-click \ [Servers], \ and \ then \ click \ [Create] >> [Server].$ 

2. In the [Create - Server] window, specify a value for each item.



### [General] tab

- [Name]: Name of the instance to be managed

### [Connection] tab

- [Host name/address]: Host name or IP address of the server where FUJITSU Enterprise Postgres is installed
- [Port]: Port number of the instance
- [Username]: User name of the instance administrator
- [Password]: Password for the user name specified in [Username]

When you add an instance using pgAdmin, the instance is automatically connected to immediately after the addition is completed.



If you select [Save password], the FUJITSU Enterprise Postgres connection password is stored in the following location. Set the appropriate access permissions for the password file to protect it from unauthorized access.

- %APPDATA%\Roaming\pgAdmin\pgadmin4.db

### 1.3.3 Connecting/Disconnecting an Instance

This section describes how to connect pgAdmin to an instance, and how to disconnect it.



To connect to an instance created using WebAdmin, you must first configure the settings in the [Client authentication] window of WebAdmin to permit connection from pgAdmin.

......



Refer to "Changing the settings" in the Installation and Setup Guide for Server for information on the [Client authentication] window of WebAdmin.

### Connecting to an instance

Starting pgAdmin does not connect it to any instance.

To connect to an instance, in the [Browser] pane, right-click the instance, and then click [Connect Server].

If a password was not saved when the instance was added, enter a password in the password entry window that is displayed.

#### Disconnecting from an instance

To disconnect from an instance, in the [Browser] pane, right-click the server, and then click [Disconnect Server].

# 1.4 Operations Using Commands

You can operate and manage the database using the following commands:

- Server commands

This group of commands includes commands for creating a database cluster and controlling the database. You can run these commands on the server where the database is operating.

To use these commands, you must configure the environment variables.



### See

- Refer to "PostgreSQL Server Applications" under "Reference" in the PostgreSQL Documentation, or "Reference" for information on server commands.

.....

- Refer to "Configure the environment variables" in the procedure to create instances in "Using the initdb Command" in the Installation and Setup Guide for Server for information on configuring the environment variables.
- Client commands

This group of commands includes the psql command and commands for extracting the database cluster to a script file. These commands can be executed on the client that can connect to the database, or on the server on which the database is running.

To use these commands, you must configure the environment variables.



- Refer to "PostgreSQL Client Applications" under "Reference" in the PostgreSQL Documentation, or "Reference" for information on client commands.

......

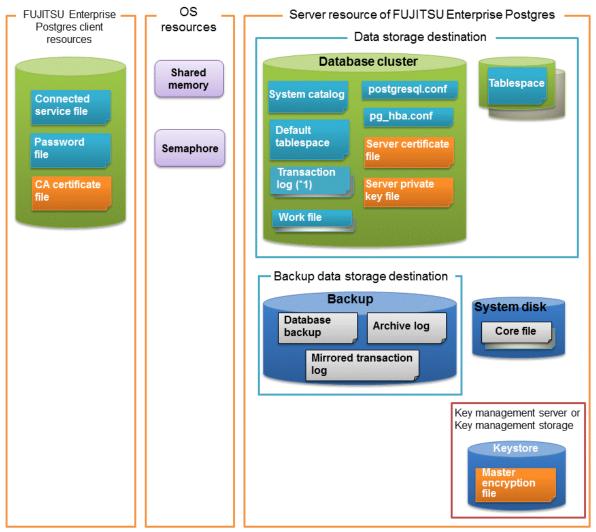
- Refer to "Configuring Environment Variables" in the Installation and Setup Guide for Client for information on the values to be set in the environment variables.

# 1.5 Operating Environment of FUJITSU Enterprise Postgres

This section describes the operating environment and the file composition of FUJITSU Enterprise Postgres.

## 1.5.1 Operating Environment

The following figure shows the configuration of the FUJITSU Enterprise Postgres operating environment. The tables given below list the roles of the OS resources and FUJITSU Enterprise Postgres resources.



<sup>\*1:</sup> To distribute the I/O load, place the transaction log on a different disk from the data storage destination.

Table 1.1 OS resources

Туре	Role
Shared memory	Used when a database process exchanges information with an external
Semaphore	process.

Table 1.2 FUJITSU Enterprise Postgres client resources

Туре	Role	
Connection service file	Specifies information, such as the host name, user name, and password, for connecting to FUJITSU Enterprise Postgres.	
Password file	Securely manages the password for connecting to FUJITSU Enterprise Postgres.	
CA certificate file	CA (certificate authority) certificate used for server authentication when encrypting communication data.	

Table 1.3 Server resources of FUJITSU Enterprise Postgres

Туре	Role	
Database cluster	Database storage area on the database storage disk. It is a collection of databases managed by an instance.	
System catalog	Contains information required for the system to run, including the database definition information and the operation information created by the user.	
Default tablespace	Contains table files and index files stored by default.	
Transaction log	Contains log information in case of a crash recovery or rollback. This is the same as the WAL (Write Ahead Log).	
Work file	Work file used when executing applications or commands.	
postgresql.conf	Contains information that defines the operating environment of FUJITSU Enterprise Postgres.	
pg_hba.conf	FUJITSU Enterprise Postgres uses this file to authenticate individual clien hosts.	
Server certificate file	Contains information about the server certificate to be used when encrypting communication data and authenticating a server.	
Server private key file	Contains information about the server private key to be used when encrypting communication data and authenticating a server	
Tablespace	Stores table files and index files in a separate area from the database cluster.  Specify a space other than that under the database cluster.	
Backup	Stores the data required for recovering the database when an error, such as disk failure, occurs.	
Database backup	Contains the backup data for the database.	
Archive log	Contains the log information for recovery.	
Mirrored transaction log (mirrored WAL)	Enables a database cluster to be restored to the state immediately before an error even if both the database cluster and transaction log fail when performing backup/recovery operations using the pgx_dmpall command or WebAdmin.	
Core file	FUJITSU Enterprise Postgres process core file that is output when an error occurs during a FUJITSU Enterprise Postgres process.	
Key management server or key management storage	Server or storage where the master encryption key file is located.	

Туре	Role
Master encryption key file	Contains the master encryption key to be used when encrypting storage data. The master encryption key file is managed on the key management server or key management storage.

# 1.5.2 File Composition

FUJITSU Enterprise Postgres consists of the following files for controlling and storing the database. The table below shows the relationship between the number of such files and their location within a single instance.

Table 1.4 Number of files within a single instance and how to specify their location

File type	Required	Quantity	How to specify the location
Duo autom filos	Y	Multiple	Note that "< <i>x</i> >" indicates the product version.
Program files	I	Multiple	/opt/fsepv< <i>x</i> >server64
Database cluster	Y	1	Specify using WebAdmin or server commands.
Tablespace	Y	Multiple	Specify a space other than that under the database cluster, using pgAdmin or the DDL statement.
Backup	Y	Multiple	Specify using WebAdmin or server commands.
Core file	Y	Multiple	Specify using WebAdmin, server commands, or postgresql.conf.
Server certificate file (*1)	N	1	Specify using postgresql.conf.
Server private key file (*1)	N	1	Specify using postgresql.conf.
Master encryption key file (*1)	N	1	Specify the directory created as the key store using postgresql.conf.
Connection service file (*1)	N	1	Specify using environment variables.
Password file (*1)	N	1	Specify using environment variables.
CA certificate file (*1)	N	1	Specify using environment variables.

Y: Mandatory N: Optional

<sup>\*1:</sup> Set manually when using the applicable feature.



- Do not use an NFS for files used in FUJITSU Enterprise Postgres except when creating a database space in a storage device on a network.
- If anti-virus software is used, set scan exception settings for directories so that none of the files that comprise FUJITSU Enterprise Postgres are scanned for viruses. Alternatively, if the files that comprise FUJITSU Enterprise Postgres are to be scanned for viruses, stop FUJITSU Enterprise Postgres and perform the scan when tasks that use FUJITSU Enterprise Postgres are not operating.

# 1.6 Notes on Compatibility of Applications Used for Operations

When you upgrade FUJITSU Enterprise Postgres to a newer version, there may be some effect on applications due to improvements or enhancements in functionality.

Take this into account when creating applications so that you can maintain compatibility after upgrading to a newer version of FUJITSU Enterprise Postgres.



Refer to "Notes on Application Compatibility" in the Application Development Guide for details.

# 1.7 Notes on Upgrading Database Instances

When upgrading FUJITSU Enterprise Postgres 9.4 or newer database instances to FUJITSU Enterprise Postgres 10 or later using pg\_upgrade, there are certain steps you need to follow.

Before using pg\_upgrade, remove the following extensions from all databases in the instance, except "template0":

- pg\_stat\_statements
- pgx\_io
- pgx\_paging
- pgx\_network
- pgx\_network\_err
- pgx\_cpu
- pgx\_memory
- pgx\_swap
- pgx\_disk
- pgx\_process
- pgx\_log
- oracle\_compatible
- pg\_dbms\_stats
- pg\_hint\_plan

For all databases except "template0", execute the following command to remove these extensions:

DROP EXTENSION extensionName;

Once the pg\_upgrade operation is complete, for all databases except "template0", execute the following command to re-create these extensions as required:

CREATE EXTENSION extensionName;



- It is strongly recommended to back up the database using pg\_dump before performing pg\_upgrade or using DROP EXTENSION.
- If there are any columns created in the user tables using a data type from these extensions, then DROP EXTENSION will also drop these columns. Therefore, it is essential that alternate upgrade mechanisms are considered instead of pg\_upgrade, in such scenarios. These may include pg\_dump/pg\_restore.

# 1.7.1 Additional Steps for upgrading to FUJITSU Enterprise Postgres with Vertical Clustered Index (VCI) Enabled

When upgrading FUJITSU Enterprise Postgres 11.0 or earlier instances that are using the VCI extension to FUJITSU Enterprise Postgres 12 using pg\_upgrade, additional steps must be performed because of the incompatibility of the VCI extension between FUJITSU Enterprise Postgres 12 and FUJITSU Enterprise Postgres 11 or earlier.

Follow the procedure below in all databases in the FUJITSU Enterprise Postgres 11 or earlier instance, except "template0".

### Before upgrading

1. Obtain the CREATE INDEX Definitions

Run the query below to list all the VCI indexes created in the database. Ensure that these indexes are re-created in the FUJITSU Enterprise Postgres 12 or later instance after pg\_upgrade has finished.

```
SELECT nspname || '.' || relname AS index_relname,* FROM pg_class, pg_namespace
WHERE relnamespace = pg_namespace.oid AND relam IN (SELECT oid FROM pg_am WHERE amname='vci');
```

For each index\_relname listed above, execute the commands below to obtain the CREATE INDEX definition (to use the same SQL syntax while re-creating the indexes on the FUJITSU Enterprise Postgres 12 or later instance).

```
SELECT pg_get_indexdef('indexName'::regclass);
```

2. Drop the VCI indexes and VCI extension along with all its dependencies.

To remove all the VCI indexes and VCI internal objects that are created in FUJITSU Enterprise Postgres, execute the commands below. VCI internal objects will be created in FUJITSU Enterprise Postgres 12 or later automatically when CREATE EXTENSION for VCI is executed.

DROP EXTENSION VCI CASCADE;



To restore the VCI extension in the FUJITSU Enterprise Postgres 11 or earlier instance, execute CREATE EXTENSION.

### After upgrading

Once the pg\_upgrade operation is complete, for all databases except "template0", execute CREATE EXTENSION to create the VCI extension, and then execute CREATE INDEX for all the VCI indexes as required.

# Chapter 2 Starting an Instance and Creating a Database

This chapter describes basic operations, from starting an instance to creating a database.

# 2.1 Starting and Stopping an Instance

This section describes how to start and stop an instance.

- 2.1.1 Using WebAdmin
- 2.1.2 Using Server Commands



To automatically start or stop an instance when the operating system on the database server is started or stopped, refer to "Configuring Automatic Start and Stop of an Instance" in the Installation and Setup Guide for Server and configure the settings.



The collected statistics are initialized if an instance is stopped in the "Immediate" mode or if it is abnormally terminated. To prepare for such initialization of statistics, consider regular collection of the statistics by using the SELECT statement. Refer to "The Statistics Collector" in "Server Administration" in the PostgreSQL Documentation for information on the statistics.

### 2.1.1 Using WebAdmin

WebAdmin enables you to start or stop an instance and check its operating status.

### Starting an instance

Start an instance by using the [Instances] tab in WebAdmin.

is displayed when an instance is stopped.

To start a stopped instance, click .

### Stopping an instance

Stop an instance by using the [Instances] tab in WebAdmin.

is displayed when an instance is active.

To stop an active instance, click .

### Stop mode

Select the mode in which to stop the instance. The following describes the operations of the modes:

Stop mode	Connected clients	Backup being executed using the command	
Smart mode (*1)	Waits for all connected clients to be disconnected.	Waits for backups being executed using the command to finish.	
Fast mode	Rolls back all transactions being executed and forcibly disconnects clients.	Terminates backups being executed using the command.	
Immediate mode	All server processes are terminated immediately. Crash recovery is executed the next time the instance is started.		

Stop mode	Connected clients	Backup being executed using the command
Kill process mode	Send SIGKILL to the process and abort all active transactions. This will lead to a crash-recovery run at the next restart.	

<sup>\*1:</sup> When the processing to stop the instance in the Smart mode has started and you want to stop immediately, use the following procedure:

- 1. Restart the Web server feature of WebAdmin.
- 2. In the [Instances] tab, click 2.
- 3. In the [Instances] tab, click , and select the Immediate mode to stop the instance.

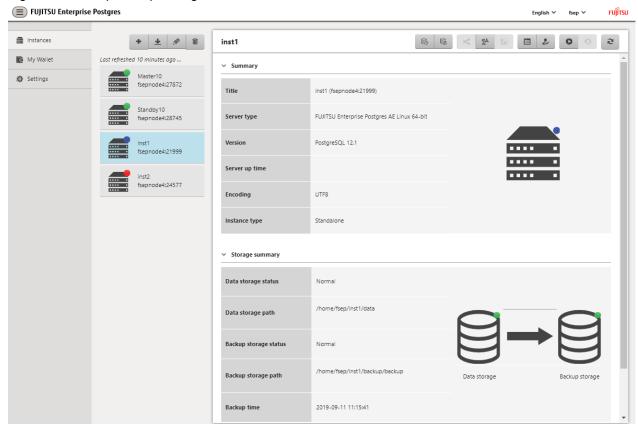
### Checking the operating status of an instance

You can check the operating status of an instance by using the [Instances] tab. The following indicators are used to show the status of a resource.

Status indicator	Explanation	
•	The resource is operating normally.	
•	The resource is stopped.	
•	There is an error in the resource.	
•	An operation is in progress on this resource or the status is being checked.	
<u> </u>	The resource is not operating optimally and needs intervention.	

If an instance stops abnormally, remove the cause of the stoppage and start the instance by using WebAdmin.

Figure 2.1 Example of operating status indicators





- When operating WebAdmin, click to update the status. WebAdmin will reflect the latest status of the operation or the instance resources from the server.
- If an error occurs while communicating with the server, there may be no response from WebAdmin. When this happens, close the browser and then log in again. If this does not resolve the issue, check the system log of the server and confirm whether a communication error has occurred.
- The following message is output during startup of an instance when the startup process is operating normally, therefore, the user does not need to be aware of this message:

FATAL: the database system is starting up

# 2.1.2 Using Server Commands

Server commands enable you to start or stop an instance and check its operating status.

To use sever commands, configure the environment variables.



Refer to "Configure the environment variables" in the procedure to create instances in "Using the initdb Command" in the Installation and Setup Guide for Server for information on configuring the environment variables.

### Starting an instance

Use the pg\_ctl command to start an instance.

Specify the following values in the pg\_ctl command:

- Specify "start" as the mode.
- Specify the data storage destination in the -D option. If the -D option is omitted, the value of the PGDATA environment variable is used by default.

If an application, command, or process tries to connect to the database while the instance is starting up, the message "FATAL:the database system is starting up(11189)" is output. However, this message may also be output if the instance is started without the -W option specified. This message is output by the pg\_ctl command to check if the instance has started successfully. Therefore, ignore this message if there are no other applications, commands, or processes that connect to the database.



> pg\_ctl start -D /database/inst1



If the -W option is specified, the command will return without waiting for the instance to start. Therefore, it may be unclear as to whether the instance startup was successful or failed.

### Stopping an instance

Use the pg\_ctl command to stop an instance.

Specify the following values in the pg\_ctl command:

- Specify "stop" as the mode.
- Specify the data storage destination in the -D option. If the -D option is omitted, the value of the PGDATA environment variable is used by default.



> pg\_ctl stop -D /database/inst1

### Checking the operating status of an instance

Use the pg\_ctl command to check the operating status of an instance.

Specify the following values in the pg\_ctl command:

- Specify "status" as the mode.
- Specify the data storage destination in the -D option. If the -D option is omitted, the value of the PGDATA environment variable is used by default.



When the instance is active:

> pg\_ctl status -D /database/inst1
pg\_ctl: server is running (PID: 1234)

When the instance is inactive:

> pg\_ctl status -D /database/instl
pg\_ctl: no server running.



Refer to "pg\_ctl" under "Reference" in the PostgreSQL Documentation for information on pg\_ctl command.

# 2.2 Creating a Database

This section explains how to create a database.

- 2.2.1 Using pgAdmin
- 2.2.2 Using Client Commands

# 2.2.1 Using pgAdmin

Follow the procedure below to define a database using pgAdmin.

- 1. In the pgAdmin window, right-click [Databases] in the [Browser] pane, and then click [Create] >> [Database] to display a [Create Database] window.
- 2. Specify appropriate values for the following items in the [Create Database] window.
  - [General] tab
    - [Database]: Name of the database to be managed
- 3. Click [Save] to create the database.

### 2.2.2 Using Client Commands

Follow the procedure below to define a database using client commands.

An example of operations on the server is shown below.

1. Use psql command to connect to the postgres database. Execute psql postgres.

```
> psql postgres
psql (12.1)
Type "help" for help.
```

2. Create the database.

To create the database, execute the CREATE DATABASE databaseName; statement.

```
postgres=# CREATE DATABASE db01;
CREATE DATABASE
```

3. Confirm that the database is created.

Execute \l+, and confirm that the name of the database created in step 2 is displayed.

```
postgres=# \l+
```

4. Disconnect from the postgres database.

Execute \q to terminate the psql command.

```
postgres=# \q
```

You can create a database using the createdb command.



Refer to "Creating a Database" in "Tutorial" in the PostgreSQL Documentation for information on creating a database using the createdb command.

# Chapter 3 Backing Up the Database

This chapter describes how to back up the database.

### Backup methods

The following backup methods enable you to recover data to a backup point or to the state immediately preceding disk physical breakdown or data logical failure.

- Backup using WebAdmin

This method enables you to back up data through intuitive window operations using the GUI.

WebAdmin is used for recovery.

- Backup using the pgx\_dmpall command

Execute the pgx\_dmpall command with a script to perform automatic backup.

To back up data automatically, you must register the process in the automation software of the operating system. Follow the procedure given in the documentation for your operating system.

The pgx\_rcvall command is used for recovery.



By using a copy command created by the user, the pgx\_dmpall command and the pgx\_rcvall command can back up database clusters and tablespaces to any destination and recover them from any destination using any copy method. Refer to "Chapter 13 Backup/Recovery Using the Copy Command" for details.

#### Approximate backup time

The formula for deriving the approximate backup time when you use WebAdmin or the pgx\_dmpall command is as follows:

backupTime = dataStorageDestinationUsage / diskWritePerformance x 1.5

- dataStorageDestinationUsage: Disk usage at the data storage destination
- diskWritePerformance: Maximum data volume (bytes/second) that can be written per second in the system environment where operation is performed
- 1.5: Coefficient to factor in tasks other than disk write (which is the most time-consuming step)

If using the copy command with the pgx\_dmpall command, the backup time will depend on the implementation of the copy command.



- Backup operation cannot be performed on an instance that is part of a streaming replication cluster in standby mode.
- Use the selected backup method continuously.

There are several differences, such as the data format, across the backup methods. For this reason, the following restrictions apply:

- It is not possible to use one method for backup and another for recovery.
- It is not possible to convert one type of backup data to a different type of backup data.
- Mirrored WALs can be used only for backup/recovery using the pgx\_dmpall command or WebAdmin.
- There are several considerations for the backup of the keystore and backup of the database in case the data stored in the database is encrypted. Refer to the following for details:
  - 5.6.4 Backing Up and Recovering the Keystore
  - 5.7 Backing Up and Restoring/Recovering the Database

- If you have defined a tablespace, back it up. If you do not back it up, directories for the tablespace are not created during recovery, which may cause the recovery to fail. If the recovery fails, refer to the system log, create the tablespace, and then perform the recovery process again.



The following methods can also be used to perform backup. Performing a backup using these methods allows you to restore to the point when the backup was performed.

- Backup using an SQL-based dump

Dump the data by using SQL. This backup method also enables data migration.

- File system level backup

This backup method requires you to stop the instance and use OS commands to backup database resources as files.

- Backup by continuous archiving

This is the standard backup method for PostgreSQL.

Refer to "Backup and Restore" in "Server Administration" in the PostgreSQL Documentation for information on these backup methods.

# 3.1 Periodic Backup

It is recommended that you perform backup periodically.

Backing up data periodically using WebAdmin or the pgx\_dmpall command has the following advantages:

- This method reduces disk usage, because obsolete archive logs (transaction logs copied to the backup data storage destination) are deleted. It also minimizes the recovery time when an error occurs.

### Backup cycle

The time interval when backup is performed periodically is called the backup cycle. For example, if backup is performed every morning, the backup cycle is 1 day.

The backup cycle depends on the jobs being run, but on FUJITSU Enterprise Postgres it is recommended that operations are run with a backup cycle of at least once per day.

# 3.2 Backup Methods

This section describes the methods for backing up the database.

- 3.2.1 Using WebAdmin
- 3.2.2 Using Server Commands

# 3.2.1 Using WebAdmin

You can use WebAdmin to perform backup and check the backup status.



- If backup is disabled for an instance, you will not be able to back up or restore the instance. Refer to "[Backup]" in "Creating an Instance" in the Installation and Setup Guide for Server for details.

- If the data to be stored in the database is to be encrypted, it is necessary to enable the automatic opening of the keystore before doing so. Refer to "5.6.3 Enabling Automatic Opening of the Keystore" for details.

- WebAdmin uses the labels "Data storage path", "Backup storage path" and "Transaction log path" to indicate "data storage destination", "backup data storage destination" and "transaction log storage destination" respectively. In this manual these terms are used interchangeably.

#### **Backup operation**

Follow the procedure below to back up the database.

1. Select the database to back up

In the [Instances] tab, select the instance to be backed up and click ...

2. Back up the database

The [Backup] dialog box is displayed. To perform backup, click [Yes]. An instance is automatically started when backup is performed.

### **Backup status**

If an error occurs and backup fails, [Error] is displayed adjacent to [Data storage status] or [Backup storage status] in the [Instances] tab. An error message is also displayed in the message list.

In this case, the backup data is not optimized. Ensure that you check the backup result whenever you perform backup. If backup fails, [Solution] appears to the right of the error message. Clicking this button displays information explaining how to resolve the cause of the error. Remove the cause of failure, and perform backup again.

### 3.2.2 Using Server Commands

Use the pgx\_dmpall command and pgx\_rcvall command to perform backup and check the backup result.

### Preparing for backup

You must prepare for backup before actually starting the backup process.

Follow the procedure below.



#### See

Refer to "Preparing Directories to Deploy Resources" in the Installation and Setup Guide for Server for information on the location of directories required for backup and for points to take into account.

1. Prepare the backup data storage disk

For backup, prepare a separate disk unit from the database storage disk and mount it using the operating system commands.

2. Create a directory where the backup data will be stored

Create an empty directory.

Set appropriate permissions so that only the instance administrator can access the directory.

### Example

```
# mkdir /backup/inst1
# chown fsepuser:fsepuser /backup/inst1
# chmod 700 /backup/inst1
```

3. Specify the settings required for backup

Stop the instance, and set the following parameters in the postgresql.conf file.

Start the instance after editing the postgresql.conf file.

Parameter name	Setting	Description
backup_destination	Name of the directory where the backup data will be stored	Specify the name of the directory where the backup data will be stored.
		Appropriate privileges that allow only the instance administrator to access the directory must already be set.
		Place the backup data storage destination directory outside the data storage destination directory, the tablespace directory, and the transaction log storage destination directory.
archive_mode	on	Specify the archive log mode.  Specify [on] (execute).
archive_command	'installationDirectory/bin/ pgx_walcopy.cmd "%p" "backupDataStorageDestinationDirectory/ archived_wal/%f"	Specify the path name of the command that will save the transaction log and the storage destination.

Refer to "Appendix A Parameters" and "Write Ahead Log" under "Server Administration" in the PostgreSQL Documentation for information on the parameters.

### Backup operation (file backup)

Use the pgx\_dmpall command to perform file backup. You can even embed the pgx\_dmpall command in OS automation software to perform backup.

The backup data is stored in the directory specified in the backup\_destination parameter of postgresql.conf.

Specify the data storage destination in the -D option. If the -D option is omitted, the value of the PGDATA environment variable is used by default.



> pgx\_dmpall -D /database/inst1



Backup stores the data obtained during the backup and the backup data of the data obtained during previous backup.

If the data to be stored in the database is encrypted, refer to the following and back up the keystore:

- 5.6.4 Backing Up and Recovering the Keystore

### **Backup status**

Use the pgx\_rcvall command to check the backup status.

Specify the following values in the pgx\_rcvall command:

- The -l option indicates backup data information.
- Specify the data storage destination in the -D option. If the -D option is omitted, the value of the PGDATA environment variable is used by default.

```
> pgx_rcvall -1 -D /database/inst1
Date Status Dir
2017-05-01 13:30:40 COMPLETE /backup/inst1/2017-05-01_13-30-40
```

If an error occurs and backup fails, a message is output to the system log.

In this case, the backup data is not optimized. Ensure that you check the backup result whenever you perform backup. If backup fails, remove the cause of failure and perform backup again.



Refer to "pgx\_dmpall" and "pgx\_rcvall" in the Reference for information on the pgx\_dmpall command and pgx\_rcvall command.

### Setting a restore point

In case you want to recover your database to a certain point in time, you can name this particular point in time, which is referred to as the restore point, by using the psql command.

By setting a restore point before executing an application, it becomes easy to identify up to which point in time the data will be reverted.

A restore point can be set to any point in time after a backup is executed. However, if a restore point is set before a backup is executed, the database cannot be recovered to that point in time. This is because restore points are recorded in the archive logs, and the archive logs are discarded when backups are executed.



The following example uses the psql command to connect to the database and execute the SQL statement to set a restore point.

However, when considering continued compatibility of applications, do not use functions directly in SQL statements. Refer to "Notes on Application Compatibility" in the Application Development Guide for details.

Refer to "14.3.2 Using the pgx\_rcvall Command" for information on using a restore point to recover the database.



- Name restore points so that they are unique within the database. Add the date and time of setting a restore point to distinguish it from other restore points, as shown below:

- YYMMDD\_HHMMSS
  - YYMMDD: Indicates the date
  - HHMMSS: Indicates the time
- There is no way to check restore points you have set. Keep a record in, for example, a file.



Refer to "System Administration Functions" under "Functions and Operators" in the PostgreSQL Documentation for information on pg\_create\_restore\_point.

# Chapter 4 Configuring Secure Communication Using Secure Sockets Layer

If communication data transferred between a client and a server contains confidential information, encrypting the communication data can protect it against threats, such as eavesdropping on the network.

# 4.1 Configuring Communication Data Encryption

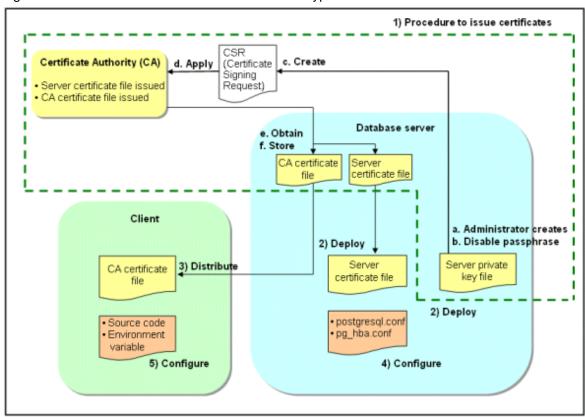
To encrypt communication data transferred between a client and a server, configure communication data encryption as described below. Communication data encryption not only protects the communication content, but it also guards against man-in-the-middle (MITM) attacks (for example, data and password theft through server impersonation).

Table 4.1 Configuration procedure

Configuration procedure		
1) Issue a certificate		
2) Deploy a server certificate file and a server private key file		
3) Distribute a CA certificate file to the client		
4) Configure the operating environment for the database server		
5) Configure the operating environment for the client		

The following figure illustrates the environment for communication data encryption.

Figure 4.1 Environment for communication data encryption



### 4.1.1 Issuing a Certificate

For authenticating servers, you must acquire a certificate issued by the certificate authority (CA).

FUJITSU Enterprise Postgres supports X.509 standard PEM format files. If the certificate authority issues a file in DER format, use a tool such as the openssl command to convert the DER format file to PEM format.

The following provides an overview of the procedure. Refer to the procedure published by the public or independent certificate authority (CA) that provides the certificate file for details.

- a. Create a server private key file
- b. Disable the passphrase for the server private key file
- c. Create a CSR (signing request for obtaining a server certificate) from the server private key file
- d. Apply to the certificate authority (CA) for a server certificate
- e. Obtain a server certificate file and a CA certificate file from the certificate authority (CA)
- f. Store the server certificate file and the CA certificate file Note: If you lose or destroy the certificates, you will need to have them re-issued.

The above procedure enables you to prepare the following files:

- Server private key file
- Server certificate file
- CA certificate file

### 4.1.2 Deploying a Server Certificate File and a Server Private Key File

Create a directory on the local disk of the database server and store the server certificate file and the server private key file in it. Use the operating system features to set access privileges for the server certificate file and the server private key file so that only the database administrator has load privileges.

Back up the server certificate file and the server private key file in the event that data corruption occurs and store them securely.

# 4.1.3 Distributing a CA Certificate File to the Client

Create a directory on the local disk of the client and place the distributed CA certificate file there. Use the operating system features to set load privileges to protect the CA certificate file against accidental deletion.

# 4.1.4 Configuring the Operating Environment for the Database Server



Refer to "Secure TCP/IP Connections with SSL" under "Server Administration" in the PostgreSQL Documentation for details.

# 4.1.5 Configuring the Operating Environment for the Client



Refer to the following sections in the Application Development Guide for details, depending on your application development environment:

- "Settings for Encrypting Communication Data" under "Setup" in "JDBC Driver"
- "Settings for Encrypting Communication Data" under "Setup" in "C Library (libpq)"
- "Settings for Encrypting Communication Data" under "Setup" in "Embedded SQL in C"

## 4.1.6 Performing Database Multiplexing

When you perform communication that uses database multiplexing and a Secure Socket Layer server certificate, certificates with the same "Common Name" must be used. To ensure this, take one of the following actions:

- Create one server certificate, replicate it, and place a copy on each server used for database multiplexing.
- Create a server certificate with the same "Common Name" for each server used for database multiplexing.



Refer to "Using the Application Connection Switch Feature" in the Application Development Guide for information on how to specify applications on the client.

......

# Chapter 5 Protecting Storage Data Using Transparent Data Encryption

This chapter describes how to encrypt data to be stored in the database.

# 5.1 Protecting Data Using Encryption

With PostgreSQL, data in a database is protected from access by unauthorized database users through the use of authentication and access controls. However, the OS file is not protected from attackers who bypass the database server's authentication and access controls.

With FUJITSU Enterprise Postgres, data inside the OS file is encrypted, so valuable information is protected even if the file or disk is stolen.

Data to be stored in a database is encrypted when it is written to the data file, and decrypted when it is read.

This is performed automatically by the instance, so the user and the application need not be aware of key management and encryption or decryption. This process is called TDE (Transparent Data Encryption).

The characteristics of TDE are described below.

### **Encryption mechanisms**

### Two-layer encryption key and the keystore

In each tablespace, there is a tablespace encryption key that encrypts and decrypts all the data within. The tablespace encryption key is encrypted by the master encryption key and saved.

Only one master encryption key exists in a database cluster. It is encrypted based on a passphrase specified by the user and stored in a keystore. FUJITSU Enterprise Postgres provides a file-based keystore. Attackers who do not know the passphrase cannot read the master encryption key from the keystore.

#### Strong encryption algorithms

TDE uses the Advanced Encryption Standard (AES) as its encryption algorithm. AES was adopted as a standard in 2002 by the United States Federal Government, and is used throughout the world.

### Faster encryption and decryption based on hardware

TDE minimizes the overhead of encryption and decryption by using the AES-NI (Advanced Encryption Standard New Instructions) built into Intel(R) Xeon(R) processors since the 5600 series. This means that even in situations where previously the minimum encryption target was selected as a tradeoff between performance and security, it is now possible to encrypt all the data of an application.

Refer to the Intel Corporation's website for information on the list of processors equipped with AES-NI.

### Zero overhead storage areas

Encryption does not change the size of data stored in tables, indexes, or WAL. There is, therefore, no need for additional estimates or disks

### Scope of encryption

#### All user data within the specified tablespace

The tablespace is the unit for specifying encryption. All tables, indexes, temporary tables, and temporary indexes created in the encrypted tablespace are encrypted. There is no need for the user to consider which tables and strings to encrypt.

Refer to "5.4 Encrypting a Tablespace" for details.

### Backup data

The pgx\_dmpall command and pg\_basebackup command create backup data by copying the OS file. Backups of the encrypted data are, therefore, also encrypted. Information is protected from leakage even if the backup medium is stolen.

#### WAL and temporary files

WAL, which is created by updating encrypted tables and indexes, is encrypted with the same security strength as the update target. When large merges and sorts are performed, the encrypted data is written to a temporary file in encrypted format.

### Streaming replication support

You can combine streaming replication and transparent data encryption. The data and WAL encrypted on the primary server is transferred to the standby server in its encrypted format and stored.



The following are not encrypted:

- pg\_dump and pg\_dumpall output files
- Files output by the COPY command
- Notification event payloads that communicate using the LISTEN or NOTIFY command
- Checksum validation is not performed on encrypted tablespaces during backup and when using the pg\_checksum utility.

# 5.2 Setting the Master Encryption Key

To use transparent data encryption, you must create a keystore and set the master encryption key.

1. In the keystore\_location parameter of postgresql.conf, specify the directory to store the keystore.

Specify a different location for each database cluster.

```
keystore_location = '/key/store/location'
```

Refer to "Appendix A Parameters" for information on postgresql.conf.

After editing the postgresql.conf file, either start or restart the instance.

- Using WebAdmin

Refer to "2.1.1 Using WebAdmin", and restart the instance.

- Using the pg\_ctl command

Specify the following in the pg\_ctl command:

- Specify "restart" as the mode.
- Specify the data storage destination directory in the -D option. If the -D option is omitted, the value of the PGDATA environment variable is used by default.
- Specify the -w option. This means that the command returns after waiting for the instance to start. If the -w option is not specified, it may not be possible to determine if the starting of the instance completed successfully or if it failed.

### Example

```
> pg_ctl restart -w -D /database/inst1
```

2. Execute an SQL function, such as the one below, to set the master encryption key. This must be performed by the superuser. Execute it as the database superuser.

```
SELECT pgx_set_master_key('passphrase');
```

The value "passphrase" is the passphrase that will be used to open the keystore. The master encryption key is protected by this passphrase, so avoid specifying a short simple string that is easy to guess.

Refer to "B.2 Transparent Data Encryption Control Functions" for information on the pgx\_set\_master\_key function.



Note that if you forget the passphrase, you will not be able to access the encrypted data. There is no method to retrieve a forgotten passphrase and decrypt data. Do not, under any circumstances, forget the passphrase.

The pgx\_set\_master\_key function creates a file with the name keystore.ks in the keystore storage destination. It also creates a master encryption key from random bit strings, encrypts it with the specified passphrase, and stores it in keystore.ks. At this point, the keystore is open.

# 5.3 Opening the Keystore

To create encrypted tablespaces and access the encrypted data, you must first open the keystore. When you open the keystore, the master encryption key is loaded into the database server memory and becomes usable for encryption and decryption.

You need to open the keystore each time you start the instance. To open the keystore, the database superuser must execute the following SQL function.

```
SELECT pgx_open_keystore('passphrase');
```

The value "passphrase" is the passphrase specified during creation of the keystore.

Refer to "B.2 Transparent Data Encryption Control Functions" for information on the pgx\_open\_keystore function.

Note that, in the following cases, the passphrase must be entered when starting the instance, because the encrypted WAL must be decrypted for recovery. In this case, the above-mentioned pgx\_open\_keystore function cannot be executed.

- If performing crash recovery at the time of starting the instance
- If performing recovery using continuous archiving

For the above cases, specify the --keystore-passphrase option in the pg\_ctl command, and then start the instance. This will display the prompt for the passphrase to be entered, as shown below.

```
> pg_ctl --keystore-passphrase start
Enter the passphrase:
The server is starting
>
```



When using an automatically opening keystore, you do not need to enter the passphrase and you can automatically open the keystore when the database server starts. Refer to "5.6.3 Enabling Automatic Opening of the Keystore" for details.

# 5.4 Encrypting a Tablespace

The keystore must be open before you can create an encrypted tablespace.

When creating a tablespace that will be encrypted, configure the encryption algorithm in the runtime parameters. For example, to create a tablespace with the name secure\_tablespace using AES with a key length of 256 bits as the encryption algorithm, configure as shown below.

```
-- Specify the encryption algorithm for the tablespace to be created below

SET tablespace_encryption_algorithm = 'AES256';

CREATE TABLESPACE secure_tablespace LOCATION '/My/Data/Dir';

-- Specify that the tablespace to be created below is not to be encrypted

SET tablespace_encryption_algorithm = 'none';
```

Or

```
CREATE TABLESPACE secure_tablespace LOCATION '/My/Data/Dir' WITH (tablespace_encryption_algorithm = 'AES256' );
```

When the tablespace is empty, the encryption algorithm can be modified with the command below.

```
ALTER TABLESPACE secure_tablespace SET (tablespace_encryption_algorithm=AES256);
```

Trying to set the encryption algorithm for a non-empty tablespace causes an error.

You can use AES with a key length of 128 bits or 256 bits as the encryption algorithm. It is recommended that you use 256-bit AES. Refer to "Appendix A Parameters" for information on how to specify the runtime parameters.

If user provides both GUC and command line options while creating the tablespace, the preference is given to the command line option.

The pg\_default and pg\_global tablespaces cannot be encrypted.

Create tables and indexes in the encrypted tablespace that you created. Relations created in the encrypted tablespace are automatically encrypted.



Example 1: Specifying an encrypted tablespace when creating it

```
CREATE TABLE my_table (...)

TABLESPACE secure_tablespace;
```

Example 2: Not explicitly specifying a tablespace when creating it and instead using the default tablespace

```
SET default_tablespace = 'secure_tablespace';
CREATE TABLE my_table (...);
```

The process is the same for encrypting temporary tables and temporary indexes. In other words, either explicitly specify the TABLESPACE clause or list encrypted tablespaces in the temp\_tablespaces parameter, and then execute CREATE TEMPORARY TABLE or CREATE INDEX.



If an encrypted tablespace is specified in the TABLESPACE clause of the CREATE DATABASE statement, relations created in the database without explicitly specifying a tablespace will be encrypted. Furthermore, the system catalog will also be encrypted, so the source code of user-defined functions is also protected.

Example: Specifying a tablespace in a database definition statement

```
CREATE DATABASE DB01 TABLESPACE=SP01 ... ;
```

Part of the data is also stored in the system catalog - to encrypt this data as well, specify an encrypted tablespace as above and create a database.



An encrypted tablespace cannot be created from the window used for creating the pgAdmin tablespace, or from the query tool. To create an encrypted tablespace, click [PSQL Console] from the [Plugins] menu and create an encrypted tablespace in the psql console window.

# 5.5 Checking an Encrypted Tablespace

The pgx\_tablespaces system view displays information about whether each tablespace has been encrypted, and about the encryption algorithm. Refer to "C.1 pgx\_tablespaces" for information on strings.

You can discover which tablespaces have been encrypted by executing the following SQL statements.

However, when considering continued compatibility of applications, do not reference system catalogs (pg\_tablespace) directly in SQL statements.

```
SELECT spcname, spcencalgo
FROM pg_tablespace ts, pgx_tablespaces tsx
WHERE ts.oid = tsx.spctablespace;
```





Refer to "Notes on Application Compatibility" in the Application Development Guide for information on how to maintain application compatibility.

# 5.6 Managing the Keystore

This section describes how to manage the keystore and the master encryption key to guard against the threat of theft.

### 5.6.1 Changing the Master Encryption Key

Using the same encryption key for an extended period gives attackers an opportunity to decipher the encrypted data. It is recommended that you change the key at regular intervals, or whenever the key is exposed to risk.

Adhere to the industry's best practices for encryption algorithms and key management when considering how often the key should be changed. For example, the NIST in the United States has published "NIST Special Publication 800-57". The PCI DSS also refers to this publication. This publication recommends changing the master encryption key once a year.

To change the master encryption key, execute the pgx\_set\_master\_key function, which is the same function used for configuring the key. Refer to "5.2 Setting the Master Encryption Key" for details.

After changing the master encryption key, you must immediately back up the keystore.

# 5.6.2 Changing the Keystore Passphrase

In security policies for organizations, it is usually a requirement that the passphrase be changed whenever a security administrator who knows the passphrase is removed from duties due to transfer or retirement. It is also recommended that the passphrase be changed if it is ever exposed to risks due to deception such as social engineering.

To change the keystore passphrase, execute the following SQL function as a superuser.

```
SELECT pgx_set_keystore_passphrase('oldPassphrase', 'newPassphrase');
```

After changing the passphrase, you must immediately back up the keystore.

Refer to "B.2 Transparent Data Encryption Control Functions" for information on the pgx\_set\_keystore\_passphrase function.

# 5.6.3 Enabling Automatic Opening of the Keystore

When using an automatically opening keystore, you do not need to enter the passphrase and you can automatically open the keystore when the instance starts. Execute the pgx\_keystore command to enable automatic opening of the keystore.

```
> pgx_keystore --enable-auto-open /key/store/location/keystore.ks
Enter the passphrase:
Automatic opening of the keystore is now enabled
>
```



Refer to "pgx\_keystore" in the Reference for information on pgx\_keystore command.

When automatic opening is enabled, an automatically opening keystore is created in the same directory as the original keystore. The file name of the automatically opening keystore is keystore.aks. The file keystore.aks is an obfuscated copy of the decrypted content of the keystore.ks file. As long as this file exists, there is no need to enter the passphrase to open the keystore when starting the instance.

Do not delete the original keystore file, keystore.ks. It is required for changing the master encryption key and the passphrase. When you change the master encryption key and the passphrase, keystore.aks is recreated from the original keystore file, keystore.ks.

Protect keystore.ks, keystore.aks, and the directory that stores the keystore so that only the user who starts the instance can access them.

Configure the permission of the files so that only the user who starts the instance can access the SQL functions and commands that create these files. Accordingly, manually configure the same permission mode if the files are restored.



```
# chown -R fsepuser:fsepuser /key/store/location
# chmod 700 /key/store/location
```

# chmod 600 /key/store/location/keystore.ks

# chmod 600 /key/store/location/keystore.aks

An automatically opening keystore will only open on the computer where it was created.

To disable automatic opening of the keystore, delete keystore.aks.



- To use WebAdmin for recovery, you must enable automatic opening of the keystore.
- Refer to "5.7 Backing Up and Restoring/Recovering the Database" after enabling or reconfiguring encryption to back up the database.

••••••

- Specify a different directory from those below as the keystore storage destination:
  - Data storage destination
  - Tablespace storage destination
  - Transaction log storage destination
  - Backup data storage destination

# 5.6.4 Backing Up and Recovering the Keystore

Back up the keystore at the following times in case it is corrupted or lost. Note that you must store the database and the keystore on separate data storage media. Storing both on the same data storage medium risks the danger of the encrypted data being deciphered if the medium is stolen. A passphrase is not required to open an automatically opening keystore, so store this type of keystore in a safe location.

- When the master encryption key is first configured
- When the master encryption key is changed
- When the database is backed up
- When the keystore passphrase is changed



Do not overwrite an old keystore when backing up a keystore. This is because during database recovery, you must restore the keystore to its state at the time of database backup. When the backup data of the database is no longer required, delete the corresponding keystore.

# **Example**

- Back up the database and the keystore on May 1, 2017.

```
> pgx_dmpall -D /database/inst1
> cp -p /key/store/location/keystore.ks /keybackup/keystore_20170501.ks
```

Specify the following in the pgx\_dmpall command:

- Specify the data storage destination in the -D option. If the -D option is omitted, the value of the PGDATA environment variable is used by default.
- Change the master encryption key, and back up the keystore on May 5, 2017.

```
> psql -c "SELECT pgx_set_master_key('passphrase')" postgres
> cp -p /key/store/location/keystore.ks /keybackup/keystore_20170505.ks
```

Specify the following in the psql command:

- Specify the SQL function that sets the master encryption key in the -c option.
- Specify the name of the database to be connected to as the argument.

If the keystore is corrupted or lost, restore the keystore containing the latest master encryption key. If there is no keystore containing the latest master encryption key, restore the keystore to its state at the time of database backup, and recover the database from the database backup. This action recovers the keystore to its latest state.



- Restore the keystore containing the latest master encryption key as of May 5, 2017.

```
> cp -p /keybackup/keystore_20170505.ks /key/store/location/keystore.ks
```

- If there is no backup of the keystore containing the latest master encryption key, recover the keystore by restoring the keystore that was backed up along with the database on 1 May 2017.

```
> cp -p /keybackup/keystore_20170501.ks /key/store/location/keystore.ks
> pgx_rcvall -B /backup/inst1 -D /database/inst1 --keystore-passphrase
```

Specify the following in the pgx\_rcvall command:

- Specify the data storage directory in the -D option. If the -D option is omitted, the value of the PGDATA environment variable is used by default.
- Specify the backup data storage directory in the -B option.
- The --keystore-passphrase option prompts you to enter the passphrase to open the keystore.

If you have restored the keystore, repeat the process of enabling automatic opening of the keystore. This ensures that the contents of the automatically opening keystore (keystore.aks) are identical to the contents of the restored keystore.

It is recommended that you do not back up the automatically opening keystore file, keystore.aks. If the database backup medium and the backup medium storing the automatically opening keystore are both stolen, the attacker will be able to read the data even without knowing the passphrase.

If the automatically opening keystore is corrupted or lost, you must again enable automatic opening. The keystore aks file will be recreated from keystore.ks at this time.



Refer to "pgx\_rcvall" and "pgx\_dmpall" in the Reference for information on the pgx\_rcvall and pgx\_dmpall commands.

Refer to "psql" under "Reference" in the PostgreSQL Documentation for information on the psql command.

Refer to "B.2 Transparent Data Encryption Control Functions" for information on the pgx\_set\_master\_key function.

Refer to "5.6.3 Enabling Automatic Opening of the Keystore" for information on how to enable automatic opening of the keystore.

# 5.7 Backing Up and Restoring/Recovering the Database

FUJITSU Enterprise Postgres enables you to use the five backup and recovery methods described below. Regardless of the method you use, you must back up the keystore at the same time.

Note that you must store the database and the keystore on separate data storage media. Storing both on the same data storage medium risks the danger of the encrypted data being deciphered if the medium is stolen.

#### Backup and recovery using WebAdmin

- Backup

WebAdmin backs up encrypted data.

Back up the key store after backing up the database.

- Recovery

Restore the keystore to its state at the time of database backup. Refer to "5.6.4 Backing Up and Recovering the Keystore" for details.

Enable automatic opening of the keystore in accordance with the procedure described in "5.6.3 Enabling Automatic Opening of the Keystore". Then, use WebAdmin to recover the database.

#### Backup and recovery using the pgx\_dmpall and pgx\_rcvall commands

- Backup

The pgx\_dmpall command backs up the encrypted data.

Back up the key store after backing up the database.

- Recovery

Restore the keystore to its state at the time of the database backup.

Configure automatic opening of the key store as necessary.

If automatic opening of the keystore is not enabled, execute the pgx\_rcvall command with the --keystore-passphrase option specified. This will display the prompt for the passphrase to be entered.



- Back up the database and the keystore on May 1, 2017.

```
> pgx_dmpall -D /database/inst1
> cp -p /key/store/location/keystore.ks /keybackup/keystore_20170501.ks
```

Specify the following in the pgx\_dmpall command:

- Specify the data storage destination in the -D option. If the -D option is omitted, the value of the PGDATA environment variable is used by default.
- Recover the database and the keystore from the backup taken on May 1, 2017.

```
> cp -p /keybackup/keystore_20170501.ks /key/store/location/keystore.ks
> pgx_keystore --enable-auto-open /key/store/location/keystore.ks (Execute only when enabling automatic opening)
> pgx_rcvall -B /backup/inst1 -D /database/inst1 --keystore-passphrase
```

Specify the following in the pgx\_rcvall command:

- Specify the data storage destination in the -D option. If the -D option is omitted, the value of the PGDATA environment variable is used by default.
- Specify the backup data storage directory in the -B option.
- The --keystore-passphrase option prompts you to enter the passphrase to open the keystore.

#### Dump and restore using SQL

- Backup

The files output by the pg\_dump and pg\_dumpall commands are not encrypted. You should, therefore, encrypt the files using OpenSSL commands or other means before saving them, as described in "5.8 Importing and Exporting the Database" below.

Back up the key store after backing up the database.

- Restore

If the backup data has been encrypted using, for example Open SSL commands, decrypt that data.

The data generated by the pg\_dumpall command includes a specification to encrypt tablespaces by default. For this reason, the psql command encrypts tablespaces during restoration.

### File system level backup and restore

- Backup

Stop the instance and backup the data directory and the tablespace directory using the file copy command of the operating system. The files of encrypted tablespaces are backed up in the encrypted state.

Back up the key store after performing the backup.

- Restore

Restore the keystore to its state at the time of the database backup.

Stop the instance and restore the data directory and the tablespace directory using the file copy command of the operating system.

### Continuous archiving and point-in-time recovery

- Backup

The pg\_basebackup command backs up the encrypted data as is.

Back up the key store after performing the backup.

- Recovery

Restore the keystore to its state at the time of the database backup.

Configure automatic opening of the key store as necessary.

If automatic opening of the keystore is not enabled, execute the pg\_ctl command to start the instance with the --keystore-passphrase option specified. This will display the prompt for the passphrase to be entered.

.....



- Refer to "pg\_ctl" under "Reference" in the PostgreSQL Documentation for information on the pg\_ctl command.
- Refer to "Reference" in the PostgreSQL Documentation for information on the following commands:
  - psql
  - pg\_dump
  - pg\_basebackup
- Refer to the Reference for information on the following commands:
  - pgx\_rcvall

- pgx\_dmpall
- pg\_dumpall

If you have restored the keystore, repeat the process of enabling automatic opening of the keystore This ensures that the contents of the automatically opening keystore (keystore.aks) are identical to the contents of the restored keystore.

Refer to "5.6.3 Enabling Automatic Opening of the Keystore" for information on how to enable automatic opening of the keystore.

## 5.8 Importing and Exporting the Database

The files output by the COPY TO command are not encrypted. Therefore, when transferring files to other systems, you should encrypt files using OpenSSL commands or other means and use scp or sftp to encrypt the data being transferred.

Use a safe method to delete obsolete plain text files.

You can use the following methods to safely delete files:

- shred command



```
# Export the contents of the table my_table to a CSV file.
> psql -c "COPY my_table TO '/tmp/my_table.csv' (FORMAT CSV)" postgres

# Encrypt the exported file.
> openssl enc -e -aes256 -in my_table.csv -out my_table.csv.enc
(The user is prompted to enter the passphrase to be used for encryption)

# Safely delete plain text files.
> shred -u -x my_table.csv
(Transfer encrypted files to other systems)

# Decrypt the encrypted files on other systems.
> openssl enc -d -aes256 -in my_table.csv.enc -out my_table.csv
(The user is prompted to enter the passphrase to be used for decryption)
```

If you use COPY FROM to import data to tables and indexes in an encrypted tablespace, the imported data is automatically encrypted before being stored.

# 5.9 Encrypting Existing Data

You cannot encrypt existing unencrypted tablespaces. In addition, you cannot change encrypted tablespaces so that they do not encrypt.

As an alternative, transfer the tables and indexes to other tablespaces. You can use the following SQL commands for this.

```
ALTER TABLE table_name SET TABLESPACE new_tablespace;
ALTER INDEX index_name SET TABLESPACE new_tablespace;
ALTER DATABASE database_name SET TABLESPACE new_tablespace;
```



Refer to "SQL Commands" under "Reference" in the PostgreSQL Documentation for information on SQL commands.

### 5.10 Operations in Cluster Systems

This section describes how to use transparent data encryption on cluster systems such as high-availability systems, streaming replication, and database multiplexing.

### 5.10.1 HA Clusters that do not Use Database Multiplexing

Take the following points into account when using transparent data encryption in an HA cluster environment that does not use database multiplexing.

#### Placement and automatic opening of the keystore file

There are two alternatives for placing the keystore file:

- Sharing the keystore file
- Placing a copy of the keystore file

#### Sharing the keystore file

This involves using the same keystore file on the primary server and the standby server.

As the standby server is not active while the primary server is running, this file would not be accessed simultaneously, and therefore, it can be shared.

To manage the keystore file in a more secure manner, place it on the key management server or the key management storage isolated in a secure location.

Enable the automatic opening of the keystore on both the primary and standby servers.

#### Placing a copy of the keystore file

This involves placing a copy of the primary server keystore file on the standby server.

You can do this if you cannot prepare a shared server or disk device that can be accessed from both the primary and standby servers.

However, if you change the master encryption key and the passphrase on the primary server, you must copy the keystore file to the standby server again.

To manage the keystore file in a more secure manner, prepare the key management server or the key management storage isolated in a secure location for both the primary and standby servers, and place the keystore files there.

Enable the automatic opening of the keystore on both the primary and standby servers. Note that copying the automatically opening keystore file (keystore.aks) to the standby server does not enable the automatic opening of the keystore.



Refer to the Cluster Operation Guide (PRIMECLUSTER) for information on building a cluster system environment for performing failover using the failover feature integrated with the cluster software.

# 5.10.2 Database Multiplexing Mode

Note the following when using transparent data encryption in environments that use streaming replication, or database multiplexing with streaming replication.

#### Placing the keystore file

Place a copy of the primary server keystore file on the standby server.

This is required as the keystore file cannot be shared, and both servers may need to access it simultaneously.



To manage the keystore file in a more secure manner, place it on the key management server or the key management storage isolated in a secure location. A keystore used by both the primary and standby servers can be managed on the same key management server or key management storage.

However, create different directories for the keystores to be used by the primary server and the standby server. Then copy the keystore for the primary server to the directory used on the standby server.

#### Automatically opening the keystore

You must enable automatic opening of the keystore.

To do this, enable automatic opening of the keystore in all servers that make up database multiplexing. The settings for automatic opening of the keystore include information unique to each server, so simply copying the file does not enable it.

#### Changing the passphrase

Changes to the passphrase are reflected in all servers that make up database multiplexing, so no special operation is required.

#### Building and starting a standby server

Before using the pg\_basebackup command or pgx\_rcvall command to build a standby server, copy the keystore file from the primary server to the standby server. When using an automatically opening keystore, use the copied keystore file to enable automatic opening on the standby server.

Open the keystore each time you start the standby server. This step is necessary for decrypting and restoring encrypted WAL received from the primary server. To open the keystore, specify the --keystore-passphrase option in the pg\_ctl command or pgx\_rcvall command and enter the passphrase, or use an automatically opening keystore.

#### Changing the master encryption key and the passphrase

Change the master encryption key and the passphrase on the primary server. You need not copy the keystore from the primary server to the standby server. You need not even restart the standby server or reopen the keystore. Changes to the master encryption key and the passphrase are reflected in the keystore on the standby server.

.....



Refer to "pgx\_rcvall" in the Reference for information on pgx\_rcvall command.

Refer to "pg\_ctl" under "Reference" in the PostgreSQL Documentation for information on pg\_ctl command.

Refer to "pg\_basebackup" under "Reference" in the PostgreSQL Documentation for information on pg\_basebackup command.

Refer to "High Availability, Load Balancing, and Replication" under "Server Administration" in the PostgreSQL Documentation for information on how to set up streaming replication.

# 5.11 Security-Related Notes

- Decrypted data is cached in the database server memory (shared buffer). As a result, unencrypted data is stored in a core file, which is a process memory dump. You should, therefore, safely delete the memory dump.

  You can safely delete files by using the following command:
  - shred command
- Unencrypted data may be written from the database server memory to the operating system's swap area. To prevent leakage of information from the swap area, consider either disabling the use of swap area or encrypting the swap area using a full-disk encryption product.
- The content of the server log file is not encrypted. Therefore, in some cases the value of a constant specified in a SQL statement is output to the server log file. To prevent this, consider setting a parameter such as log\_min\_error\_statement.
- When executing an SQL function that opens the keystore and modifies the master encryption key, ensure that the SQL statement containing the passphrase is not output to the server log file. To prevent this, consider setting a parameter such as log\_min\_error\_statement. If you are executing this type of SQL function on a different computer from the database server, encrypt the communication between the client and the database server with SSL.
- Starting with FEP 10, logical replication is available, which allows non-backed up clusters to subscribe to databases where transparent data encryption is enabled. Logical replication does not need to have the same encryption strategy between publisher and subscriber.
  - In this scenario, if the user wants to encrypt the subscribed copy of data as well, then it is the user's responsibility to create encryption policies to the subscribed databases. By default, published encrypted tablespace data will not be encrypted in the subscriber side.

# 5.12 Tips for Installing Built Applications

With transparent data encryption, you can easily encrypt all the data in an application without modifying the application. Database administrators install built applications in the following manner. However, this procedure stores data to the default tablespace, so take necessary action if processing differs from the original design.

1. (Normal procedure) Create an owner and a database for the built application.

```
CREATE USER crm_admin ...;
CREATE DATABASE crm_db ...;
```

2. (Procedure for encryption) Create an encrypted tablespace to store the data for the built application.

```
SET tablespace_encryption_algorithm = 'AES256';
CREATE TABLESPACE crm_tablespace LOCATION '/crm/data';
```

3. (Procedure for encryption) Configure an encrypted tablespace as the default tablespace for the owner of the built application.

```
ALTER USER crm_admin SET default_tablespace = 'crm_tablespace';
ALTER USER crm_admin SET temp_tablespaces = 'crm_tablespace';
```

4. (Normal procedure) Install the built application. The application installer prompts you to enter the host name and the port number of the database server, the user name, and the database name. The installer uses the entered information to connect to the database server and execute the SQL script. For applications that do not have an installer, the database administrator must manually execute the SQL script.

Normally, the application's SQL script includes logic definition SQL statements, such as CREATE TABLE, CREATE INDEX, and GRANT or REVOKE, converted from the entity-relationship diagram. It does not include SQL statements that create databases, users, and tablespaces. Configuring the default tablespace of the users who will execute the SQL script deploys the objects generated by the SQL script to the tablespace.

# Chapter 6 Data Masking

Data masking is a feature that can change the returned data for queries generated by applications, so that it can be referenced by users. For example, for a query of employee data, digits except the last four digits of an eight-digit employee number can be changed to "\*" so that it can be used for reference.

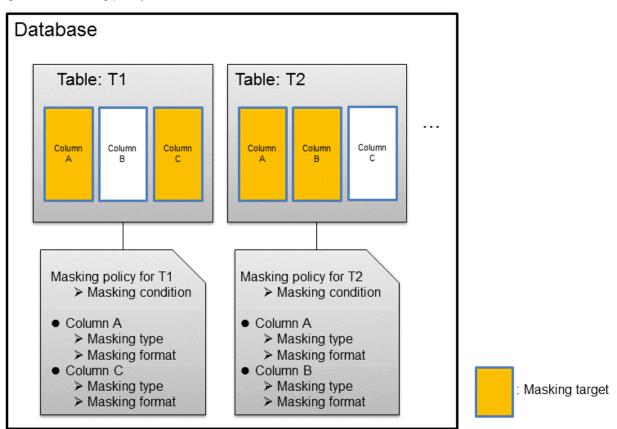


When using this feature, it is recommended that the changed data be transferred to another medium for users to reference. This is because, if users directly access the database to extract the masked data, there is a possibility that they can deduce the original data by analyzing the masking policy or query result to the masking target column.

# 6.1 Masking Policy

Masking policy is a method of changing data under specific conditions when it is returned for a query from an application. One masking policy can be created per table. You can configure masking target, masking type, masking condition and masking format in a masking policy.

Figure 6.1 Masking policy





When a masking policy is defined, the search performance for the corresponding table may deteriorate.

### 6.1.1 Masking Target

Masking target refers to a column to which a masking policy will be applied. When referring to a masking target or a function that includes a masking target, the execution result will be changed and obtained.

The following commands can change the execution result:

- SELECT
- COPY
- pg\_dump
- pg\_dumpall



- If a masking target is specified to INSERT...SELECT target columns, processing will be performed using data before change.
- If a masking target other than SELECT target columns is specified, processing will be performed using data before change.
- If a masking target is specified in a function where the data type will be converted, an error will occur.

### 6.1.2 Masking Type

Masking type is a method to change column data that is returned from queries. Specify the masking type in the function\_type parameter. The following masking types can be specified and selected depending on the masking target data type.

#### **Full masking**

All the data in the specified column is changed. The changed value returned to the application that made the query varies depending on the column data type.

For example, 0 is used for a numeric type column and a space is used for a character type column.

### Partial masking

The data in the specified column is partially changed.

For example, digits except the last four digits of an employee number can be changed to "\*".

#### Regular expression masking

The data in the specified column is changed via a search that uses a regular expression.

For example, for strings such as email address that can have variable length, "\*" can be used to change characters preceding "@" by using a regular expression. Regular expression masking can only be used for character type data.



- If multiple valid masking targets are specified for a function, the masking type for the left-most masking target will be applied. For example, if "SELECT GREATEST(c1, c2) FROM t1" is executed for numeric type masking target c1 and c2, the masking type for c1 will be applied.

- When masking the data that includes multibyte characters, do not specify partial masking for masking type. The result may not be as expected.

### **6.1.3 Masking Condition**

Masking condition refers to the conditions configured to perform masking. Specify the masking condition in the expression parameter. Changed or actual data can be displayed for different users by defining masking condition. An expression that returns a boolean type result needs to be specified in masking condition and masking is performed only when TRUE is returned. Refer to "Value Expressions" in the PostgreSQL Documentation for information on the expressions that can be specified. Note that expressions that include a column cannot

be specified.

For example, when masking data only for "postgres" users, specify 'current\_user = "postgres" in the masking condition.



Specify '1=1' so the masking condition is always evaluated to be TRUE and masking is performed all the time.

### 6.1.4 Masking Format

Masking format is a combination of change method and displayed characters when the masking condition is met. Masking format varies depending on the masking type. The following describes the masking format.

......

#### Full masking

With full masking, all characters are changed to values as determined by the database. Changed characters can be referenced in the pgx\_confidential\_values table. Also, replacement characters can be changed using the pgx\_update\_confidential\_values system management function.



Refer to "6.3 Data Types for Masking" for information on the data types for which data masking can be performed.

#### Partial masking

With partial masking, data is changed according to the content in the function\_parameters parameter. The method of specifying function\_parameters varies depending on the data type.

Category	Method of specifying function_parameters
Numeric type	'replacementCharacter, startPosition, endPosition'
	- replacementCharacter. Specify the number to display. Specify a value from 0 to 9.
	- startPosition: Specify the start position of masking. Specify a positive integer.
	- endPosition. Specify the end position of masking. Specify a positive integer that is greater than startPosition.
	Example
	Specify as below to change the values from the 1st to 5th digits to 9.
	function_parameters := '9, 1, 5'
	In this example, if the original data is "123456789", it will be changed to "999996789".
Character type	'inputFormat, outputFormat, replacementCharacter, startPosition, endPosition'
	- <i>inputFormat</i> : Specify the current format of the data. Specify "V" for characters that will potentially be masked, and specify "F" for values such as spaces or hyphens that will not be masked.
	- <i>outputFormat</i> : Define the method to format the displayed data. Specify "V" for characters that will potentially be masked. Any character to be output can be specified for each character "F" in <i>inputFormat</i> . If you want to output a single quotation mark, specify two of them consecutively.
	- replacementCharacter. Specify any single character. If you want to output a single quotation mark, specify two of them consecutively.
	- <i>startPosition</i> : Specify the position of "V" as the start position of masking. For example, to specify the position of the 4th "V" from the left, specify 4. Specify a positive integer.

Category	Method of specifying function_parameters
	- <i>endPosition</i> . Specify the position of "V" as an end position of masking. When working out the end position, do not include positions of "F". For example, to specify the position of the 11th "V" from the left, specify 11. Specify a positive integer that is greater than <i>startPosition</i> .
	Example
	Specify as below to mask a telephone number other than the first three digits using *.
	function_parameters := 'VVVFVVVVVVVVVVVVVVVVVVVVVVVVVVVVVVVVV
	In this example, if the original data is "012-3156-7890", it will be changed to "012-***-*."
Date/timestamp type	'MDYHMS'
	- M: Masks month. To mask month, enter the month from 1 to 12 after a lowercase letter m. Specify an uppercase letter M to not mask month.
	- D: Masks date. To mask date, enter the date from 1 to 31 after a lowercase letter d. If a value bigger than the last day of the month is entered, the last day of the month will be displayed. Specify an uppercase letter D to not mask date.
	- Y: Masks year. To mask year, enter the year from 1 to 9999 after a lowercase letter y. Specify an uppercase letter Y to not mask year.
	- H: Masks hour. To mask hour, enter the hour from 0 to 23 after a lowercase letter h. Specify an uppercase letter H to not mask hour.
	- M: Masks minute. To mask minute, enter the minute from 0 to 59 after a lowercase letter m. Specify an uppercase letter M to not mask minute.
	- S: Masks second. To mask second, enter the second from 0 to 59 after a lowercase letter s. Specify an uppercase letter S to not mask second.
	Example
	Specify as below to mask hour, minute, and second and display 00:00:00.
	function_parameters := 'MDYh0m0s0'
	In this example, if the original data is "2010-10-10 10:10:10", it will be changed to "2010-10-10 00:00:00".



- Refer to "B.3.2 pgx\_create\_confidential\_policy" for information on function\_parameters.
- Refer to "6.3 Data Types for Masking" for information on the data types for which masking can be performed.

### Regular expression masking

With regular expression masking, data is changed according to the content of the regexp\_pattern, regexp\_replacement and regexp\_flags parameters. For regexp\_pattern, specify the search pattern using a regular expression. For regexp\_replacement, specify the replacement character to use when data matches the search pattern. For regexp\_flags, specify the regular expression flags.

.....



Specify as below to change all three characters starting from b to X.

regexp\_pattern := 'b..'

regexp\_replacement:= 'X'

regexp\_flags := 'g'

In this example, if the original data is "foobarbaz", it will be changed to "fooXX".



#### See

- Refer to "POSIX Regular Expressions" in the PostgreSQL Documentation and check pattern, replacement, and flags for information on the values that can be specified for regexp\_pattern, regexp\_replacement, and regexp\_flags.
- Refer to "6.3 Data Types for Masking" for information on the data types for which masking can be performed.



- When column data type is character(*n*) or char(*n*) and if the string length after change exceeds n, the extra characters will be truncated and only characters up to the nth character will be displayed.
- When column data type is character varying(*n*) or varchar(*n*) and if the string length after change exceeds the length before the change, the extra characters will be truncated and only characters up to the length before change will be displayed.

## 6.2 Usage Method

#### Preparation

The following preparation is required to use this feature.

- 1. Set the postgresql.conf file parameters.
  - Prepend "pgx\_datamasking" to the shared\_preload\_libraries parameter.
- 2. Restart the instance.
- 3. Execute CREATE EXTENSION for the database that will use this feature.

The target database is described as "postgres" here.

Use the psql command to connect to the "postgres" database.



#### Example

postgres=# CREATE EXTENSION pgx\_datamasking;
CREATE EXTENSION



You must always prepend "pgx\_datamasking" to the "shared\_preload\_libraries" parameter.



- Specify "false" for pgx\_datamasking.enable to not use this feature. Data will not be masked even if a masking policy is configured. This feature becomes available again once "true" is specified for pgx\_datamasking.enable. This setting can be made

by specifying a SET statement or specifying a parameter in the postgresql.conf file. Example

```
postgres=# SET pgx_datamasking.enable=false;
```

- Hereafter, also perform this preparatory task for the "template1" database, so that this feature can be used by default when creating a new database.

#### **Usage**

To perform masking, a masking policy needs to be configured. The masking policy can be created, changed, confirmed, enabled, disabled or deleted during operation.

The procedures to perform these tasks are explained below with examples.

- 1. Creating a masking policy
- 2. Changing a masking policy
- 3. Confirming a masking policy
- 4. Enabling and disabling a masking policy
- 5. Deleting a masking policy



Only database superusers can configure masking policies.

### 6.2.1 Creating a Masking Policy

An example of the operation on the server is shown below.

1. Create a masking policy

Execute the pgx\_create\_confidential\_policy system management function to create a masking policy.

The following values are configured in this example.

- Masking target: Numeric type c1
- Masking type: FULL
- Masking condition: '1=1'

```
postgres=# select pgx_create_confidential_policy(table_name := 't1', policy_name := 'p1',
expression := '1=1', column_name := 'c1', function_type := 'FULL');
pgx_create_confidential_policy
------t
(1 row)
```

2. Confirm the displayed data

Confirm that the masking target data (column c1) has been correctly changed.



- Refer to "B.3.2 pgx\_create\_confidential\_policy" for information on the pgx\_create\_confidential\_policy system management function.



- Only one masking policy can be created per table.
- All users can view the masking policy created, so do not grant the login privilege of the database where this feature is set to the users who refer to the changed data. Masking policies are defined in the "pgx\_confidential\_columns", "pgx\_confidential\_policies" and "pgx\_confidential\_values" tables.

### 6.2.2 Changing a Masking Policy

- 1. An example of the operation on the server is shown below.
- 2. Change a masking policy

Execute the pgx\_alter\_confidential\_policy system management function to change a masking policy.

The following values are changed in this example.

- Content of change: Add a masking target
- Masking target: Character type c2
- Masking type: PARTIAL
- Masking condition: 'VVVFVVVVFVVVV, VVV-VVVV-VVVV, \*, 4, 11'

3. Confirm the displayed data

Confirm that the masking target data has been correctly changed.

```
postgres=# select * from t1;
c1 | c2
----+----
0 | 012-***-***
0 | 012-***-***
0 | 012-***-***
(3 row)
```



- Refer to "B.3.1 pgx\_alter\_confidential\_policy" for information on the pgx\_alter\_confidential\_policy system management function.

# 6.2.3 Confirming a Masking Policy

An example of the operation on the server is shown below.

Confirm information about a masking target where a masking policy is set
 Refer to the pgx\_confidential\_columns table to confirm the masking target where the masking policy is set.

2. Confirm information about the masking policy content

Refer to pgx\_confidential\_policies to confirm the masking policy content.



- Refer to "D.1 pgx\_confidential\_columns" for information on the pgx\_confidential\_columns table.
- Refer to "D.2 pgx\_confidential\_policies" for information on the pgx\_confidential\_policies table.

### 6.2.4 Enabling and Disabling a Masking Policy

An example of the operation on the server is shown below.

1. Disable a masking policy

Execute the pgx\_enable\_confidential\_policy system management function to disable a masking policy.

```
postgres=# select pgx_enable_confidential_policy(table_name := 't1', policy_name := 'p1',
enable := 'f');
pgx_enable_confidential_policy
-----t
(1 row)
```

2. Confirm the displayed data

Confirm that the original data is displayed by disabling the masking policy.

3. Enable a masking policy

Execute the pgx\_enable\_confidential\_policy system management function to enable a masking policy.

```
postgres=# select pgx_enable_confidential_policy(table_name := 't1', policy_name := 'p1',
enable := 't');
pgx_enable_confidential_policy
-----t
(1 row)
```

4. Confirm the displayed data

Confirm that the masking target data has been correctly changed.

```
postgres=# select * from t1;
c1 | c2
---+----
0 | 012-***-***
0 | 012-***-***
```

```
0 | 012-***-***
(3 row)
```



- Refer to "B.3.4 pgx\_enable\_confidential\_policy" for information on the pgx\_enable\_confidential\_policy system management function.

### 6.2.5 Deleting a Masking Policy

An example of the operation on the server is shown below.

1. Delete a masking policy

Execute the pgx\_drop\_confidential\_policy system management function to delete a masking policy.

```
postgres=# select pgx_drop_confidential_policy(table_name := 't1', policy_name := 'p1');
pgx_drop_confidential_policy
-----t
t
(1 row)
```

2. Confirm the displayed data

Confirm that the original data is displayed by deleting the masking policy.



- Refer to "B.3.3 pgx\_drop\_confidential\_policy" for information on the gx\_drop\_confidential\_policy function.

# 6.3 Data Types for Masking

The data types for which data masking can be performed are shown below.

Category	Data type	Masking type		
		Full masking	Partial masking	Regular expression masking
Numeric type	smallint	Y	Y	N
	integer	Y	Y	N
	bigint	Y	Y	N
	decimal	Y	Y	N
	numeric	Y	Y	N
	float	Y	Y	N
	real	Y	Y	N
	double precision	Y	Y	N

Category	Category Data type Masking type		Masking type	e	
		Full masking	Partial masking	Regular expression masking	
Character type	character varying(n)	Y	Y	Y	
	varchar(n)	Y	Y	Y	
	character(n)	Y	Y	Y	
	char(n)	Y	Y	Y	
Date/timestamp type	date	Y	Y	N	
	timestamp	Y	Y	N	

# **6.4 Security Notes**

- Starting with FEP 10, logical replication is available, which allows non-backed up clusters to subscribe to databases where data masking policies are enabled. Logical replication allows publisher and subscriber databases to have their own or the same data masking policies.
  - In this scenario, the user must disable data masking on the publisher database whenever a subscription is created. This ensures that subscribers are able to obtain the original data (initial copy) instead of the masked version. Then, it is the user's responsibility to set masking policies to each subscribed database.
- Take strong caution in publishing data masking's confidential tables (pgx\_confidential\_policies, pgx\_confidential\_columns, etc.) unless the user is publishing all tables of the database and wants to apply the same data masking's policies on the subscribed database for all of them.

Otherwise, as these confidential tables contain the masking policies for all tables of the database, confidential policies of unpublished tables may be unintentionally published. Additionally, it is not possible to apply different data masking policies on the subscriber database.

# Chapter 7 Periodic Operations

This chapter describes the operations that must be performed periodically when running daily database jobs.

### 7.1 Configuring and Monitoring the Log

FUJITSU Enterprise Postgres enables you to output database errors and warnings to a log file.

This information is useful for identifying if errors have occurred and the causes of those errors.

By default, this information is output to the system log. It is recommended that you configure FUJITSU Enterprise Postgres to collect logs from its log files (for example, log\_destination) before operating FUJITSU Enterprise Postgres.

Periodically monitor the log files to check if any errors have occurred.



- Refer to "Error Reporting and Logging" under "Server Administration" in the PostgreSQL Documentation for information on logs.

......

- Refer to "Configuring Parameters" in the Installation and Setup Guide for Server for information on log settings when operating with WebAdmin.

### 7.2 Monitoring Disk Usage and Securing Free Space

When a database is used for an extended period, free space on the disk is continuously consumed and in some cases the disk space runs out. When this happens, database jobs may stop and no longer run.

You should, therefore, periodically monitor the usage of disk space, and delete obsolete files located in the disk.

Monitor the disk usage of the disk where the following directories are located:

- Data storage destination directory
- Transaction log storage destination (if the transaction log is stored in a different directory from the data storage destination directory)
- Backup data storage destination directory
- Tablespace storage destination directory

### 7.2.1 Monitoring Disk Usage

To check the disk usage, use the following operating system commands:

- df command

You can even use SQL statements to check tables and indexes individually.

Refer to "Determining Disk Usage" under "Server Administration" in the PostgreSQL Documentation for information on this method.



If you are using WebAdmin for operations, a warning is displayed when disk usage reaches 80%

### 7.2.2 Securing Free Disk Space

Secure free disk space by using the following operating system commands to delete unnecessary files, other than the database, from the same disk unit.

- rm command

You can also secure disk space by performing the following tasks periodically:

- To secure space on the data storage destination disk:
  - Execute the REINDEX statement. Refer to "7.5 Reorganizing Indexes" for details.
- To secure space on the backup data storage destination disk:
  - Execute backup using WebAdmin or the pgx\_dmpall command.

# 7.3 Automatically Closing Connections

If an application stops responding and abnormally terminates for any reason, the connection from the application may remain active on the database server. If this situation continues for an extended period, other applications attempting to connect to the database server may encounter an error, or an error indicating that the tables are unavailable may occur.

It is, therefore, recommended that idle connections be closed automatically at regular intervals.

Set the following parameters in the postgresql.conf file to indicate the time permitted to elapse before a connection is closed.

Parameter name	Setting	Description
tcp_keepalives_idle	Time until keepalive is sent (seconds)  If 0, the default value of the system is used.	Sends keepalive to an idle connection at the specified interval in seconds  It is recommended to specify 30 seconds.
tcp_keepalives_interval	keepalive send interval (seconds)  If 0, the default value of the system is used.	Sends keepalive at the specified interval  It is recommended to specify 10 seconds.
tcp_user_timeout	Time to wait for a response from the server (milliseconds)  If 0, the default value of the system is used.  If not set, the behavior is the same as if 0 were specified.	After establishing the connection, when sending from the client to the server, if the TCP resend process operates, specify the time until it is considered to be disconnected.  If a value other than 0 is specified in this parameter, the time until automatic disconnection is determined by the waiting time specified in this parameter. The actual wait time is until the timing of the first keepalive retransmission after the time specified by this parameter has elapsed.



If a value other than 0 is specified for the tcp\_user\_timeout parameter, the waiting time set by the tcp\_keepalives\_idle parameter and tcp\_keepalives\_interval parameter will be invalid and the waiting time specified by the tcp\_user\_timeout parameter will be used.



Refer to "Connection Settings" under "Server Administration" in the PostgreSQL Documentation for information on the parameters.

# 7.4 Monitoring the Connection State of an Application

FUJITSU Enterprise Postgres does not immediately delete the updated or deleted data. If the VACUUM determines there are no transactions that reference the database, FUJITSU Enterprise Postgres collects obsolete data.

However, obsolete data is not collected if there are connections that have remained active for an extended period or connections occupying resources. In this case the database may expand, causing performance degradation.



Refer to "Routine Vacuuming" under "Server Administration" in the PostgreSQL Documentation for information on the VACUUM command.

In such cases, you can minimize performance degradation of the database by monitoring problematic connections.

The following methods are supported for monitoring connections that have been in the waiting status for an extended period:

- 7.4.1 Using the View (pg\_stat\_activity)
- 7.4.2 Using pgAdmin

### 7.4.1 Using the View (pg\_stat\_activity)

Use the view (pg\_stat\_activity) to identify and monitor connections where the client has been in the waiting status for an extended period.

......



The example below shows connections where the client has been in the waiting status for at least 60 minutes.

However, when considering continued compatibility of applications, do not reference system catalogs directly in the following SQL statements.

```
postgres=# select * from pg_stat_activity where state='idle in transaction' and current_timestamp >
cast(query_start + interval '60 minutes' as timestamp);
-[ RECORD 1 ]----+
             | 13003
datid
datname
             db01
             4638
pid
             | 10
usesysid
             | fsep
usename
application_name | apl01
client_addr | 192.33.44.15
client_hostname
             27500
client_port
backend_start | 2018-02-24 09:09:21.730641+09
          2018-02-24 09:09:23.858727+09
xact start
            | 2018-02-24 09:09:23.858727+09
query_start
state_change
            2018-02-24 09:09:23.858834+09
wait_event_type | Client
wait_event
            ClientRead
             | idle in transaction
state
backend xid
backend_xmin
query
             | begin;
backend type
           | client backend
.....
```



- Refer to "Notes on Application Compatibility" in the Application Development Guide for information on maintaining application compatibility.

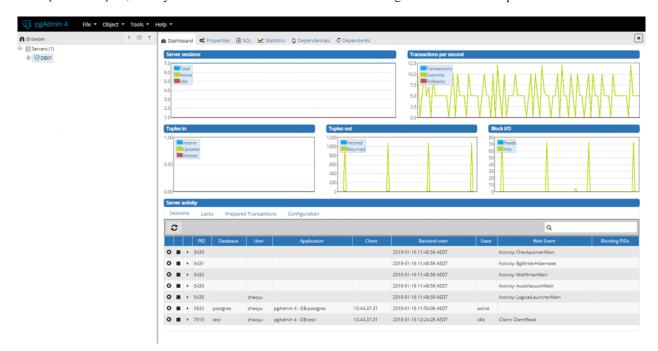
......

- Refer to "The Statistics Collector" under "Server Administration" in the PostgreSQL Documentation for information on pg\_stat\_activity.

# 7.4.2 Using pgAdmin

This section describes the procedure for monitoring connections using [Server Status] in pgAdmin.

- 1. In the [Browser] pane, click the database server for monitoring.
- 2. In the [Dashboard] tab, identify client connections that have been in the waiting state for an extended period.



# 7.5 Reorganizing Indexes

Normally, a database defines indexes in tables, but if data is frequently updated, indexes can no longer use free space in the disk efficiently. This situation can also cause a gradual decline in database access performance.

To rearrange used space on the disk and prevent the database access performance from declining, it is recommended that you periodically execute the REINDEX command to reorganize indexes.

Check the disk usage of the data storage destination using the method described in "7.2 Monitoring Disk Usage and Securing Free Space".



Because the REINDEX command retrieves the exclusive lock for an index being processed and locks writing of tables that are the source of the index, other processes that access these may stop while waiting to be locked.

Therefore, it is necessary to consider measures such as executing the command after the task is completed.



Refer to "Routine Reindexing" under "Server Administration" in the PostgreSQL Documentation for information on reorganizing indexes by periodically executing the REINDEX command.



Typically, reorganize indexes once a month at a suitable time such as when conducting database maintenance. Use SQL statements to check index usage. If this usage is increasing on a daily basis, adjust the frequency of recreating the index as compared to the free disk space.

The following example shows the SQL statements and the output.

However, when considering continued compatibility of applications, do not reference system catalogs and functions directly in the following SQL statements. Refer to "Notes on Application Compatibility" in the Application Development Guide for details.

#### [SQL statements]

```
SELECT
  nspname AS schema_name,
  relname AS index_name,
  round(100 * pg_relation_size(indexrelid) / pg_relation_size(indrelid)) / 100 AS index_ratio,
  pg_size_pretty(pg_relation_size(indexrelid)) AS index_size,
  pg_size_pretty(pg_relation_size(indrelid)) AS table_size
FROM pg_index I
  LEFT JOIN pg_class C ON (C.oid = I.indexrelid)
  LEFT JOIN pg_namespace N ON (N.oid = C.relnamespace)
WHERE
  C.relkind = 'i' AND
  pg_relation_size(indrelid) > 0
ORDER BY pg_relation_size(indexrelid) DESC, index_ratio DESC;
```

#### [Output]

schema_name	index_name	index_ratio	_ '	_
public pg_catalog	pgbench_accounts_pkey   pg_depend_depender_index   pg_depend_reference_index	0.16	2208 KB   224 KB   216 KB	13 MB 368 KB 368 KB



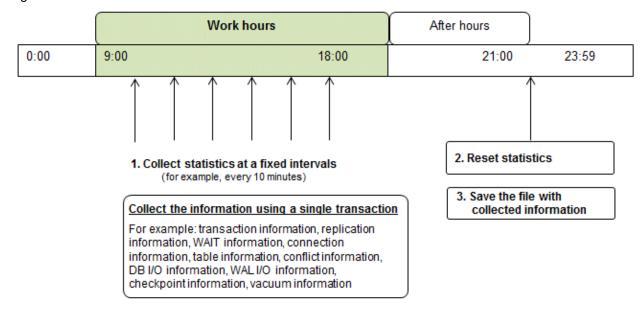
Refer to "Notes on Application Compatibility" in the Application Development Guide for information on maintaining application compatibility.

# 7.6 Monitoring Database Activity

FUJITSU Enterprise Postgres enables you to collect information related to database activity. By monitoring this information, you can check changes in the database status.

This information includes wait information for resources such as internal locks, and is useful for detecting performance bottlenecks. Furthermore, you should collect this information in case you need to request Fujitsu technical support for an investigation.

Figure 7.1 Overview of information collection



1. Collect statistics at fixed intervals during work hours.

Accumulate the collected information into a file.

Wherever possible, collect data from the various statistics views using a single transaction, because it enables you to take a snapshot of system performance at a given moment.

Refer to "7.6.1 Information that can be Collected" for information on the system views that can be collected.

2. Reset statistics after work hours, that is, after jobs have finished.

Refer to "7.6.3 Information Reset" for information on how to reset statistics.

3. Save the file with collected information.

Keep the file with collected information for at least two days, in order to check daily changes in performance and to ensure that the information is not deleted until you have sent a query to Fujitsu technical support.

Where jobs run 24 hours a day, reset statistics and save the file with collected information when the workload is low, for example, at night.



Statistics cumulatively add the daily database value, so if you do not reset them, the values will exceed the upper limit, and therefore will not provide accurate information.

The subsections below explain the following:

- Information that can be collected
- Collection configuration
- Information reset

### 7.6.1 Information that can be Collected

Information that can be collected is categorized into the following types:

- Information common to PostgreSQL
- Information added by FUJITSU Enterprise Postgres

### Information common to PostgreSQL



Refer to "Monitoring Database Activity" under "Server Administration" in the PostgreSQL Documentation for information common to PostgreSQL.

#### Information added by FUJITSU Enterprise Postgres

You can collect the following information added by FUJITSU Enterprise Postgres.

Table 7.1 Information added by FUJITSU Enterprise Postgres

View name	Description
pgx_stat_lwlock	Displays statistic related to lightweight lock, with each type of content displayed on a separate line. This information helps to detect bottlenecks.  Refer to "C.2 pgx_stat_lwlock" for details.
pgx_stat_latch	Displays statistics related latches, with each type of wait information within FUJITSU Enterprise Postgres displayed on a separate line. This information helps to detect bottlenecks.  Refer to "C.3 pgx_stat_latch" for details.

View name	Description
pgx_stat_walwriter	Displays statistics related to WAL writing, in a single line.
	Refer to "C.4 pgx_stat_walwriter" for details.
pgx_stat_sql	Displays statistics related to SQL statement executions, with each type of SQL statement displayed on a separate line.
	Refer to "C.5 pgx_stat_sql" for details.
pgx_stat_gmc	Displays statistics related to Global Meta Cache hit ration and used memory size.
	Refer to "C.6 pgx_stat_gmc" for detail. Also refer to Chapter 12 Global Meta Cache" for information on the Global Meta Cache.

## 7.6.2 Collection Configuration

The procedure for configuring collection depends on the information content.

- Information common to PostgreSQL
- Information added by FUJITSU Enterprise Postgres

#### Information common to PostgreSQL



Refer to "The Statistics Collector" in "Monitoring Database Activity" under "Server Administration" in the PostgreSQL Documentation for information on information common to PostgreSQL.

......

#### Information added by FUJITSU Enterprise Postgres

Information added by FUJITSU Enterprise Postgres is collected by default.

To enable or disable information collection, change the configuration parameters in postgresql.conf. The following table lists the views for which you can enable or disable information collection, and the configuration parameters.

View name	Parameter
pgx_stat_lwlock	track_waits (*1)
pgx_stat_latch	
pgx_stat_sql	track_sql
pgx_stat_gmc	track_gmc

Remarks: You cannot change the collection status for  $pgx\_stat\_walwriter$ .

\*1: When executing the SQL statement with EXPLAIN ANALYZE, processing time may increase because of this information collection. It is recommended to set this parameter to "off" when executing EXPLAIN ANALYZE to check the processing time.

Refer to "Appendix A Parameters" for information on the parameters.

### 7.6.3 Information Reset

This section describes how to reset information.

#### Information added by FUJITSU Enterprise Postgres

You can reset information added by FUJITSU Enterprise Postgres by using the pg\_stat\_reset\_shared function in the same way as for information common to PostgreSQL.

Configure the following parameters in the pg\_stat\_reset\_shared function:

Function	Type of return value	Description
pg_stat_reset_shared(text)	void	Reset some cluster-wide statistics counters to zero, depending on the argument (requires superuser privileges).
		Calling pg_stat_reset_shared('lwlock') will zero all counters shown in pgx_stat_lwlock.
		Similarly, in the following cases, all values of the pertinent statistics counter are reset:
		- If pg_stat_reset_shared('latch') is called:
		All values displayed in pgx_stat_latch
		- If pg_stat_reset_shared('walwriter') is called:
		All values displayed in pgx_stat_walwriter
		- If pg_stat_reset_shared('sql') is called:
		All values displayed in pgx_stat_sql
		- If pg_stat_reset_shared('gmc') is called:
		All values except size column in pgx_stat_gmc



Refer to "Statistics Functions" in "Monitoring Database Activity" under "Server Administration" in the PostgreSQL Documentation for information on other parameters of the pg\_stat\_reset\_shared function.

......

# Chapter 8 Streaming Replication Using WebAdmin

This chapter describes how to create a streaming replication cluster using WebAdmin.

Streaming replication allows the creation of one or more standby instances, which connect to the master instances and replicate the data using WAL records. The standby instance can be used for read-only operations.

WebAdmin can be used to create a streaming replication cluster. WebAdmin allows the creation of a cluster in the following configurations:

- Master-Standby Configuration: This configuration creates a master and standby instance together.
- Standby Only Configuration: This configuration creates a standby instance from an already existing instance.

# Point

- A standby instance can be created from a standalone instance, a master instance, or even from another standby instance.
- If a streaming replication cluster is created using WebAdmin, the network with the host name (or IP address) specified in [Host name] will be used across sessions of WebAdmin, and also used as the log transfer network.

- To use a network other than the job network as the log transfer network, specify the host name other than the job network one in [Host name].

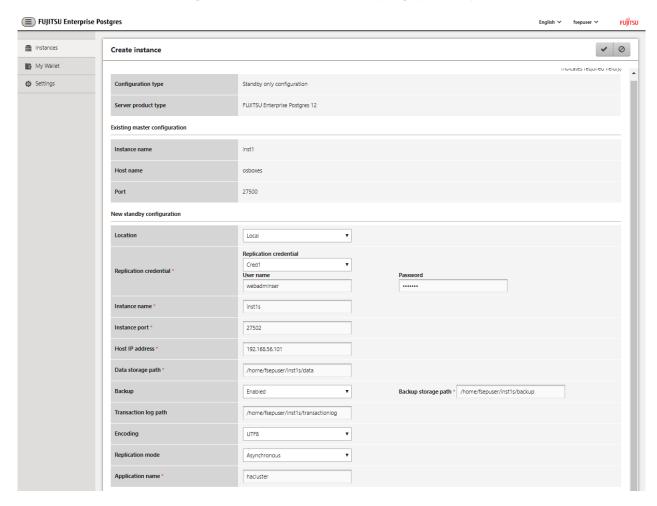
## 8.1 Creating a Standby Instance

Follow the procedure below to create a standby instance.

- 1. In the [Instances] tab, select the instance from which a standby instance is to be created.
- 2. Click

3. Enter the information for the standby instance to be created. In the example below, a standby instance is created from instance "inst1".

The instance name, host address and port of the selected instance are already displayed for easy reference.



#### Enter the following items:

- [Location]: Whether to create the instance in the server that the current user is logged in to, or in a remote server. The default is "Local", which will create the instance in the server machine where WebAdmin is currently running.
- [Replication credential]: The user name and password required for the standby instance to connect to the master instance. The user name and password can be entered or selected from the Wallet. Refer to "Appendix G WebAdmin Wallet" for information on creating wallet entries.
- [Instance name]: Name of the standby database instance to create.

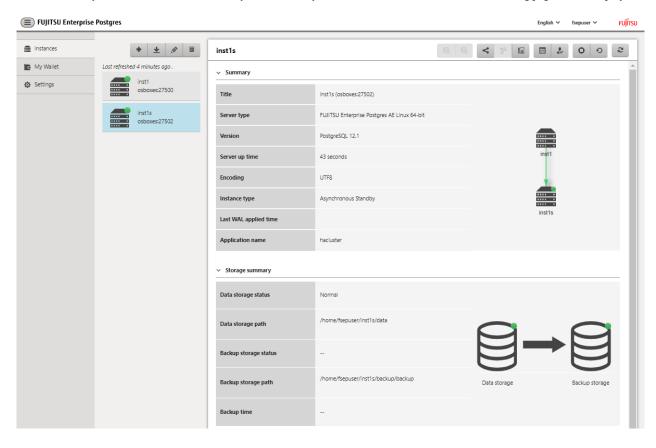
The name must meet the conditions below:

- Maximum of 16 characters
- The first character must be an ASCII alphabetic character
- The other characters must be ASCII alphanumeric characters
- [Instance port]: Port number of the standby database instance.
- [Host IP address]: The IP address of the server machine where the standby instance is to be created. This information is needed to configure the standby instance to be connected to the master.
- [Data storage path]: Directory where the database data will be stored
- [Backup storage path]: Directory where the database backup will be stored
- [Transaction log path]: Directory where the transaction log will be stored

- [Encoding]: Database encoding system
- [Replication mode]: Replication mode of the standby instance to be created ("Asynchronous" or "Synchronous")
- [Application name]: The reference name of the standby instance used to identify it to the master instance.

The name must meet the conditions below:

- Maximum of 16 characters
- The first character must be an ASCII alphabetic character
- The other characters must be ASCII alphanumeric characters
- 4. Click to create the standby instance.
- 5. Once the standby instance is created successfully, select standby instance in the [Instances] tab. The following page will be displayed:





- Backups are not possible for standby instances in WebAdmin. As a result, and are disabled and no value is shown for [Backup storage status] and [Backup time].
- If using WebAdmin to manage Mirroring Controller, the message below may be output to the server log or system log in the standby instance. No action is required, as the instance is running normally.

```
ERROR: pgx_rcvall failed (16491)
ERROR: pgx_rcvall: backup of the database has not yet been performed, or an incorrect backup storage directory was specified
```

- Replication credential (user name and password) should not contain hazardous characters. Refer to "Appendix H WebAdmin Disallow User Inputs Containing Hazardous Characters".

### 8.2 Promoting a Standby Instance

Streaming replication between a master and standby instance can be discontinued using WebAdmin.

Follow the procedure below to promote a standby instance to a standalone instance, thereby discontinuing the streaming replication.

- 1. In the [Instances] tab, select the standby instance that needs to be promoted.
- 2. Click Ta.
- 3. Click [Yes] from the confirmation dialog box.

The standby instance will be promoted and will become a standalone instance, which is not part of a streaming replication cluster.

Once the standby instance is promoted to become a standalone instance, the backup storage status will be "Error". This is because no backups are available when the instance is newly promoted to a standalone instance. The status will be reset if a new backup is performed by clicking [Solution] or ...

# 8.3 Converting an Asynchronous Replication to Synchronous

Streaming replication between a master and standby instance can be configured to be in Asynchronous or Synchronous mode. This mode can be changed even after the standby instance was successfully created.

Follow the procedure below to convert an Asynchronous standby instance to Synchronous.

- 1. In the [Instances] tab, select the master instance of the relevant cluster.
- 2. Click
- 3. In the [Streaming replication] section, edit the value for [Synchronous standby names].
  - Add the "Application name" of the standby instance you want to be in Synchronous mode.
- 4. Click .
- 5. Select the master instance and click
- 6. Select the standby instance. [Instance type] will now show the updated status.



- Converting an Asynchronous standby instance to Synchronous can cause the master instance to queue the incoming transactions until
  the standby instance is ready. For this reason, it is recommended that this operation be performed during a scheduled maintenance
  period.
- When adding a synchronous standby instance, FUJITSU Enterprise Postgres will only keep the first entry in [Synchronous standby names] in synchronous state.
- To learn more about the differences between synchronous and asynchronous standby modes and their behavior, refer to "Streaming Replication" in "High Availability, Load Balancing, and Replication" in the PostgreSQL Documentation.

# 8.4 Converting a Synchronous Replication to Asynchronous

Streaming replication between a master and standby instance can be configured to be in Asynchronous or Synchronous Mode. This mode can be changed even after the standby instance was successfully created.

Follow the procedure below to convert a Synchronous standby instance to Asynchronous.

- 1. In the [Instances] tab, select the master instance of the relevant cluster.
- 2. Click

- 3. In the [Streaming replication] section, edit the value for [Synchronous standby names].
  - Remove the "Application name" of the standby instance you want to be in Asynchronous mode.
- 4. Click .
- 5. Select the master instance and click .
- 6. Select the standby instance. [Instance type] will now show the updated status.



To learn more about the differences between synchronous and asynchronous standby modes and their behavior, refer to "Streaming Replication" in "High Availability, Load Balancing, and Replication" in the PostgreSQL Documentation.

# 8.5 Joining a Replication Cluster

WebAdmin facilitates the joining of an old master of the cluster as a standby node.

- 1. In the [Instances] tab, select the remote instance (from where the new cluster node will stream WAL entries), and then click 💂

- 2. Configure the node to accept streaming requests from the new node.
- 3. In the [Instances] tab, select the new standby instance (which needs to be connected to the cluster), and then click 🌠.



- 4. Set [Replication host name] to the remote instance.
- 5. Enter [Replication credential].

Specify the user name and password required for the standby instance to connect to the remote instance. The user name and password can be entered or selected from the Wallet. Refer to "Appendix G WebAdmin Wallet" for information on creating wallet entries. Replication credential (user name and password) should not contain hazardous characters. Refer to "Appendix H WebAdmin Disallow User Inputs Containing Hazardous Characters".

6. Enter [Host IP address].

Specify the IP address of the node where the standby instance was created.

- 7. Click v to open the [Join replication cluster] dialog box.
- 8. Select [Restart later] or [Restart now], and then click [Yes] to set up the standby instance.
- 9. Upon successful completion, the confirmation dialog box will be displayed.
- 10. Click [Close] to return to the instance details window.

The instance will become a standby instance, and will be part of the streaming replication cluster. The replication diagram will display the relationship between the standby instance and the remote instance. The user can change the replication relationship of the remote instance from asynchronous to synchronous (and vice versa) using the [Configuration] window.

# Chapter 9 Installing and Operating the In-memory Feature

The in-memory feature enables fast aggregation using Vertical Clustered Index (VCI) and memory-resident feature.

VCI has a data structure suitable for aggregation, and features parallel scan and disk compression, which enable faster aggregation through reduced disk I/O.

The memory-resident feature reduces disk I/O that occurs during aggregation. It consists of the preload feature that reads VCI data to memory in advance, and the stable buffer feature that suppresses VCI data eviction from memory. The stable buffer feature secures the proportion specified by parameter in the shared memory for VCI.

This chapter describes how to install and operate the in-memory feature.



This feature can only be used in Advanced Edition.

# 9.1 Installing Vertical Clustered Index (VCI)

This section describes the installation of VCI.

- 1. Evaluating whether to Install VCI
- 2. Estimating Resources
- 3. Setting up

### 9.1.1 Evaluating whether to Install VCI

VCI uses available resources within the server to increase scan performance.

It speeds up processing in many situations, and can be more effective in the following situations:

- Single table processing
- Processing that handles many rows in the table
- Processing that handles some columns in the table
- Processing that performs very heavy aggregation such as simultaneous sum and average aggregation

VCI will not be used in the following cases, so it is necessary to determine its effectiveness in advance:

- The data type of the target table or column contains VCI restrictions.
- The SQL statement does not meet the VCI operating conditions
- VCI is determined to be slower based on cost estimation



If performing operations that use VCI, the full\_page\_writes parameter setting in postgresql.conf must be enabled (on). For this reason, if this parameter is disabled (off), operations that use VCI return an error. In addition, to perform operations for tables that do not create a VCI when the full\_page\_writes parameter setting is temporarily disabled (off), do not create a VCI or perform operations to tables that created a VCI during that time.

......



- Refer to "9.1.4 Data that can Use VCI" for information on VCI restrictions.

- Refer to "Scan Using a Vertical Clustered Index (VCI)" - "Operating Conditions" in the Application Development Guide for information on VCI operating conditions.

### 9.1.2 Estimating Resources

Estimate resources before setting up VCI.

Select the aggregation that you want to speed up and identify the required column data. The additional resources below are required according to the number of columns.

- Memory

Secure additional capacity required for the disk space for the column for which VCI is to be created.

- Disk

Secure additional disks based on the disk space required for the column for which VCI is to be created, as VCI stores column data as well as existing table data on the disk. It is recommended to provide a separate disk in addition to the existing one, and specify it as the tablespace to avoid impact on any other jobs caused by I/O.



The operations on VCI can continue even if the memory configured for VCI is insufficient by using VCI data on the disk.



Refer to "Estimating Memory Requirements" and "Estimating Database Disk Space Requirements" in the Installation and Setup Guide for Server for information on how to estimate required memory and disk space.

### 9.1.3 Setting up

This section describes how to set up VCI.

#### Setup flow

- 1. Setting Parameters
- 2. Installing the Extensions
- 3. Creating VCI
- 4. Confirming that VCI has been Created

### 9.1.3.1 Setting Parameters

Edit postgresql.conf to set the required parameters for VCI. After that, start or restart the instance.

The following table lists the parameters that need or are recommended to be configured in advance:

Parameter name	Setting value	Description	Required
shared_preload_libraries	Literal 'vci, pg_prewarm'	VCI and shared library to be preloaded at server start.	Y
session_preload_libraries	Literal 'vci, pg_prewarm'	VCI and shared library to be preloaded at connection start.	Y
reserve_buffer_ratio	Percentage of shared memory to be used for stable buffer table	Proportion of shared memory to be used for a stable buffer table.	N

Parameter name	Setting value	Description	Required
vci.control_max_workers	Number of background workers that manage VCI	Number of background workers that manage VCI.  Add this value to max_worker_processes.	N
vci.max_parallel_degree	Maximum number of background workers used for parallel scan	Maximum number of background workers used for parallel scan.  Add this value to max_worker_processes.	N



```
shared_preload_libraries = 'vci, pg_prewarm'
session_preload_libraries = 'vci, pg_prewarm'
reserve_buffer_ratio = 20
vci.control_max_workers = 8
vci.max_parallel_degree = 4
max_worker_processes = 18 # Example: If the initial value was 6, 6 + 8 + 4 = 18
```



An error occurs if you use VCI to start instances when procfs is not mounted. Ensure that procfs is mounted before starting instances.



- Refer to "Appendix A Parameters" for information on all parameters for VCI. Refer also to default value for each parameter and details such as specification range in the same chapter. Refer to "Server Configuration" under "Server Administration" in the PostgreSQL documentation for information on shared\_preload\_libraries, session\_preload\_libraries, and max\_worker\_processes.

### 9.1.3.2 Installing the Extensions

Execute CREATE EXTENSION to install the VCI and pg\_prewarm extensions. Both extensions need to be installed for each database.

- Installing VCI

```
db01=# CREATE EXTENSION vci;
```

- Installing pg\_prewarm

db01=# CREATE EXTENSION pg\_prewarm;



- Only superusers can install VCI extensions.
- VCI extensions can only be installed in public schema.
- Some operations cannot be performed for VCI extensions. Refer to "9.2.1 Commands that cannot be Used for VCI" for details.

### 9.1.3.3 Creating a VCI

Execute the CREATE INDEX statement with the "USING vci" clause to create a VCI for the desired columns and the "WITH (stable\_buffer=true)" clause to enable the stable buffer feature.

To use a separate disk for the VCI, specify the TABLESPACE clause.

db01=# CREATE INDEX idx\_vci ON table01 USING vci (col01, col02) WITH (stable\_buffer=true);



- Some table types cannot be specified on the ON clause of CREATE INDEX. Refer to "9.1.4.1 Relation Types" for details.
- Some data types cannot be specified on the column specification of CREATE INDEX. Refer to "9.1.4.2 Data Types" for details.

- Some operations cannot be performed for VCI. Refer to "9.2.1 Commands that cannot be Used for VCI" for details.
- The same column cannot be specified more than once on the column specification of CREATE INDEX.
- VCI cannot be created for table columns that belong to the template database.
- CREATE INDEX creates multiple views named vci\_10digitRelOid\_5digitRelAttr\_1charRelType alongside VCI itself. These are called VCI internal relations. Do not update or delete them as they are used for VCI aggregation.
- All data for the specified column will be replaced in columnar format when VCI is created, so executing CREATE INDEX on an existing table with data inserted takes more time compared with a general index (B-tree). Jobs can continue while CREATE INDEX is running.
- When CREATE INDEX USING VCI is invoked on a partitioned table, the default behavior is to recurse to all partitions to ensure they all have matching indexes. Each partition is first checked to determine whether an equivalent index already exists, and if so, that index will become attached as a partition index to the index being created, which will become its parent index. If no matching index exists, a new index will be created and automatically attached; the name of the new index in each partition will be determined as if no index name had been specified in the command. If the ONLY option is specified, no recursion is done, and the index is marked invalid. (ALTER INDEX ... ATTACH PARTITION marks the index valid, once all partitions acquire matching indexes.) Note, however, that any partition that is created in the future using CREATE TABLE ... PARTITION OF will automatically have a matching index, regardless of whether ONLY is specified.
- Parallel index build is not supported on VCI indexes.

### 9.1.3.4 Confirming that the VCI has been Created

Execute the SELECT statement to reference the pg\_indexes catalog, and confirm that the VCI was created for the target columns.



### 9.1.4 Data that can Use VCI

This section describes on which relation types and for which data types VCIs can be created.

### 9.1.4.1 Relation Types

VCIs cannot be created on some relation types.

The ON clause of CREATE INDEX described in "9.1.3.3 Creating a VCI" cannot specify relations on which VCIs cannot be created.

- Relations on which VCIs can be created
  - Normal tables
  - UNLOGGED TABLES

- Relations on which VCIs cannot be created
  - Materialized views
  - Temporary tables
  - Views
  - Temporary views
  - Foreign tables

# 9.1.4.2 Data Types

VCIs cannot be created for some data types.

The column specification of CREATE INDEX described in "9.1.3.3 Creating a VCI" cannot specify a column with data type on which VCIs cannot be created.

Category	Data type	Supported by VCI?
Numeric	smallint	Y
	integer	Y
	bigint	Y
	decimal	Y
	numeric	Y
	real	Y
	double precision	Y
	serial	Y
	bigserial	Y
Monetary	money	Y
Character	varchar(n)	Y
	char(n)	Y
	nchar	Y
	nvarchar(n)	Y
	text	Y
Binary	bytea	Y
Date/time	timestamp	Y
	timestamp with time zone	Y
	date	Y
	time	Y
	time with time zone	Y
	interval	Y
Boolean	boolean	Y
Geometric	point	N
	line	N
	lseg	N
	box	N
	path	N

Category	Data type	Supported by VCI?
	polygon	N
	circle	N
Network address	cidr	N
	inet	N
	macaddr	N
	macaddr8	N
Bit string	bit(n)	Y
	bit varying(n)	Y
Text search	tsvector	N
	tsquery	N
UUID	uuid	Y
XML	xml	N
JSON	json	N
	jbson	N
Range	int4range	N
	int8range	N
	numrange	N
	tsrange	N
	tstzrange	N
	daterange	N
Object identifier	oid	N
	regproc	N
	regprocedure	N
	regoper	N
	regoperator	N
	regclass	N
	regtype	N
	regconfig	N
	regdictionary	N
pg_lsn type	pg_lsn	N
Array type	-	N
User-defined type	-	N
(Basic type, enumerated type, composite type, and range type)		

# 9.2 Operating VCI

This section describes how to operate VCI.

## 9.2.1 Commands that cannot be Used for VCI

Some operations cannot be performed for VCI extensions and VCI itself.

This section describes SQL commands that cannot be executed for the VCI extensions and VCI itself, and client application commands.

### SQL commands

- Operations that cannot be performed for the VCI extension

Command	Subcommand	Description
ALTER EXTENSION	UPDATE	The VCI extension cannot be specified.
	SET SCHEMA	This operation is not required for VCI.
	ADD	
	DROP	
CREATE EXTENSION	SCHEMA	The subcommands on the left cannot be performed if the VCI extension is specified.
		This operation is not required for VCI.

- Operations that cannot be performed on relations containing a VCI

Command	Subcommand	Description	
ALTER INDEX	SET	The subcommands on the left cannot be performed	
	SET TABLESPACE	if a VCI is specified.	
	ALL IN TABLESPACE	If the operation is required, delete the VCI using DROP INDEX, and re-create it using CREATE INDEX after completing the operation.	
ALTER OPERATOR	RENAME TO	The subcommands on the left cannot be performed	
CLASS	OWNER TO	if a VCI is specified.	
	SET SCHEMA	This operation is not supported in VCI.	
ALTER OPERATOR	ADD		
FAMILY	DROP		
	RENAME TO		
	OWNER TO		
	SET SCHEMA		
ALTER TABLE	ALL IN TABLESPACE name [ OWNED BY roleName ] SET	A tablespace that contains a VCI cannot be specified.	
	TABLESPACE newTablespace	If the operation is required, delete the VCI using DROP INDEX, and re-create it using CREATE INDEX after completing the operation.	
	DROP [ COLUMN ] [ IF EXISTS ] colName [ RESTRICT   CASCADE ]	A column that contains a VCI cannot be specified.  If the operation is required, delete the VCI using	
	ALTER [ COLUMN ] colName [ SET DATA ] TYPE dataType [ COLLATE collation ] [ USING expr ]	DROP INDEX, and re-create it using CREATE INDEX after completing the operation.	
	CLUSTER ON indexName	A VCI cannot be specified.	
	REPLICA IDENTITY {DEFAULT   USING INDEX indexName   FULL   NOTHING}	This operation is not supported in VCI.	

Command	Subcommand	Description
CLUSTER	-	A table that contains a VCI and VCI cannot be specified.
		If the operation is required, delete the VCI using DROP INDEX, and re-create it using CREATE INDEX after completing the operation.
CREATE INDEX	UNIQUE	The subcommands on the left cannot be performed
	CONCURRENTLY	if a VCI is specified.
	[ ASC   DESC ]	This operation is not supported in VCI.
	[ NULLS { FIRST   LAST } ]	
	WITH	
	WHERE	
	INCLUDE	
CREATE OPERATOR CLASS	-	A VCI cannot be specified.  This operation is not supported in VCI.
CREATE OPERATOR FAMILY	-	
CREATE TABLE	EXCLUDE	
DROP INDEX	CONCURRENTLY	The subcommands on the left cannot be performed if a VCI is specified.
		This operation is not supported in VCI.
REINDEX	-	A VCI cannot be specified.
		This command is not required as VCI uses daemon's automatic maintenance to prevent disk space from increasing.

### Client application command

- Operations that cannot be performed on relations containing a VCI

Command	Description	
clusterdb	Clustering cannot be performed for tables that contain a VCI.	
reindexdb	VCIs cannot be specified on theindex option.	

### 9.2.2 Data Preload Feature

The first aggregation using VCI immediately after an instance is started may take time, because the VCI data has not been loaded to buffer. Therefore, use the preload feature to load the VCI data to buffer in advance when performing VCI aggregation after an instance is started.

When using the preload feature, execute the function pgx\_prewarm\_vci to each VCI created with CREATE INDEX.

db01=# SELECT pgx\_prewarm\_vci('idx\_vci');



Refer to "B.4 VCI Data Load Control Function" for information on pgx\_prewarm\_vci.

## Chapter 10 Parallel Query

FUJITSU Enterprise Postgres enhances parallel queries, by taking into consideration the aspects below:

- CPU load calculation
- Increase of workers during runtime

### 10.1 CPU Load Calculation

There may be a case when the user tries to execute a parallel query but there is not enough CPU available.

Adding dynamic workers at this stage will provide no benefits - instead, it may add overhead due to context switching.

FUJITSU Enterprise Postgres takes into consideration the current CPU load when deciding on the number of workers for parallel query.

## 10.2 Increase of Workers during Runtime

This FUJITSU Enterprise Postgres enhancement allows systems to allocate additional workers during query execution (if there are workers available at the time). This improves query performance, which could otherwise starve of CPU if there were fewer or no workers when the query started.



The ability to increase workers during runtime is available only with parallel sequence scan.

## Chapter 11 High-Speed Data Load

High-speed data load uses the pgx\_loader command to load data from files at high speed into FUJITSU Enterprise Postgres.



- This feature is available only in the Advanced Edition.
- This feature is not available in single-user mode. This is because in single-user mode instances run in a single process, and it cannot start parallel workers.

## 11.1 Installing High-Speed Data Load

This section describes how to install high-speed data load.

#### Installation flow

- 1. Deciding whether to Install
- 2. Estimating Resources
- 3. Setup

### 11.1.1 Deciding whether to Install

The feature achieves high speed data load by executing the COPY FROM command in parallel. If the database system is unable to use sufficient resources due to the feature using more resources than the COPY FROM command of PostgreSQL, load performance may be inferior to that of the COPY FROM command of PostgreSQL. Therefore, determine if the feature will be effective by considering the factors below before deciding to install.

#### Database server memory

If the value of shared\_buffers in postgresql.conf is small, fewer data pages are cached to the shared memory of the database server. This will result in multiple parallel workers more often having to wait for write exclusive locks to the same data page. Moreover, the smaller the number of data pages, the more often the table expands. During table expansion, access to the table is exclusive (standby event name: extend), so write time increases. To cater for that, increase the value of shared\_buffers.



The standby event name is stored in the wait\_event column of the pg\_stat\_activity view. Refer to "wait\_event Description" in "The Statistics Collector" in the PostgreSQL Documentation for details.

#### Frequency of checkpoints

If checkpoints are issued at short intervals, write performance is reduced. If the messages below are output to the server log during data writes, increase the values of max\_wal\_size and checkpoint\_timeout in postgresql.conf to reduce the frequency of checkpoints.



LOG: checkpoints are occurring too frequently (19 seconds apart)
HINT: Consider increasing the configuration parameter "max\_wal\_size".

## 11.1.2 Estimating Resources

Estimate the memory requirements for high-speed data load.

Up to 128 parallel workers to perform data load can be specified for this feature. The additional resources below are required depending on the number of parallel workers.

- Dynamic shared memory created during data load

The feature creates shared memory and shared memory message queues during data load. These are used to send external data from the back end to the parallel workers, and for error notifications.



If the value of shared\_buffers in postgresql.conf is small, the system will often have to wait for write exclusive locks to the same data page (as described in "Database server memory" in "11.1.1 Deciding whether to Install"). Since input data cannot be loaded from the shared memory message queues during such waits, they will often be full. In these cases, it will not be possible to write to the shared memory message queues, resulting in degraded data load performance.



Refer to "High-Speed Data Load Memory Requirements" in the Installation and Setup Guide for Server for information on the formula for estimating memory requirements.

### 11.1.3 Setup

This section describes how to set up high-speed data load.

### Setup flow

- 1. Setting Parameters
- 2. Installing the Extension

### 11.1.3.1 Setting Parameters

Set the parameters required for high-speed data load in postgresql.conf. After that, start or restart the instance.

The table below lists the postgresql.conf parameters that must be changed, and the values that must be added to their current values. After editing postgresql.conf, start or restart the instance.

Parameter	Setting	Required
max_prepared_transactions	Add the number of transactions that can be prepared by parallel workers during data load to the parameter's current value.	Mandatory
	The resulting value must be equal to or greater than the maximum number of parallel workers used with this feature.	
max_worker_processes Number of parallel workers to perform data load.		Mandatory
max_parallel_workers	Add the maximum number of parallel workers to be used in a parallel query by this feature to the parameter's current value.	Mandatory
	The resulting value must be equal to or greater than the number of parallel workers used with this feature.	



The example below shows how to configure 2 instances of high-speed data load being executed simultaneously using a degree of parallelism of 4.

```
max_prepared_transactions = 13 #Example if the initial value was 5: 5 + 2 x 4 = 13
max_worker_processes = 16 #Example if the initial value was 8: 8 + 2 x 4 = 16
max_parallel_workers = 12 #Example if the initial value was 4: 4 + 2 x 4 = 12
```



As shown in the example above, set the value of max\_prepared\_transactions, max\_worker\_processes and max\_parallel\_workers multiplied by the number of instances of this feature executed simultaneously.

The table below lists the postgresql.conf parameters that must also be checked.

Parameter Setting		Required
dynamic_shared_memory_type	Implementation of dynamic shared memory to be used by the instance.	Mandatory
	The default value is recommended.	



Refer to "Resource Consumption" in the PostgreSQL Documentation for information on the parameters.

### 11.1.3.2 Installing the Extension

Execute CREATE EXTENSION to install the high-speed data load extension. The extension needs to be installed on each database.



The example below installs the extension on the "postgres" database.

postgres=# CREATE EXTENSION pgx\_loader; CREATE EXTENSION



- Only superusers can install the high-speed data load extension.
- The high-speed data load extension can only be installed on the public schema.

## 11.2 Using High-Speed Data Load

This section describes how to use high-speed data load.

## 11.2.1 Loading Data

To load data from a file into a FUJITSU Enterprise Postgres table, execute the pgx\_loader command in load mode.



The example below loads the file /path/to/data.csv (2000 records) into table tbl using a degree of parallelism of 3.

```
$ pgx_loader load -j 3 -c "COPY tbl FROM '/path/to/data.csv' WITH CSV"
LOAD 2000
```



If an external file contains data that violates the format or constraints, the data load may fail partway through, resulting in delays for routine tasks such as nightly batch processing. Therefore, it is recommended to remove the invalid data before executing the data load.



The data inserted using this feature is dumped as a COPY command by the pg\_dump command and the pg\_dumpall command.



- Refer to "pgx\_loader" in the Reference for information on the command.
- Refer to "COPY" in the PostgreSQL Documentation for information on the deployment destination and access privileges for external files

### 11.2.2 Recovering from a Data Load that Ended Abnormally

If a system interruption such as a server failure occurs while high-speed data load is being performed, transactions prepared using this feature may be changed to the in-doubt state. At that point, resources occupied by the transaction will be locked, and access to the relevant resources from other transactions will be blocked, rendering them unusable.

In such cases, check transactions that are in an in-doubt state, and resolve them.

### Checking for in-doubt transactions

This section describes how to check for in-doubt transactions.

1. Refer to the pgx\_loader\_state table in the pgx\_loader schema.

Retrieve the global transaction identifier (gid column) of in-doubt transactions. In-doubt transactions will contain "rollback" in the column "state".



The example below retrieves the global transaction identifier (gid) of in-doubt transactions performed by the database role myrole and that used table tbl. The retrieved global transaction identifiers pgx\_loader:9589 and pgx\_loader:9590 identify in-doubt transactions.

2. Refer to the pg\_prepared\_xacts system view.

Check if the in-doubt transactions retrieved above exist.



The example below checks if in-doubt transactions with the global transaction identifiers pgx\_loader:9589 and pgx\_loader:9590 exist.



Refer to "E.1 pgx\_loader\_state" for information on the pgx\_loader\_state table.

### Resolving in-doubt transactions

Execute the pgx\_loader command in recovery mode to resolve in-doubt transactions.

After executing the pgx\_loader command in recovery mode, perform the procedure described in "Checking for in-doubt transactions" to check if the in-doubt transactions have been resolved.



The example below completes the in-doubt transactions prepared for table tbl.

\$ pgx\_loader recovery -t tbl



The recovery mode of the pgx\_loader command only resolves transactions prepared by high-speed data load. For transactions prepared by an application using distributed transactions other than this feature, follow the procedure described in "14.13 Actions in Response to Error in a Distributed Transaction".

## 11.3 Removing High-Speed Data Load

This section describes how to remove high-speed data load.

## 11.3.1 Removing the Extension

Execute DROP EXTENSION to remove the high-speed data load extension. The extension needs to be removed on each database.



The example below removes the extension on the "postgres" database.

postgres=# DROP EXTENSION pgx\_loader;
DROP EXTENSION



- The information required for operation of high-speed data load is stored in the pgx\_loader\_state table of the pgx\_loader schema. Do not remove the high-speed data load extension if the pgx\_loader\_state table is not empty.
- Only superusers can remove the high-speed data load extension.
- The high-speed data load extension can only be removed on the public schema.

## Chapter 12 Global Meta Cache

The Global Meta Cache (GMC) feature loads a meta cache into shared memory using the pgx\_global\_metacache parameter. This reduces the amount of memory required throughout the system.



This feature can only be used in Advanced Edition.

## **12.1 Usage**

Describes how to use the Global Meta Cache feature.

### 12.1.1 Deciding Whether to Enable the Global Meta Cache Feature

Global Meta Cache is a mechanism for sharing meta caches between processes, so it works well on systems with a high number of resources accessed and SQL connections. The number of resources is primarily the number of tables accessed by a process, the number of indexes, or the total number of all columns in all tables accessed.

In particular, consider using Global Meta Cache if the total size of the meta cache for each process exceeds the amount of installed memory, or takes up a large portion of that memory, thereby squeezing memory allocations to the database cache or the Operating system file cache. Using Global Meta Cache may slightly increase the time it takes to execute a single SQL to reference a meta cache on shared memory (For example, it will not exceed 5% of the execution time in a situation where all the data is located in the database cache), but you can expect a greater benefit from being able to allocate more memory, such as for the database cache.

If performance degradation using Global Meta Cache is not acceptable, you may want to limit the number of tables accessed by a process.

### 12.1.2 Estimating Memory for Global Meta Cache

To enable the Global Meta Cache feature, the pgx\_global\_metacache parameter must specify an upper limit on the size of the shared memory (Hereinafter, the GMC area) dedicated to Global Meta Cahche. Ideally, this upper limit should be the size estimated in "Appendix A Parameters". Values lower than this can still work, but refer to "12.1.3 How the GMC Memory Area Is Used" on using the GMC area to understand the disadvantages.

## 12.1.3 How the GMC Memory Area Is Used

At startup, the memory for the GMC area is not used much, but the GMC area grows as new meta caches are placed in the GMC area. If it does, it discards any meta caches that the system determines are not heavily used and places a new one in the GMC area.

Therefore, the GMC area will work even if it is smaller than the estimate, but the meta cache will be regenerated if the discarded meta cache needs to be reused. Note that if this happens frequently, it will degrade overall performance.

With this in mind, it may not be a problem if, for example, the tables to be accessed are different depending on the time zone, and the degradation of the time zone immediately after the change is acceptable.

In any case, be sure to test and tune the system thoroughly before running it.

## 12.1.4 Enabling the Global Meta Cache Feature

To Enable the Global Meta Cache feature edit the postgresql.conf file and set the pgx\_global\_metacahe parameter. Restarting the instance after editing the postgresql.conf file is required. Refer to "Appendix A Parameters" for information on the parameters.

Parameter Name	Description	
pgx_global_metacache	Specify the maximum amount of memory for the GMC area on shared memory. When it's s to 0 (default value), the Global Meta Cache feature is disabled.	
	When enabled, the minimum value allowed is 10MB.	

When the cache is created, if the total amount of meta caches on shared memory exceeds the value specified by pgx\_global\_metacache, the inactive, unreferenced meta caches are removed from the GMC area. Note that if all GMC are in use and the cache cannot be created in the GMC area, the cache is temporarily created in the local memory of the backend process.



Here is an example postgresql.conf configuration:

pgx\_global\_metacache = 800 MB

#### Wait Events

The Global Meta Cache feature may cause wait events. Wait events are identified in the wait\_event column of the pg\_stat\_activity view. GMC specific wait events are described below.

[GMC Feature Wait Events]

Wait Event Type	Wait Event Name	Description
LWLock	global_metacache	Waiting to find, add, and remove meta caches in the GMC area.
	global_metacache_sweep	Waiting to select a meta cache that can be deleted when GMC space is low.
		If the GMC is fully referencing and there is no deletable meta cache, it is waiting for the reference to be removed and a deletable meta cache to be selected.



If global\_metacache\_sweep is happened frequently, increase the pgx\_global\_metacache setting.



See

Refer to "Viewing Statistics" in the PostgreSQL Documentation for information on the pg\_stat\_activity view.

### 12.1.5 Estimating Resources

Refer to "Global Meta Cache Memory Requirements" in the Installation and Setup Guide for Server for formulas to estimate the amount of memory used by the Global Meta Cache feature.

### 12.2 Statistics

Describes the statistics for the Global Meta Cache feature.

## 12.2.1 System View

You can check the cache hit ratio and size of the GMC area in the system view pgx\_stat\_gmc. Refer to "C.6 pgx\_stat\_gmc" for information on the columns.

If the cache hit ratio is low and the current memory usage is close to pgx\_global\_metacache, increase the pgx\_global\_metacache setting because performance may be degraded.

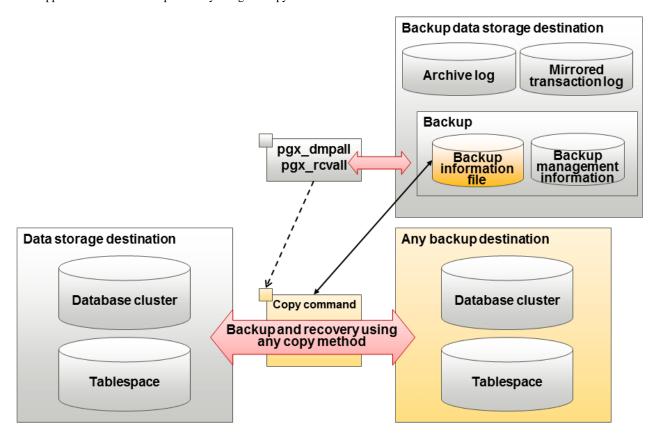
Refer to "7.6 Monitoring Database Activity" in the Operations Guide for information on the statistics.

## Chapter 13 Backup/Recovery Using the Copy Command

By using a copy command created by the user, the pgx\_dmpall command and the pgx\_rcvall command can perform backup to any destination and can perform recovery from any destination using any copy method.

Copy commands must be created in advance as executable scripts for the user to implement the copy process on database clusters and tablespaces, and are called when executing the pgx\_dmpall and pgx\_rcvall commands.

This appendix describes backup/recovery using the copy command.





- By using the high-speed copy feature of the storage device to copy the data storage destination, the processing time for backup of large databases can be greatly reduced.
- It is also possible to back up only some tablespaces using the copy command. However, database resources not backed up using the copy command are still backed up to the backup data storage destination.



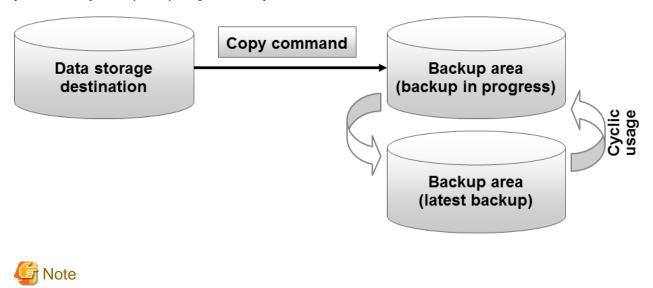
Both the backup data storage destination and the optional backup destination are necessary for recovery - if they are located in secondary media, combined management of these is necessary.

## 13.1 Configuration of the Copy Command

This section describes the configuration of the copy command for backup and recovery.

### Cyclic usage of the backup area

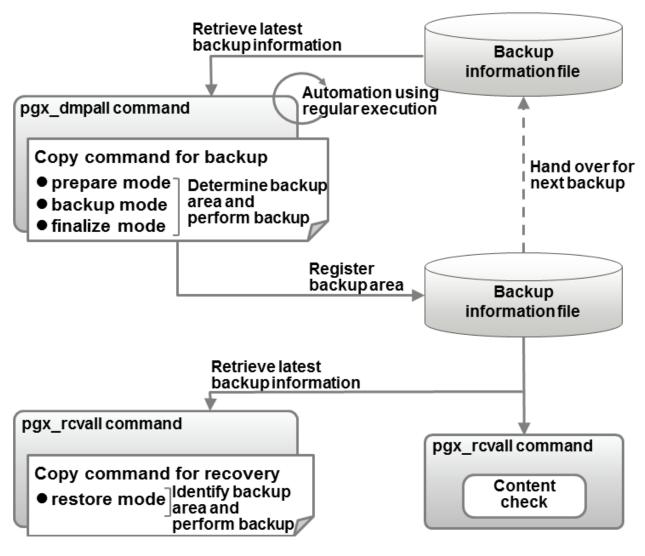
Prepare two backup areas for the copy command in case an issue affects the data storage destination during backup. The copy command performs backup while cyclically using these backup areas.



The backup data storage destination cannot be used as these backup areas used by the copy command.

### Backup using the backup information file

The copy command must determine the backup destination on each backup, as it is necessary to cycle through the backup areas. Backup can be automated by using the backup information file, which contains information about the backup destination.





The backup information file is prepared in the backup data storage destination by the pgx\_dmpall command, and contains information that can be read or updated by the copy command. This file is managed by associating it with the latest backup successfully completed by the pgx\_dmpall command, so the latest backup information relating to the copy command registered by the user can be retrieved. Additionally, the content of the backup information file can be displayed using the pgx\_revall command.

### Configuration of the copy command for backup

The pgx\_dmpall command calls the copy command for backup after execution for the three modes below. It is therefore necessary for the copy command for backup to implement the required processing for each of the modes.

- prepare mode

Determines which of the two backup areas will be used for the current backup.

The backup area to be used for the current backup is determined by reading the information relating to the latest backup destination where the backup information file was written to during the previous backup.

- backup mode

Performs backup on the backup area determined by prepare mode, using any copy method.

- finalize mode

Writes information relating to the destination of the current backup to the backup information file.

This enables the prepare mode to check the destination of the previous backup during the next backup.



The user can use any method to hand over backup information between modes within the copy command, such as creating temporary files.

### Configuration of the copy command for recovery

The pgx\_rcvall command calls the copy command for recovery for the mode below. It is therefore necessary for the copy command for recovery to implement the required processing for the mode.

- restore mode

Any copy method can be used to implement restore from the backup destination retrieved using the copy command for backup.



By referring to the mode assigned to the copy command as an argument, backup and recovery can be implemented using a single copy command.

......



### Using a bash script

```
case $1 in
   prepare)
   processingRequiredForPrepareMode
   ;;
  backup)
   processingRequiredForBackupMode
   ;;
  finalize)
   processingRequiredForFinalizeMode
   ;;
  restore)
   processingRequiredForRestoreMode
  ;;
  restore)
   processingRequiredForRestoreMode
  ;;
  restore)
```

## Point

- A sample script that backs up the database cluster and tablespace directory to a specific directory is supplied to demonstrate how to write a copy command.

The sample is stored in the directory below:

```
/installDir/share/copy_command.archive.sh.sample
```

- A sample script that uses OPC (an advanced copy feature of Fujitsu Storage ETERNUS disk array) is supplied. Refer to "Appendix I Copy Command Samples that Use the Advanced Copy Feature of the ETERNUS Disk Array" for details.

The samples cannot be used on SLES 12.

## 13.2 Backup Using the Copy Command

To perform backup using the copy command, in addition to performing the standard backup procedure, it is also necessary to create a copy command, and then execute the pgx\_dmpall command specifying it. This section describes the procedure specific to using the copy command.

### Preparing for backup

You must prepare for backup before actually starting the backup process.

Perform the following procedure:

1. Determine the database resources to be backed up

Determine the database resources to be backed up using the copy command. The copy command can back up the following resources:

- Database cluster
- Tablespace

To back up only some tablespaces, create a file listing them. This file is not necessary to back up all tablespaces.

#### Example

To back up only tablespaces tblspc1 and tblspc2

tblspc1 tblspc2

2. Prepare a backup area

Prepare a backup area to save the database resources to be backed up, as determined in step 1.

3. Create the copy command

Create the copy commands for backup and recovery. Refer to "13.4 Copy Command Interface" for details.

### Performing backup

Execute the pgx\_dmpall command with the -Y option specifying the full path of the copy command for backup created in step 3 of preparation for backup.

The example below backs up only some tablespaces, but not the database cluster, using the copy command.



```
$ pgx_dmpall -D /database/inst1 -Y '/database/command/backup.sh'
--exclude-copy-cluster -P '/database/command/tablespace_list.txt'
```



- To exclude up the database cluster from backup using the copy command, specify the --exclude-copy-cluster option.
- To back up only some tablespaces using the copy command, use the -P option specifying the full path of the file created in step 1 of preparation for backup.



- Refer to "pgx\_dmpall" in the Reference for information on the command.

### Checking backup status

Use the pgx\_rcvall command to check the backup status.

Execute the pgx\_rcvall command with the -l option specified to output backup data information. If backup was performed using the copy command, the resources backed up using the copy command will also be output.



## 13.3 Recovery Using the Copy Command

To perform recovery using the copy command, in addition to performing the standard recovery procedure, it is also necessary to create a copy command, and then execute the pgx\_rcvall command specifying it. This section describes the procedure specific to using the copy command.

### Determining the backup area of the latest backup

Check the backup information file to determine the backup area used for the latest backup, and confirm that it is in a recoverable state.

Execute the pgx\_rcvall command with the --view-results-of-copying option to output the content of the backup information file.



\$ pgx\_rcvall -D /database/inst1 --view-results-of-copying

### Perform recovery

Execute the pgx\_rcvall command with the -Y option specifying the full path of the copy command for recovery created in step 3 of the preparation for backup described in "13.2 Backup Using the Copy Command".

The example below recover only some tablespaces, but not the database cluster, using the copy command.



\$ pgx\_rcvall -D /database/inst1 -B /backup/inst1 -Y '/database/command/recovery.sh'



If the latest backup was performed using the copy command, the pgx\_rcvall command automatically recognizes which database resources were backed up using the copy command, or whether resources were backed up to the backup data storage destination. Therefore, recovery can be performed by simply executing the pgx\_rcvall command specifying the copy command for recovery.



Refer to "pgx\_rcvall" in the Reference for information on the command.

## 13.4 Copy Command Interface

The following types of copy command are available:

- Copy command for backup
- Copy command for recovery

This appendix describes the interface of each copy command.

### 13.4.1 Copy Command for Backup

### **Feature**

User exit (for the copy command) called from the pgx\_dmpall command.

#### **Format**

The syntax for calling the copy command from the pgx\_dmpall command is described below.

If the operation mode is "prepare"

```
copyCommandName prepare 'pathOfBackupInfoFile' 'pathOfBackupTargetListFile'
```

If the operation mode is "backup"

```
copyCommandName backup
```

If the operation mode is "finalize"

```
copyCommandName finalize 'pathOfBackupInfoFile'
```

### Argument

- Operation mode

Mode	Description	
prepare	Implements the preparation process for backing up using the copy command.	
	Called before the PostgreSQL online backup mode is started.	
backup	Implements the backup process.	
	Called during the PostgreSQL online backup mode.	
finalize	Implements the backup completion process.	
	Called after the PostgreSQL online backup mode is completed.	



If using high-speed copy of the storage device (which performs high-speed retrieval of snapshots and copies data to different physical areas), it is possible to invoke the snapshot process in backup mode, and the copy process in finalize mode.

- Full path of the backup information file

Full path of the backup information file of the latest backup, enclosed in single quotation marks. If a backup has not been performed, specify '-'.

- Full path of the backup target list file

Full path of the file containing the resources to be backed up using the copy command, enclosed in single quotation marks. One of the following is described in each resource name.

Resource	Description	
Database cluster	pg_data	
Tablespace	Tablespace name	



To back up the database cluster and the tablespaces dbspace and indexspace using the copy command, the file should contain the following:

pg_data	
dbspace	
indexspace	



The encoding of resource names output to the backup target list file by the pgx\_dmpall command is the encoding used when this command connects to the database with auto specified for the client\_encoding parameter, and is dependent on the locale at the time of command execution.

The number of arguments vary depending on operation mode. The argument of each operation mode is as follows.

Operation mode	First argument	Second argument	Third argument
prepare	Operation mode	Backup information file path name	Backup target list file path name
backup		None	None
finalize		Backup information file path name	

Additionally, the access permissions for the backup information file and backup target list file are different depending on the operation mode. The access permissions of each operation mode are as follows.

Operation mode	Backup information file	Backup target list file
prepare	Can be viewed by the instance administrator only	Can be viewed by the instance administrator only
backup	-	-
finalize	Can be viewed and updated by the instance administrator only	-

### Return value

Return value	Description	
0	Normal end	
	The pgx_dmpall command continues processing.	
Other than 0	Abnormal end	
	The pgx_dmpall command terminates in error.	

### Description

- The copy command operates with the privileges of the operating system user who executed the pgx\_dmpall command. Therefore, grant copy command execution privileges to users who will execute the pgx\_dmpall command. Additionally, have the copy command change users as necessary.

- To write to the backup information file, use a method such as redirection from the copy command.
- Because the copy command is called for each mode, implement all processing for each one.
- To copy multiple resources simultaneously, have the copy command copy them in parallel.



- The backup information file and backup target list file cannot be deleted. Additionally, the privileges cannot be changed.
- Standard output and standard error output of the copy command are output to the terminal where the pgx\_dmpall command was executed.
- If the copy command becomes unresponsive, the pgx\_dmpall command will also become unresponsive. If the copy command is deemed to be unresponsive by the operating system, use an operating system command to forcibly stop it.
- Output the copy command execution trace and the result to a temporary file, so that if it terminates in error, the cause can be investigated at a later time.
- For prepare mode only, it is possible to use the PostgreSQL client application to access the database using the copy command. For all other modes, do not execute FUJITSU Enterprise Postgres commands or PostgreSQL applications.
- Enable the fsync parameter in postgresql.conf, because data on the shared memory buffer needs to have been already written to disk when backup starts.

### 13.4.2 Copy Command for Recovery

#### **Feature**

User exit (for the copy command) called from the pgx\_rcvall command.

### **Format**

The syntax for calling the copy command from the pgx\_rcvall command is described below.

copyCommandName restore 'pathOfBackupInfoFile' 'pathOfBackupTargetListFile'

### Argument

- Operation mode

Mode	Description	
restore	Performs restore.	

- Full path of the backup information file

Full path of the backup information file, enclosed in single quotation marks.

- Full path of the backup target list file

Full path of the file containing the resources to be restored using the copy command, enclosed in single quotation marks.

The access permissions for the backup information file and backup target list file are as below.

Backup information file	Backup target list file	
Can be viewed by the instance administrator only	Can be viewed by the instance administrator only	

#### Return value

Return value	Description	
0	Normal end	
	The pgx_rcvall command continues processing.	
Other than 0	Abnormal end	
	The pgx_rcvall command terminates in error.	

### Description

- The copy command operates with the privileges of the operating system user who executed the pgx\_rcvall command. Therefore, grant copy command execution privileges to users who will execute the pgx\_rcvall command. Additionally, have the copy command change users as necessary.
- The copy command is called once only in restore mode.
- To copy multiple resources simultaneously, have the copy command copy them in parallel.



- The backup information file and backup target list file cannot be deleted. Additionally, the privileges cannot be changed.
- Standard output and standard error output of the copy command are output to the terminal where the pgx\_rcvall command was executed.

- If the copy command becomes unresponsive, the pgx\_rcvall command will also become unresponsive. If the status of the copy command is deemed to be unresponsive by the operating system, use an operating system command to forcibly stop it.
- Output the copy command execution trace and the result to a temporary file, so that if it terminates in error, the cause can be investigated at a later time.
- Do not execute FUJITSU Enterprise Postgres commands or PostgreSQL applications in the copy command.
- There may be files and directories not required for recovery using the archive log included in the backup, such as postmaster.pid, pg\_wal/subdirectory and pg\_replslot in the database cluster. If such unnecessary files and directories exist, have the copy command delete them after the restore.

## Chapter 14 Actions when an Error Occurs

This chapter describes the actions to take when an error occurs in the database or an application, while FUJITSU Enterprise Postgres is operating.

Depending on the type of error, it may be necessary to recover the database cluster. The recovery process recovers the following resources:

- Data storage destination
- Transaction log storage destination (if the transaction log is stored in a separate disk from the data storage destination)
- Backup data storage destination



Even if a disk is not defective, the same input-output error messages, as those generated when the disk is defective, may be output. The recovery actions differ for these error messages.

Check the status of the disk, and select one of the following actions:

- If the disk is defective

Refer to "14.1 Recovering from Disk Failure (Hardware)", and take actions accordingly.

- If the disk is not defective

Refer to "14.14 I/O Errors Other than Disk Failure", and take actions accordingly.

A few examples of errors generated even if the disk is not defective include:

- Network error with an external disk
- Errors caused by power failure or mounting issues

### Determining the cause of an error

If an error occurs, refer to the WebAdmin message and the server log, and determine the cause of the error.



Refer to "Configuring Parameters" in the Installation and Setup Guide for Server for information on server logs.

### Approximate recovery time

The formulas for deriving the approximate recovery time of resources in each directory are given below.

If using the copy command with the pgx\_rcvall command, the recovery time will depend on the implementation of the copy command.

- Data storage destination or transaction log storage destination

Recovery time = (usageByTheDataStorageDestination + usageByTheTransactionLogStorageDestination) /  $diskWritePerformance \times 1.5$ 

- usageByTheDataStorageDestination: Disk space used by the database cluster
- usageByTheTransactionLogStorageDestination. Disk space used by the transaction log stored outside the database cluster
- *diskWritePerformance*: Measured maximum data volume (bytes/second) that can be written per second in the system environment where the operation is performed
- 1.5: Coefficient assuming the time excluding disk write, which is the most time-consuming step
- Backup data storage destination

 ${\tt Recovery\ time = usageByTheBackupDataStorageDestination\ /\ diskWritePerformance\ x\ 1.5}$ 

- usageByTheBackupDataStorageDestination: Disk space used by the backup data
- *diskWritePerformance*: Measured maximum data volume (bytes/second) that can be written per second in the system environment where the operation is performed
- 1.5: Coefficient assuming the time excluding disk write, which is the most time-consuming step

## 14.1 Recovering from Disk Failure (Hardware)

This section describes how to recover database clusters to a point immediately before failure, if a hardware failure occurs in the data storage disk or the backup data storage disk.

There are two methods of recovery:

- 14.1.1 Using WebAdmin
- 14.1.2 Using Server Command



Back up the database cluster after recovering it. Backup deletes obsolete archive logs (transaction logs copied to the backup data storage destination), freeing up disk space and reducing the recovery time.

### 14.1.1 Using WebAdmin

Recover the database cluster by following the appropriate recovery procedure below for the disk where the failure occurred.



Recovery operation cannot be performed on an instance that is part of a streaming replication cluster in standby mode.

If disk failure occurs on a standby instance, it may be necessary to delete and re-create the instance.

Recovery operation can be performed on an instance that is part of a streaming replication cluster in "Master" mode. If a recovery operation is performed on a master instance, it will break the replication cluster and streaming replication will stop between the master instance and all its standby instances. In such an event, the standby instances can be promoted to standalone instances or can be deleted and re-created.

### If failure occurred in the data storage disk or the transaction log storage disk

Follow the procedure below to recover the data storage disk or the transaction log storage disk.

1. Stop applications

Stop applications that are using the database.

2. Stop the instance

Stop the instance. Refer to "2.1.1 Using WebAdmin" for information on how to stop an instance. WebAdmin automatically stops instances if recovery of the database cluster is performed without stopping the instance.

3. Recover the failed disk

Replace the disk, and then recover the volume configuration information.

4. Create a tablespace directory

If a tablespace was defined after backup, create a directory for it.

5. Recover the keystore, and enable automatic opening of the keystore

Do the following if the data in the database has been encrypted:

- Restore the keystore to its state at the time of the database backup.
- Enable automatic opening of the keystore.

#### 6. Recover the database cluster

Log in to WebAdmin, and in the [Instances] tab, click [Solution] for the error message in the lower-right corner.

#### 7. Run recovery

In the [Restore Instance] dialog box, click [Yes].

Instance restore is performed. An instance is automatically started when recovery is successful.



WebAdmin does not support recovery of hash index. If you are using a hash index, then after recovery, execute the REINDEX command to rebuild it. Use of hash indexes is not recommended.

#### 8. Resume applications

Resume applications that are using the database.



WebAdmin may be unable to detect disk errors, depending on how the error occurred.

If this happens, refer to "14.10.3 Other Errors" to perform recovery.

### If failure occurred on the backup data storage disk

Follow the procedure below to recover the backup data storage disk.

1. Recover the failed disk

Replace the disk, and then recover the volume configuration information.

2. Recover the backup data

Log in to WebAdmin, and in the [Instances] tab, click [Solution] for the error message.

3. Run backup

 $Perform\ backup\ to\ enable\ recovery\ of\ the\ backup\ data.\ In\ the\ [Backup]\ dialog\ box,\ click\ [Yes].\ The\ backup\ is\ performed.\ An\ instance\ is\ automatically\ started\ when\ backup\ is\ performed.$ 

.....



If you click [Recheck the status], the resources in the data storage destination and the backup data storage destination are reconfirmed. As a result, the following occurs:

- If an error is not detected

The status of the data storage destination and the backup data storage destination returns to normal, and it is possible to perform operations as usual.

- If an error is detected

An error message is displayed in the message list again. Click [Solution], and resolve the problem by following the resolution for the cause of the error displayed in the dialog box.

## 14.1.2 Using Server Command

Recover the database cluster by following the appropriate recovery procedure below for the disk where the failure occurred.

### If failure occurred on the data storage disk or the transaction log storage directory

Follow the procedure below to recover the data storage disk or the transaction log storage directory.

### 1. Stop applications

Stop applications that are using the database.

### 2. Stop the instance

Stop the instance, refer to "2.1.2 Using Server Commands" for details.

If the instance fails to stop, refer to "14.11 Actions in Response to Failure to Stop an Instance".

3. Recover the failed disk

Replace the disk, and then recover the volume configuration information.

- 4. Create a storage destination directory
  - If failure occurred on the data storage disk
    Create a data storage destination directory. If a tablespace was defined, also create a directory for it.
  - If failure occurred on the translation log storage disk Create a transaction log storage destination directory.

#### Example

To create a data storage destination directory:

```
$ mkdir /database/inst1
$ chown fsepuser:fsepuser /database/inst1
$ chmod 700 /database/inst1
```



#### See

Refer to "Preparing Directories to Deploy Resources" under "Setup" in the Installation and Setup Guide for Server for information on how to create a storage directory.

5. Recover the keystore, and enable automatic opening of the keystore

When the data in the database has been encrypted, restore the keystore to its state at the time of the database backup. Configure automatic opening of the keystore as necessary.

6. Recover the database cluster

Recover the database cluster using the backup data.

Specify the following in the pgx\_rcvall command:

- Specify the data storage location in the -D option. If the -D option is omitted, the value of the PGDATA environment variable is used by default.
- Specify the backup data storage location in the -B option.

#### Example

> pgx\_rcvall -D /database/inst1 -B /backup/inst1



If recovery fails, remove the cause of the error in accordance with the displayed error message and then re-execute the pgx\_rcvall command.

If the message "pgx\_rcvall: an error occurred during recovery" is displayed, then the log recorded when recovery was executed is output after this message. The cause of the error is output in around the last fifteen lines of the log, so remove the cause of the error in accordance with the message and then re-execute the pgx\_rcvall command.

The following message displayed during recovery is output as part of normal operation of pgx\_rcvall command (therefore the user does not need not be concerned).

```
FATAL: The database system is starting
```

7. Start the instance

Start the instance.

Refer to "2.1.2 Using Server Commands" for information on how to start an instance.

8. Resume applications

Resume applications that are using the database.

### If failure occurred on the backup data storage disk

The procedure for recovering the backup data storage disk is described below.

There are two methods of taking action:

- Performing recovery while the instance is active
- Stopping the instance before performing recovery

The following table shows the different steps to be performed depending on whether you stop the instance.

No	Step	Instance stopped	
		No	Yes
1	Confirm that transaction log mirroring has stopped	Y	N
2	Stop output of archive logs	Y	N
3	Stop applications	N	Y
4	Stop the instance	N	Y
5	Recover the failed disk	Y	Y
6	Create a backup data storage destination directory	Y	Y
7	Resume output of archive logs	Y	N
8	Resume transaction log mirroring	Y	N
9	Start the instance	N	Y
10	Run backup	Y	Y
11	Resume applications	N	Y

Y: Required

N: Not required

The procedure is as follows:

If an instance has not been stopped

1. Confirm that transaction log mirroring has stopped

Use the following SQL function to confirm that transaction log mirroring has stopped.

```
postgres=# SELECT pgx_is_wal_multiplexing_paused();
pgx_is_wal_multiplexing_paused
-----t
t
(1 row)
```

If transaction log mirroring has not stopped, then stop it using the following SQL function.

```
postgres=# SELECT pgx_pause_wal_multiplexing();
LOG: multiplexing of transaction log files has been stopped
```

```
pgx_pause_wal_multiplexing
------
(1 row)
```

### 2. Stop output of archive logs

Transaction logs may accumulate during replacement of backup storage disk, and if the data storage disk or the transaction log storage disk becomes full, there is a risk that operations may not be able to continue.

To prevent this, use the following methods to stop output of archive logs.

- Changing archive\_command

Specify a command that will surely complete normally, such as "echo skipped archiving WAL file %f" or "/bin/true", so that archive logs will be regarded as having been output.

If you specify echo, a message is output to the server log, so it may be used as a reference when you conduct investigations.

- Reload the configuration file

Execute the pg\_ctl reload command or the pg\_reload\_conf SQL function to reload the configuration file.

If you simply want to stop output of errors without the risk that operations will not be able to continue, specify an empty string (") in archive\_command and reload the configuration file.

3. Recover the failed disk

Replace the disk, and then recover the volume configuration information.

4. Create a backup data storage destination

Create a backup data storage destination.

#### Example

```
$ mkdir /database/inst1
$ chown fsepuser:fsepuser /database/inst1
$ chmod 700 /database/inst1
```

Refer to "3.2.2 Using Server Commands" for information on how to create a backup data storage destination.

5. Resume output of archive logs

Return the archive\_command setting to its original value, and reload the configuration file.

6. Resume transaction log mirroring

Execute the pgx\_resume\_wal\_multiplexing SQL function.

#### Example

```
SELECT pgx_resume_wal_multiplexing()
```

### 7. Run backup

Use the pgx\_dmpall command to back up the database cluster.

Specify the following value in the pgx\_dmpall command:

- Specify the data storage destination in the -D option. If the -D option is omitted, the value of the PGDATA environment variable is used by default.

### Example

```
> pgx_dmpall -D /database/inst1
```

### If an instance has been stopped

1. Stop applications

Stop applications that are using the database.

2. Stop the instance

Stop the instance. Refer to "2.1.2 Using Server Commands" for details.

If the instance fails to stop, refer to "14.11 Actions in Response to Failure to Stop an Instance".

3. Recover the failed disk

Replace the disk, and then recover the volume configuration information.

4. Create a backup data storage destination

Create a backup data storage destination.

### Example

```
# mkdir /backup/inst1
# chown fsepuser:fsepuser /backup/inst1
# chmod 700 /backup/inst1
```

Refer to "3.2.2 Using Server Commands" for details.

5. Start the instance

Start the instance. Refer to "2.1.2 Using Server Commands" for information on how to start an instance.

6. Run backup

Use the pgx\_dmpall command to back up the database cluster.

Specify the following value in the pgx\_dmpall command:

- Specify the data storage destination in the -D option. If the -D option is omitted, the value of the PGDATA environment variable is used by default.

### Example

```
> pgx_dmpall -D /database/inst1
```

7. Resume applications

Resume applications that are using the database.



- Refer to "pgx\_rcvall" and "pgx\_dmpall" in the Reference for information on the pgx\_rcvall command and pgx\_dmpall command.
- Refer to "Write Ahead Log" under "Server Administration" in the PostgreSQL Documentation for information on archive\_command.
- Refer to "B.1 WAL Mirroring Control Functions" for information on pgx\_resume\_wal\_multiplexing.

## 14.2 Recovering from Data Corruption

If data in a disk is logically corrupted and the database does not operate properly, you can recover the database cluster to its state at the time of backup.

There are two methods of recovery:

- 14.2.1 Using WebAdmin
- 14.2.2 Using the pgx\_rcvall Command



- Back up the database cluster after recovering it. Backup deletes obsolete archive logs (transaction logs copied to the backup data storage destination), freeing up disk space and reducing the recovery time.
- If you recover data to a point in the past, a new time series (database update history) will start from that recovery point. When recovery is complete, the recovery point is the latest point in the new time series. When you subsequently recover data to the latest state, the database update is re-executed on the new time series.

### 14.2.1 Using WebAdmin

If using WebAdmin, recover the data to the point immediately prior to data corruption by using the backup data.

Refer to "14.1.1 Using WebAdmin" for details.

### 14.2.2 Using the pgx\_rcvall Command

Recover the database cluster by specifying in the pgx\_rcvall command the date and time of the backup you want to read from. Then reexecute the transaction as required to recover the data.

Follow the procedure below to recover the data storage disk.

1. Stop applications

Stop applications that are using the database.

2. Stop the instance

Stop the instance. Refer to "2.1.2 Using Server Commands" for information on how to stop an instance.

If the instance fails to stop, refer to "14.11 Actions in Response to Failure to Stop an Instance".

3. Confirm the backup date and time

Execute the pgx\_rcvall command to confirm the backup data saved in the backup data storage destination, and determine a date and time prior to data corruption.

Specify the following values in the pgx\_rcvall command:

- Specify the data storage destination in the -D option. If the -D option is omitted, the value of the PGDATA environment variable is used by default.
- Specify the backup storage directory in the -B option.
- The -l option displays the backup data information.

### Example

```
> pgx_rcvall -D /database/inst1 -B /backup/inst1 -l
Date Status Dir
2017-05-20 10:00:00 COMPLETE /backup/inst1/2017-05-20_10-00-00
```

4. Recover the keystore, and enable automatic opening of the keystore

When the data in the database has been encrypted, restore the keystore to its state at the time of the database backup. Configure automatic opening of the keystore as necessary.

5. Recover the database cluster

Use the pgx\_rcvall command to recover the database cluster.

Specify the following values in the pgx\_rcvall command:

- Specify the data storage destination in the -D option. If the -D option is omitted, the value of the PGDATA environment variable is used by default.
- Specify the backup storage directory in the -B option.

- Specify the recovery date and time in the -e option.

#### Example

In the following examples, "May 20, 2017 10:00:00" is specified as the recovery time.

pgx\_rcvall -D /database/inst1 -B /backup/inst1 -e '2017-05-20 10:00:00'



If recovery fails, remove the cause of the error in accordance with the displayed error message and then re-execute the pgx\_rcvall command.

If the message "pgx\_rcvall: an error occurred during recovery" is displayed, then the log recorded when recovery was executed is output after this message. The cause of the error is output in around the last fifteen lines of the log, so remove the cause of the error in accordance with the message and then re-execute the pgx\_rcvall command.

The following message displayed during recovery is output as part of normal operation of pgx\_rcvall command (therefore the user does not need not be concerned).

FATAL: The database system is starting

#### 6. Start the instance

Start the instance. Refer to "2.1.2 Using Server Commands" for information on how to start an instance.

If necessary, re-execute transaction processing from the specified recovery time, and then resume database operations.



The pgx\_rcvall command cannot accurately recover a hash index. If you are using a hash index, wait for the instance to start and then execute the REINDEX command for the appropriate index.

### 7. Resume applications

Resume applications that are using the database.



See

Refer to "pgx\_rcvall" in the Reference for information on the pgx\_rcvall command.

## 14.3 Recovering from an Incorrect User Operation

This section describes how to recover database clusters when data has been corrupted due to erroneous user operations.

There are two methods of recovery:

- 14.3.1 Using WebAdmin
- 14.3.2 Using the pgx\_rcvall Command



- Back up the database cluster after recovering it. Backup deletes obsolete archive logs (transaction logs copied to the backup data storage destination), freeing up disk space and reducing the recovery time.

- If you recover data to a point in the past, a new time series (database update history) will start from that recovery point. When recovery is complete, the recovery point is the latest point in the new time series. When you subsequently recover data to the latest state, the database update is re-executed on the new time series.

- An effective restore point is one created on a time series for which you have made a backup. That is, if you recover data to a point in the past, you cannot use any restore points set after that recovery point. Therefore, once you manage to recover your target past data, make a backup.

### 14.3.1 Using WebAdmin

You can use WebAdmin to recover data to a backup point.



Recovery operation cannot be performed on an instance that is part of a streaming replication cluster in standby mode.

If disk failure occurs on a standby instance, it may be necessary to delete and re-create the instance.

Recovery operation can be performed on an instance that is part of a streaming replication cluster in "Master" mode. If a recovery operation is performed on a master instance, it will break the replication cluster and streaming replication will stop between the master instance and all its standby instances. In such an event, the standby instances can be promoted to standalone instances or can be deleted and re-created.

Follow the procedure below to recover the data in the data storage disk.

1. Stop applications

Stop applications that are using the database.

2. Stop the instance

Stop the instance. Refer to "2.1.1 Using WebAdmin" for information on how to stop an instance.

3. Recover the keystore, and enable automatic opening of the keystore

Do the following if the data in the database has been encrypted:

- Restore the keystore to its state at the time of the database backup.
- Enable automatic opening of the keystore.
- 4. Recover the database cluster

Log in to WebAdmin, and in the [Instances] tab, select the instance to be recovered and click 🔜.



5. Recover to the backup point

In the [Restore Instance] dialog box, click [Yes].

Recovery is performed. An instance is automatically started when recovery is successful.



WebAdmin cannot accurately recover a hash index. If you are using a hash index, then after recovery, execute the REINDEX command for the appropriate index.

6. Resume database operations

If necessary, re-execute transaction processing from the backup point to when an erroneous operation was performed, and then resume database operations.

## 14.3.2 Using the pgx\_rcvall Command

The pgx\_rcvall command recovers database clusters to the restore point created with the server command. Refer to "Setting a restore point" in "3.2.2 Using Server Commands" for information on how to create a restore point.

Follow the procedure below to recover the data in the data storage disk.

### 1. Stop applications

Stop applications that are using the database.

#### 2. Stop the instance

Stop the instance. Refer to "2.1.2 Using Server Commands" for information on how to stop an instance.

If the instance fails to stop, refer to "14.11 Actions in Response to Failure to Stop an Instance".

### 3. Confirm the restore point

Execute the pgx\_rcvall command to confirm the backup data saved in the backup data storage destination, and use a restore point recorded in an arbitrary file, as explained in "3.2.2 Using Server Commands", to determine a restore point prior to the erroneous operation.

Specify the following values in the pgx\_rcvall command:

- Specify the data storage destination in the -D option. If the -D option is omitted, the value of the PGDATA environment variable is used by default.
- Specify the backup data storage destination in the -B option.
- The -l option displays the backup data information.

### Example

```
> pgx_rcvall -D /database/inst1 -B /backup/inst1 -l
Date Status Dir
2017-05-01 10:00:00 COMPLETE /backup/inst1/2017-05-01_10-00-00
```

#### 4. Recover the keystore, and enable automatic opening of the keystore

When the data in the database has been encrypted, restore the keystore to its state at the time of the database backup. Configure automatic opening of the keystore as necessary.

### 5. Recover the database cluster

Use the pgx\_rcvall command to recover the database cluster.

Specify the following values in the pgx\_rcvall command:

- Specify the data storage destination in the -D option. If the -D option is omitted, the value of the PGDATA environment variable is used by default.
- Specify the backup data storage destination in the -B option.
- The -n option recovers the data to the specified restore point.

### Example

The following example executes the pgx\_rcvall command with the restore point "batch\_20170503\_1".

```
> pgx_rcvall -D /database/inst1 -B /backup/inst1 -n batch_20170503_1
```



If recovery fails, remove the cause of the error in accordance with the displayed error message and then re-execute the pgx\_rcvall command.

If the message "pgx\_rcvall: an error occurred during recovery" is displayed, then the log recorded when recovery was executed is output after this message. The cause of the error is output in around the last fifteen lines of the log, so remove the cause of the error in accordance with the message and then re-execute the pgx\_rcvall command.

The following message displayed during recovery is output as part of normal operation of pgx\_rcvall (therefore the user does not need not be concerned).

FATAL: The database system is starting

#### 6. Start the instance

Start the instance.

Refer to "2.1.2 Using Server Commands" for information on how to start an instance.



The pgx\_rcvall command cannot accurately recover a hash index. If you are using a hash index, wait for the instance to start and then execute the REINDEX command for the appropriate index.

7. Restart operation of the database

If necessary, re-execute transaction processing from the specified recovery time to the point when an erroneous operation was performed, and then resume database operations.



Refer to "pgx\_rcvall" in the Reference for information on the pgx\_rcvall command.

## 14.4 Actions in Response to an Application Error

If there is a connection from a client that has been in the waiting state for an extended period, you can minimize performance degradation of the database by closing the problematic connection.

The following methods are available for identifying a connection to be closed:

- view(pg\_stat\_activity) (refer to "14.4.1 When using the view (pg\_stat\_activity)")
- ps command (refer to "14.4.2 Using the ps Command")
- pgAdmin (refer to "14.4.3 Using pgAdmin")

Use the system management function (pg\_terminate\_backend) to disconnect connections.

## 14.4.1 When using the view (pg\_stat\_activity)

When using the view (pg\_stat\_activity), follow the procedure below to close a connection.

1. Use psql command to connect to the postgres database.

```
> psql postgres
psql (12.1)
Type "help" for help.
```

2. Close connections from clients that have been in the waiting state for an extended period.

Use pg\_terminate\_backend() to close connections that have been trying to connect for an extended period.

However, when considering continued compatibility of applications, do not reference or use system catalogs and functions directly in SQL statements. Refer to "Notes on Application Compatibility" in the Application Development Guide for details.

#### Example

The following example closes connections where the client has been in the waiting state for at least 60 minutes.

```
client_addr | 192.11.11.1
pg_terminate_backend | t
```



- Refer to "System Administration Functions" under "The SQL Language" in the PostgreSQL Documentation for information on pg\_terminate\_backend.
- Refer to "Notes on Application Compatibility" in the Application Development Guide for information on how to maintain application compatibility.

### 14.4.2 Using the ps Command

Follow the procedure below to close a connection using a standard Unix tool (ps command).

1. Execute the ps command.

Note that "<x>" indicates the product version.

```
> ps axwfo user,pid,ppid,tty,command | grep postgres
fsepuser 19174 18027 pts/1
                                      \_ grep postgres
fsepuser 20517
                          /opt/fsepv<x>server64/bin/postgres -D /disk01/data
fsepuser 20518 20517 ?
                           \_ postgres: logger process
fsepuser 20520 20517 ?
                           \_ postgres: checkpointer process
fsepuser 20521 20517 ?
                           \_ postgres: writer process
fsepuser 20522 20517 ?
                           \_ postgres: wal writer process
fsepuser 20523 20517 ?
                           \_ postgres: autovacuum launcher process
fsepuser 20524 20517 ?
                           \_ postgres: archiver process
fsepuser 20525 20517 ?
                           fsepuser 18673 20517 ?
                           \_ postgres: fsepuser db01 192.168.100.11(49449) UPDATE waiting
fsepuser 16643 20517 ?
                          \_ postgres: fsepuser db01 192.168.100.12(49450) idle in transaction
fsepuser 16644 20517 ?
```

Process ID 16643 may be a connection that was established a considerable time ago by the UPDATE statement, or a connection that has occupied resources (waiting).

2. Close connections from clients that have been in the waiting state for an extended period.

Use pg\_terminate\_backend() to close the connection with the process ID identified in step 1 above.

The example below disconnects the process with ID 16643.

However, when considering continued compatibility of applications, do not reference or use system catalogs and functions directly in SQL statements.

```
postgres=# SELECT pg_terminate_backend (16643);
  pg_terminate_backend
-----t
t (1 row)
```



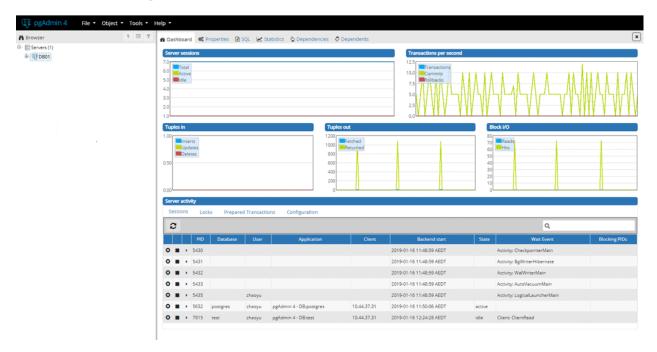
- Refer to "System Administration Functions" under "The SQL Language" in the PostgreSQL Documentation for information on pg\_terminate\_backend.
- Refer to "Notes on Application Compatibility" in the Application Development Guide for information on how to maintain application compatibility.

### 14.4.3 Using pgAdmin

If using pgAdmin, follow the procedure below to close connections.

- 1. In the [Browser] pane, click the relevant database server.
- 2. Close the client connections that have been in a wait state for an extended period.

Click the [Dashboard] tab. In the [Server activity] section, select the connections that have been in an "idle" or "idle in transaction" state for an extended period. For each of these connections, click 3 to close the session.



## 14.5 Actions in Response to an Access Error

If access is denied, grant privileges allowing the instance administrator to operate the following directories, and then re-execute the operation. Also, refer to the event log and the server log, and confirm that the file system has not been mounted as read-only due to a disk error. If the file system has been mounted as read-only, mount it properly and then re-execute the operation.

- Data storage destination
- Tablespace storage destination
- Transaction log storage destination
- Backup data storage destination



Refer to "Preparing Directories to Deploy Resources" under "Setup" in the Installation and Setup Guide for Server for information on the privileges required for the directory.

# 14.6 Actions in Response to Insufficient Space on the Data Storage Destination

If the data storage destination runs out of space, check if the disk contains any unnecessary files and delete them so that operations can continue

If deleting unnecessary files does not solve the problem, you must migrate data to a disk with larger capacity.

There are two methods of migrating data:

- 14.6.1 Using a Tablespace
- 14.6.2 Replacing the Disk with a Larger Capacity Disk

### 14.6.1 Using a Tablespace

FUJITSU Enterprise Postgres enables you to use a tablespace to change the storage destination of database objects, such as tables and indexes, to a different disk.

The procedure is as follows:

1. Create a tablespace

Use the CREATE TABLESPACE command to create a new tablespace in the prepared disk.

2. Modify the tablespace

Use the ALTER TABLE command to modify tables for the newly defined tablespace.



Refer to "SQL Commands" under "Reference" in the PostgreSQL Documentation for information on the CREATE TABLESPACE command and ALTER TABLE command.

### 14.6.2 Replacing the Disk with a Larger Capacity Disk

Before replacing the disk with a larger capacity disk, migrate resources at the data storage destination using the backup and recovery features.

There are two methods of performing backup and recovery:

- 14.6.2.1 Using WebAdmin
- 14.6.2.2 Using Server Commands

The following sections describe procedures that use each of these methods to replace the disk and migrate resources at the data storage destination.



- Before replacing the disk, stop applications and instances that are using the database.
- It is recommended that you back up the database cluster following recovery. Backup deletes obsolete archive logs (transaction logs copied to the backup data storage destination), freeing up disk space and reducing the recovery time.

### 14.6.2.1 Using WebAdmin

Follow the procedure below to replace the disk and migrate resources at the data storage destination by using WebAdmin.

1. Back up files

If the disk at the data storage destination contains any required files, back up the files. It is not necessary to back up the data storage destination.

2. Stop applications

Stop applications that are using the database.

3. Back up the database cluster

Back up the latest resources at the data storage destination. Refer to "3.2.1 Using WebAdmin" for details.

4. Stop the instance

Stop the instance. Refer to "2.1.1 Using WebAdmin" for information on how to stop an instance.

5. Replace with a larger capacity disk

Replace the disk. Then, recover the volume configuration information.

6. Recover the database cluster

Log in to WebAdmin, and perform recovery operations. Refer to steps 4 ("Create a tablespace directory") to 7 ("Run recovery") under "If failure occurred in the data storage disk or the transaction log storage disk" in "14.1.1 Using WebAdmin" for information on the procedure. An instance is automatically started when recovery is successful.

7. Resume applications

Resume applications that are using the database.

8. Restore the files

Restore the files backed up in step 1.

### 14.6.2.2 Using Server Commands

Follow the procedure below to replace the disk and migrate resources at the data storage destination by using server commands.

1. Back up files

If the disk at the data storage destination contains any required files, back up the files. It is not necessary to back up the data storage destination.

2. Stop applications

Stop applications that are using the database.

3. Back up the database cluster

Back up the latest resources at the data storage destination. Refer to "3.2.2 Using Server Commands" for details.

4. Stop the instance

After backup is complete, stop the instance. Refer to "2.1.2 Using Server Commands" for information on how to stop an instance. If the instance fails to stop, refer to "14.11 Actions in Response to Failure to Stop an Instance".

5. Replace with a larger capacity disk

Replace the disk. Then, recover the volume configuration information.

6. Create a data storage destination

Create a data storage destination. If a tablespace was defined, also create a directory for it.

### Example

```
$ mkdir /database/inst1
$ chown fsepuser:fsepuser /database/inst1
$ chmod 700 /database/inst1
```

7. Recover the keystore, and enable automatic opening of the keystore

When the data in the database has been encrypted, restore the keystore to its state at the time of the database backup. Configure automatic opening of the keystore as necessary.

8. Recover the database cluster

Use the pgx\_rcvall command to recover the database cluster.

- Specify the data storage destination in the -D option. If the -D option is omitted, the value of the PGDATA environment variable is used by default.
- Specify the backup storage directory in the -B option.

### Example

pgx\_rcvall -D /database/inst1 -B /backup/inst1



If recovery fails, remove the cause of the error in accordance with the displayed error message and then re-execute the pgx\_rcvall command.

If the message "pgx\_rcvall: an error occurred during recovery" is displayed, then the log recorded when recovery was executed is output after this message. The cause of the error is output in around the last fifteen lines of the log, so remove the cause of the error in accordance with the message and then re-execute the pgx\_rcvall command.

The following message displayed during recovery is output as part of normal operation of pgx\_rcvall (therefore the user does not need not be concerned).

FATAL: The database system is starting



Refer to "pgx\_rcvall" in the Reference for information on the pgx\_rcvall command.

9. Start the instance

Start the instance.

Refer to "2.1.2 Using Server Commands" for information on how to start an instance.



The pgx\_rcvall command cannot accurately recover a hash index. If you are using a hash index, wait for the pgx\_rcvall command to end and then execute the REINDEX command for the appropriate index.

10. Resume applications

Resume applications that are using the database.

11. Restore files

Restore the files backed up in step 1.

# 14.7 Actions in Response to Insufficient Space on the Backup Data Storage Destination

If space runs out on the backup data storage destination, check if the disk contains any unnecessary files and delete them, and then make a backup as required.

If deleting unnecessary files does not solve the problem, take the following action:

- 14.7.1 Temporarily Saving Backup Data
- 14.7.2 Replacing the Disk with a Larger Capacity Disk

# 14.7.1 Temporarily Saving Backup Data

This method involves temporarily moving backup data to a different directory, saving it there, and securing disk space on the backup data storage destination so that a backup can be made normally.

Use this method if you need time to prepare a larger capacity disk.

If space runs out on the backup data storage destination, archive logs can no longer be stored in the backup data storage destination. As a result, transaction logs continue to accumulate in the data storage destination or the transaction log storage destination.

If action is not taken soon, the transaction log storage destination will become full, and operations may not be able to continue.

To prevent this, secure space in the backup data storage destination, so that archive logs can be stored.

There are two methods of taking action:

- 14.7.1.1 Using WebAdmin
- 14.7.1.2 Using Server Commands

### 14.7.1.1 Using WebAdmin

Follow the procedure below to recover the backup data storage disk.

1. Temporarily save backup data

Move backup data to a different directory and temporarily save it, and secure space in the backup data storage destination directory.

The reason for saving the backup data is so that the data in the data storage destination can be recovered even if it is corrupted before you perform recovery. If there is no disk at the save destination and you consider that there is no risk of corruption at the data storage destination, delete the backup data.

The following example saves backup data from the backup data storage destination directory (/backup/inst1) under/mnt/usb/backup.

### Example

```
> mkdir /mnt/usb/backup/
> mv /backup/inst1/* /mnt/usb/backup/
```

2. Back up the database cluster

Back up the latest resources at the data storage destination. Refer to "3.2.1 Using WebAdmin" for details.

3. Delete temporarily saved backup data

If backup completes normally, the temporarily saved backup data becomes unnecessary and is deleted.

The following example deletes backup data that was temporarily saved in /mnt/usb.

### Example

```
> rm -rf /mnt/usb/backup
```

### 14.7.1.2 Using Server Commands

The following describes the procedure for recovering the backup storage disk.

There are two methods of taking action:

- Performing recovery while the instance is active
- Stopping the instance before performing recovery

The following table shows the different steps to be performed depending on whether you stop the instance.

No	Step	Instance stopped	
		No	Yes
1	Stop transaction log mirroring	Y	N
2	Stop output of archive logs	Y	N
3	Stop applications	N	Y
4	Stop the instance	N	Y
5	Temporarily save backup data	Y	Y
6	Resume output of archive logs	Y	N
7	Resume transaction log mirroring	Y	N

No	Step	Instance stopped	
		No	Yes
8	Start an instance	N	Y
9	Run backup	Y	Y
10	Resume applications	N	Y
11	Delete temporarily saved backup data	Y	Y

Y: Required N: Not required

The procedure is as follows:

### Performing recovery while the instance is active

1. Stop transaction log mirroring

Stop transaction log mirroring.

#### 2. Stop output of archive logs

Transaction logs may accumulate during replacement of backup storage disk, and if the data storage disk or the transaction log storage disk becomes full, there is a risk that operations may not be able to continue.

To prevent this, use the following methods to stop output of archive logs.

- Changing the archive\_command parameter

Specify a command that will surely complete normally, such as "echo skipped archiving WAL file %f" or "/bin/true", so that archive logs will be regarded as having been output.

If you specify echo, a message is output to the server log, so it may be used as a reference when you conduct investigations.

- Reloading the configuration file

Run the pg\_ctl reload command or the pg\_reload\_conf SQL function.

If you simply want to stop output of errors without the risk that operations will not be able to continue, specify an empty string (") in archive\_command and reload the configuration file.

### 3. Temporarily save backup data

Move backup data to a different directory and temporarily save it, and secure space in the backup data storage destination directory.

The reason for saving the backup data is so that the data in the data storage destination can be recovered even if it is corrupted before you perform the next step. If there is no disk at the save destination and you consider that there is no risk of corruption at the data storage destination, delete the backup data.

 $The following \ example \ saves \ backup \ data \ from \ the \ backup \ data \ storage \ destination \ directory \ (/backup/inst1) \ under/mnt/usb/backup.$ 

### Example

```
> mkdir /mnt/usb/backup/
> mv /backup/inst1/* /mnt/usb/backup/
```

### 4. Resume output of archive logs

Return the archive\_command setting to its original value, and reload the configuration file.

### 5. Resume transaction log mirroring

Execute the pgx\_resume\_wal\_multiplexing SQL function.

#### Example

```
SELECT pgx_resume_wal_multiplexing()
```

#### 6. Run backup

Use the pgx\_dmpall command to back up the database cluster.

Specify the following option in the pgx\_dmpall command:

- Specify the directory of the data storage destination in the -D option. If the -D option is omitted, the value of the PGDATA environment variable is used by default.

### Example

```
> pgx_dmpall -D /database/instl
```

### 7. Delete temporarily saved backup data

If backup completes normally, the temporarily saved backup data becomes unnecessary and is deleted.

The following example deletes backup data that was temporarily saved in /mnt/usb.

#### Example

```
> rm -rf /mnt/usb/backup
```

### If an instance has been stopped

1. Stop applications

Stop applications that are using the database.

2. Stop the instance

Stop the instance. Refer to "2.1.2 Using Server Commands" for details.

If the instance fails to stop, refer to "14.11 Actions in Response to Failure to Stop an Instance".

3. Temporarily save backup data

Move backup data to a different directory and temporarily save it, and secure space in the backup data storage destination directory.

The reason for saving the backup data is so that the data in the data storage destination can be recovered even if it is corrupted before you perform recovery. If there is no disk at the save destination and you consider that there is no risk of corruption at the data storage destination, delete the backup data.

The following example saves backup data from the backup data storage destination directory (/backup/inst1) under/mnt/usb/backup.

### Example

```
> mkdir /mnt/usb/backup/
> mv /backup/inst1/* /mnt/usb/backup/
```

### 4. Start the instance

Start the instance. Refer to "2.1.2 Using Server Commands" for information on how to start an instance.

5. Run backup

Use the pgx\_dmpall command to back up the database cluster.

Specify the following value in the pgx\_dmpall command:

- Specify the data storage destination in the -D option. If the -D option is omitted, the value of the PGDATA environment variable is used by default.

### Example

> pgx\_dmpall -D /database/inst1

### 6. Resume applications

Resume applications that are using the database.

7. Delete temporarily saved backup data

If backup completes normally, the temporarily saved backup data becomes unnecessary and is deleted.

The following example deletes backup data that was temporarily saved in /mnt/usb.

### Example

rm -rf /mnt/usb/backup



#### See

- Refer to "pgx\_rcvall" and "pgx\_dmpall" in the Reference for information on the pgx\_rcvall command and pgx\_dmpall command.
- Refer to "Write Ahead Log" under "Server Administration" in the PostgreSQL Documentation for information on archive\_command.
- Refer to "B.1 WAL Mirroring Control Functions" for information on the pgx\_is\_wal\_multiplexing\_paused and pgx\_resume\_wal\_multiplexing.

# 14.7.2 Replacing the Disk with a Larger Capacity Disk

This method involves replacing the disk at the backup data storage destination with a larger capacity disk, so that it does not run out of free space again. After replacing the disk, back up data to obtain a proper backup.

There are two methods of performing backup:

- 14.7.2.1 Using WebAdmin
- 14.7.2.2 Using Server Commands



Before replacing the disk, stop applications that are using the database.

### 14.7.2.1 Using WebAdmin

Follow the procedure below to recover the backup storage disk.

1. Back up files

If the disk at the backup data storage destination contains any required files, back up the files. It is not necessary to back up the backup data storage destination.

2. Temporarily save backup data

Save the backup data to a different directory.

The reason for saving the backup data is so that the data in the data storage destination can be recovered even if it is corrupted before you perform the next step. If there is no disk at the save destination and you consider that there is no risk of corruption at the data storage destination, delete the backup data.

The following example saves backup data from the backup data storage destination directory (/backup/inst1) under/mnt/usb/backup.

### Example

```
> mkdir /mnt/usb/backup/
```

> mv /backup/inst1/\* /mnt/usb/backup/

3. Replace with a larger capacity disk

Replace the disk. Then, recover the volume configuration information.

4. Run backup

Log in to WebAdmin, and perform recovery operations. Refer to steps 2 ("Recover the backup data") and 3 ("Run backup") under "If failure occurred on the backup storage disk" in "14.1.1 Using WebAdmin".

5. Restore files

Restore the files backed up in step 1.

6. Delete temporarily saved backup data

If backup completes normally, the temporarily saved backup data becomes unnecessary and is deleted.

The following example deletes backup data that was temporarily saved in /mnt/usb.

### Example

> rm -rf /mnt/usb/backup

### 14.7.2.2 Using Server Commands

The procedure for recovering the backup data storage disk is described below.

There are two methods of taking action:

- Performing recovery while the instance is active
- Stopping the instance before performing recovery

The following table shows the different steps to be performed depending on whether you stop the instance.

No	Step	Instance	stopped
		No	Yes
1	Back up files	Y	Y
2	Temporarily save backup data	Y	Y
3	Confirm that transaction log mirroring has stopped	Y	N
4	Stop output of archive logs	Y	N
5	Stop applications	N	Y
6	Stop the instance	N	Y
7	Replace with a larger capacity disk	Y	Y
8	Create a backup storage directory	Y	Y
9	Resume output of archive logs	Y	N
10	Resume transaction log mirroring	Y	N
11	Start the instance	N	Y
12	Run backup	Y	Y
13	Resume applications	N	Y
14	Restore files	Y	Y
15	Delete temporarily saved backup data	Y	Y

Y: Required

N: Not required

The procedure is as follows:

### If an instance has not been stopped

#### 1. Back up files

If the disk at the backup data storage destination contains any required files, back up the files. It is not necessary to back up the backup data storage destination.

#### 2. Temporarily save backup data

Save the backup data to a different directory.

The reason for saving the backup data is so that the data in the data storage destination can be recovered even if it is corrupted before you perform the next step. If there is no disk at the save destination and you consider that there is no risk of corruption at the data storage destination, delete the backup data.

The following example saves backup data from the backup data storage destination directory (/backup/inst1) under /mnt/usb/backup.

#### Example

```
> mkdir /mnt/usb/backup/
> mv /backup/inst1/* /mnt/usb/backup/
```

#### 3. Confirm that transaction log mirroring has stopped

Use the following SQL function to confirm that transaction log mirroring has stopped.

```
postgres=# SELECT pgx_is_wal_multiplexing_paused();
pgx_is_wal_multiplexing_paused
-----t
t
(1 row)
```

If transaction log mirroring has not stopped, then stop it using the following SQL function.

### 4. Stop output of archive logs

Transaction logs may accumulate during replacement of backup storage disk, and if the data storage destination disk or the transaction log storage destination disk becomes full, there is a risk that operations may not be able to continue.

To prevent this, use the following methods to stop output of archive logs.

- Changing the archive\_command parameter

Specify a command that will surely complete normally, such as "echo skipped archiving WAL file %f" or "/bin/true", so that archive logs will be regarded as having been output.

If you specify echo, a message is output to the server log, so it may be used as a reference when you conduct investigations.

- Reloading the configuration file

Run the pg\_ctl reload command or the pg\_reload\_conf SQL function.

If you simply want to stop output of errors without the risk that operations will not be able to continue, specify an empty string (") in archive\_command and reload the configuration file.

### 5. Replace with a larger capacity disk

Replace the disk. Then, recover the volume configuration information.

### 6. Create a backup data storage destination

Create a backup data storage destination.

### Example

```
# mkdir /backup/inst1
# chown fsepuser:fsepuser /backup/inst1
# chmod 700 /backup/inst1
```

Refer to "3.2.2 Using Server Commands" for details.

### 7. Resume output of archive logs

Return the archive\_command setting to its original value, and reload the configuration file.

8. Resume transaction log mirroring

Execute the pgx\_resume\_wal\_multiplexing SQL function.

#### Example

```
SELECT pgx_resume_wal_multiplexing()
```

#### 9. Run backup

Use the pgx\_dmpall command to back up the database cluster.

Specify the following value in the pgx\_dmpall command:

- Specify the data storage destination in the -D option. If the -D option is omitted, the value of the PGDATA environment variable is used by default.

### Example

```
> pgx_dmpall -D /database/inst1
```

#### 10. Restore files

Restore the files backed up in step 1.

11. Delete temporarily saved backup data

If backup completes normally, the temporarily saved backup data becomes unnecessary and is deleted.

The following example deletes backup data that was temporarily saved in /mnt/usb.

### Example

```
> rm -rf /mnt/usb/backup
```

### If an instance has been stopped

1. Back up files

If the disk at the backup data storage destination contains any required files, back up the files. It is not necessary to back up the backup data storage destination.

2. Temporarily save backup data

Save the backup data to a different directory.

The reason for saving the backup data is so that the data in the data storage destination can be recovered even if it is corrupted before you perform the next step. If there is no disk at the save destination and you consider that there is no risk of corruption at the data storage destination, delete the backup data.

The following example saves backup data from the backup data storage destination directory (/backup/inst1) under /mnt/usb/backup.

### Example

```
> mkdir /mnt/usb/backup/
> mv /backup/inst1/* /mnt/usb/backup/
```

3. Stop applications

Stop applications that are using the database.

4. Stop the instance

Stop the instance. Refer to "2.1.2 Using Server Commands" for information on how to stop an instance.

If the instance fails to stop, refer to "14.11 Actions in Response to Failure to Stop an Instance".

5. Replace with a larger capacity disk

Replace the disk. Then, recover the volume configuration information.

6. Create a backup data storage destination

Create a backup data storage destination.

### Example

```
# mkdir /backup/inst1
# chown fsepuser:fsepuser /backup/inst1
# chmod 700 /backup/inst1
```

Refer to "3.2.2 Using Server Commands" for details.

7. Start the instance

Start the instance. Refer to "2.1.2 Using Server Commands" for information on how to start an instance.

8. Run backup

Use the pgx\_dmpall command to back up the database cluster.

Specify the following value in the pgx\_dmpall command:

- Specify the data storage destination in the -D option. If the -D option is omitted, the value of the PGDATA environment variable is used by default.

### Example

```
> pgx_dmpall -D /database/inst1
```

9. Resume applications

Resume applications that are using the database.

10. Restore files

Restore the files backed up in step 1.

11. Delete temporarily saved backup data

If backup completes normally, the temporarily saved backup data becomes unnecessary and is deleted.

The following example deletes backup data that was temporarily saved in /mnt/usb.

#### Example

```
> rm -rf /mnt/usb/backup
```

......



- Refer to "pgx\_rcvall" and "pgx\_dmpall" in the Reference for information on the pgx\_rcvall command and pgx\_dmpall command.
- Refer to "Write Ahead Log" under "Server Administration" in the PostgreSQL Documentation for information on archive\_command.
- Refer to "B.1 WAL Mirroring Control Functions" for information on the pgx\_is\_wal\_multiplexing\_paused and pgx\_resume\_wal\_multiplexing.

# 14.8 Actions in Response to Insufficient Space on the Transaction Log Storage Destination

If the transaction log storage destination runs out of space, check if the disk contains any unnecessary files and delete them so that operations can continue.

If deleting unnecessary files does not solve the problem, you must migrate data to a disk with larger capacity.

### 14.8.1 Replacing the Disk with a Larger Capacity Disk

Before replacing the disk with a larger capacity disk, migrate resources at the transaction log storage destination using the backup and recovery features.

There are two methods of performing backup and recovery:

- 14.8.1.1 Using WebAdmin
- 14.8.1.2 Using Server Commands

The following sections describe procedures that use each of these methods to replace the disk and migrate resources at the transaction log storage destination.



- Before replacing the disk, stop applications that are using the database.
- It is recommended that you back up the database cluster following recovery. Backup deletes obsolete archive logs (transaction logs copied to the backup data storage destination), freeing up disk space and reducing the recovery time.

### 14.8.1.1 Using WebAdmin

Follow the procedure below to replace the disk and migrate resources at the transaction log storage destination by using WebAdmin.

1. Back up files

If the disk at the transaction log storage destination contains any required files, back up the files. It is not necessary to back up the transaction log storage destination.

2. Back up the database cluster

Back up the latest data storage destination resources and transaction log storage destination resources (refer to "3.2.1 Using WebAdmin" for details).

3. Stop applications

Stop applications that are using the database.

4. Stop the instance

Stop the instance. Refer to "2.1.1 Using WebAdmin" for information on how to stop an instance. WebAdmin automatically stops instances if recovery of the database cluster is performed without stopping the instance.

5. Replace with a larger capacity disk

Replace the disk. Then, recover the volume configuration information.

6. Create a tablespace directory

If a tablespace was defined after backing up, create a directory for it.

7. Recover the keystore, and enable automatic opening of the keystore

Do the following if the data in the database has been encrypted:

- Restore the keystore to its state at the time of the database backup.
- Enable automatic opening of the keystore.

#### 8. Recover the database cluster

Log in to WebAdmin, and perform recovery operations. Refer to steps 4 ("Create a tablespace directory") to 7 ("Run Recovery") under "If failure occurred in the data storage disk or the transaction log storage disk "in "14.1.1 Using WebAdmin" for information on the procedure. An instance is automatically started when recovery is successful.

### 9. Resume applications

Resume applications that are using the database.

### 10. Restore files

Restore the files backed up in step 1.

### 14.8.1.2 Using Server Commands

Follow the procedure below to replace the disk and migrate resources at the transaction log storage destination by using server commands.

### 1. Back up files

If the disk at the transaction log storage destination contains any required files, back up the files. It is not necessary to back up the transaction log storage destination.

### 2. Back up the database cluster

Use server commands to back up the latest data storage destination resources and transaction log storage destination resources. Refer to "3.2.2 Using Server Commands" for information on how to perform backup.

#### 3. Stop applications

Stop applications that are using the database.

### 4. Stop the instance

After backup is complete, stop the instance. Refer to "2.1.2 Using Server Commands" for information on how to stop an instance.

If the instance fails to stop, refer to "14.11 Actions in Response to Failure to Stop an Instance".

### 5. Replace with a larger capacity disk

Replace the disk. Then, recover the volume configuration information.

### 6. Create a transaction log storage destination

Create a transaction log storage destination. If a tablespace was defined, also create a directory for it.

### Example

```
# mkdir /tranlog/inst1
# chown fsepuser:fsepuser /tranlog/inst1
# chmod 700 /tranlog/inst1
```

### 7. Recover the keystore, and enable automatic opening of the keystore

When the data in the database has been encrypted, restore the keystore to its state at the time of the database backup. Configure automatic opening of the keystore as necessary.

### 8. Recover the database cluster

Use the pgx\_rcvall command to recover the database cluster.

- Specify the data storage destination in the -D option. If the -D option is omitted, the value of the PGDATA environment variable is used by default.
- Specify the backup storage directory in the -B option.

#### Example

```
> pgx_rcvall -D /database/inst1 -B /backup/inst1
```



If recovery fails, remove the cause of the error in accordance with the displayed error message and then re-execute the pgx\_rcvall command.

If the message "pgx\_rcvall: an error occurred during recovery" is displayed, then the log recorded when recovery was executed is output after this message. The cause of the error is output in around the last fifteen lines of the log, so remove the cause of the error in accordance with the message and then re-execute the pgx\_rcvall command.

The following message displayed during recovery is output as part of normal operation of pgx\_rcvall command (therefore the user does not need not be concerned).

FATAL: The database system is starting



### See

Refer to "pgx\_rcvall" in the Reference for information on the pgx\_rcvall command.

9. Start the instance

Start the instance.

Refer to "2.1.2 Using Server Commands" for information on how to start an instance.



The pgx\_rcvall command cannot accurately recover a hash index. If you are using a hash index, wait for the instance to start and then execute the REINDEX command for the appropriate index.

10. Resume applications

Resume applications that are using the database.

11. Restore files

Restore the files backed up in step 1.

# 14.9 Errors in More Than One Storage Disk

If an error occurs in the storage destination disks or resources are corrupted, determine the cause of the error from system logs and server logs and remove the cause.

If errors occur in either of the following combinations, you cannot recover the database.

Recreate the instance, and rebuild the runtime environment.

Data storage disk	Transaction log storage disk	Backup data storage disk
Error	-	Error
-	Error	Error



#### See

Refer to "Setup" in the Installation and Setup Guide for Server for information on how to create an instance and build the runtime environment.

# 14.10 Actions in Response to Instance Startup Failure

If an instance fails to start, refer to the system log and the server log, and determine the cause of the failure.

If using WebAdmin, remove the cause of the error. Then, click [Solution] and [Recheck the status] and confirm that the instance is in the normal state.

The following sections describe common causes of errors and the actions to take.

### 14.10.1 Errors in the Configuration File

If you have directly edited the configuration file using a text editor or changed the settings using WebAdmin, refer to the system log and the server log, confirm that no messages relating to the files below have been output.

- postgresql.conf
- pg\_hba.conf



Refer to the following for information on the parameters in the configuration file:

- "Configuring Parameters" in the Installation and Setup Guide for Server
- "Appendix A Parameters"
- "Server Configuration" and "Client Authentication" under "Server Administration" in the PostgreSQL Documentation

### 14.10.2 Errors Caused by Power Failure or Mounting Issues

If mounting is cancelled after restarting the server, for example, because the disk device for each storage destination disk was not turned on, or because automatic mounting has not been set, then starting an instance will fail.

Refer to "14.14.2 Errors Caused by Power Failure or Mounting Issues", and take actions accordingly.

### 14.10.3 Other Errors

This section describes the recovery procedure to be used if you cannot take any action or the instance cannot start even after you have referred to the system log and the server log.

There are two methods of recovery:

- 14.10.3.1 Using WebAdmin
- 14.10.3.2 Using Server Commands

Note that recovery will not be possible if there is an error at the backup data storage destination. If the problem cannot be resolved, contact Fujitsu technical support.

### 14.10.3.1 Using WebAdmin

Follow the procedure below to perform recovery.

- Delete the data storage destination directory and the transaction log storage destination directory
   Back up the data storage destination directory and the transaction log storage destination directory before deleting them.
- 2. Reconfirm the status

Log in to WebAdmin, and in the [Instances] tab, click [Solution] for the error message.

Click [Recheck the status] to reconfirm the storage destination resources.

3. Run recovery

Restore the database cluster after WebAdmin detects an error.

Refer to "14.2.1 Using WebAdmin" for details.

### 14.10.3.2 Using Server Commands

Follow the procedure below to recover the database.

- Delete the data storage destination directory and the transaction log storage destination directory
   Save the data storage destination directory and the transaction log storage destination directory, and then delete them.
- 2. Execute recovery

Use the pgx\_rcvall command to recover the database cluster.

Refer to "14.2.2 Using the pgx\_rcvall Command" for details.

# 14.11 Actions in Response to Failure to Stop an Instance

If an instance fails to stop, refer to the system log and the server log, and determine the cause of the failure.

If the instance cannot stop despite taking action, perform the following operation to stop the instance.

There are two methods of recovery:

- 14.11.1 Using WebAdmin
- 14.11.2 Using Server Commands

### 14.11.1 Using WebAdmin

In the [Instances] tab, click and select the Fast stop mode or the Immediate stop mode to stop the instance. Forcibly terminate the server process from WebAdmin if the instance cannot be stopped.

Refer to "2.1.1 Using WebAdmin" for information on the stop modes.

# 14.11.2 Using Server Commands

There are three methods:

- Stopping the Instance Using the Fast Mode

If backup is in progress, then terminate it, roll back all executing transactions, forcibly close client connections, and then stop the instance.

- Stopping the Instance Using the Immediate Mode

Forcibly terminate the instance immediately. A crash recovery is run when the instance is restarted.

- Forcibly Stopping the Server Process

Reliably stops the server process when the other methods are unsuccessful.

### 14.11.2.1 Stopping the Instance Using the Fast Mode

Specify "-m fast" in the pg\_ctl command to stop the instance.

If the instance fails to stop when you use this method, stop the instance as described in "14.11.2.2 Stopping the Instance Using the Immediate Mode" or "14.11.2.3 Forcibly Stopping the Server Process".



pg\_ctl stop -D /database/instl -m fast

### 14.11.2.2 Stopping the Instance Using the Immediate Mode

Specify "-m immediate " in the pg\_ctl command to stop the instance.

If the instance fails to stop when you use this method, stop the instance as described in "14.11.2.3 Forcibly Stopping the Server Process".



```
> pg_ctl stop -D /database/inst1 -m immediate
```

### 14.11.2.3 Forcibly Stopping the Server Process

If both the Fast mode and the Immediate mode fail to stop the instance, use the kill command or the kill parameter of the pg\_ctl command to forcibly stop the server process.

The procedure is as follows:

1. Execute the ps command

Note that "<*x*>" indicates the product version.

```
> ps axwfo user,pid,ppid,tty,command | grep postgres
fsepuser 19174 18027 pts/1
                                          \_ grep postgres
fsepuser 20517
                1 ?
                             /opt/fsepv<x>server64/bin/postgres -D /database/inst1
fsepuser 20518 20517 ?
                             \_ postgres: logger process
fsepuser 20520 20517 ?
                              \_ postgres: checkpointer process
fsepuser 20521 20517 ?
                              \_ postgres: writer process
fsepuser 20522 20517 ?
                              \_ postgres: wal writer process
fsepuser 20523 20517 ?
                              \_ postgres: autovacuum launcher process
fsepuser 20524 20517 ?
                              \_ postgres: archiver process
fsepuser 20525 20517 ?
                              \_ postgres: stats collector process
```

The process ID (20517) indicates the server process.

2. Forcibly stop the server process

As instance manager, forcibly stop the server process.

Using the pg\_ctl command

```
> pg_ctl kill SIGQUIT 20517
```

Using the kill command

```
> kill -s SIGQUIT 20517
```

# 14.12 Actions in Response to Failure to Create a Streaming Replication Standby Instance

When creating a streaming replication standby instance using WebAdmin, if the instance creation fails, refer to the system log and the server log, and determine the cause of the failure.

When an error occurs in the creation of the standby instance using WebAdmin, it is unlikely that the partially created standby instance can be resumed to complete the operation.

In such a scenario, fix the cause of the error, delete the partially created standby instance, and then create a new standby instance. This recommendation is based on the following assumptions:

- As the instance is yet to be created completely, there are no applications connecting to the database.
- The standby instance is in error state and is not running.
- There are no backups for the standby instance and as a result, it cannot be recovered.



Refer to "Deleting Instances" in the Installation and Setup Guide for details on how to delete an instance.

# 14.13 Actions in Response to Error in a Distributed Transaction

If a system failure (such as server failure) occurs in an application that uses distributed transactions (such as .NET TransactionScope), then transactions may be changed to the in-doubt state. At that point, resources accessed by the transaction will be locked, and rendered unusable by other transactions.

The following describes how to check for in-doubt transactions, and how to resolve them.

### How to check for in-doubt transactions

The following shows how to check for them:

If the server fails

1. An in-doubt transaction will have occurred if a message similar to the one below is output to the log when the server is restarted.

#### Example

```
LOG: Restoring prepared transaction 2103.
```

2. Refer to system view pg\_prepared\_xacts to obtain information about the prepared transaction.

If the transaction identifier of the prepared transaction in the list (in the transaction column of pg\_prepared\_xacts) is the same as the identifier of the in-doubt transaction obtained from the log output when the server was restarted, then that row is the information about the in-doubt transaction.

#### Example

Information about the in-doubt transaction is output to the row with the transaction ID 2103 in the transaction column.

### If the client fails

If there are no clients connected and there is a prepared transaction in pg\_prepared\_xacts, then you can determine that the transaction is in the in-doubt state.

If at least one client is connected and there is a prepared transaction in pg\_prepared\_xacts, you cannot determine whether there is a transaction in the in-doubt state. In this case, use the following query to determine the in-doubt transaction from the acquired database name, user name, the time PREPARE TRANSACTION was executed, and the information about the table name accessed.

```
select gid,x.database,owner,prepared,l.relation::regclass as relation from pg_prepared_xacts x left join pg_locks l on l.virtualtransaction = '-1/'||x.transaction and l.locktype='relation';
```

If it still cannot be determined from this information, wait a few moments and then check pg\_prepared\_xacts again.

If there is a transaction that has continued since the last time you checked, then it is likely that it is the one in the in-doubt state.



As you can see from the explanations in this section, there is no one way to definitively determine in-doubt transactions.

Consider collecting other supplementary information (for example, logging on the client) or performing other operations (for example, allocating database users per job).

#### How to resolve in-doubt transactions

From the system view pg\_prepared\_xacts mentioned above, obtain the global transaction identifier (in the gid column of pg\_prepared\_xacts) for the in-doubt transaction, and issue either a ROLLBACK PREPARED statement or COMMIT PREPARED statement to resolve the in-doubt transaction.



- Rolling back in-doubt transactions

```
postgres=# rollback prepared '374cc221-f6dc-4b73-9d62-d4fec9b430cd';
ROLLBACK PREPARED
```

- Committing in-doubt transactions

```
postgres=# commit prepared '374cc221-f6dc-4b73-9d62-d4fec9b430cd';
COMMIT PREPARED
```

### 14.14 I/O Errors Other than Disk Failure

Even if a disk is not defective, the same input-output error messages, as those generated when the disk is defective, may be output.

A few examples of such errors are given below. The appropriate action for each error is explained respectively.

- 14.14.1 Network Error with an External Disk
- 14.14.2 Errors Caused by Power Failure or Mounting Issues

### 14.14.1 Network Error with an External Disk

This is an error that occurs in the network path to/from an external disk.

Determine the cause of the error by checking the information in the system log and the server log, the disk access LED, network wiring, and network card status. Take appropriate action to remove the cause of the error, for example, replace problematic devices.

# 14.14.2 Errors Caused by Power Failure or Mounting Issues

These are errors that occur when the disk device is not turned on, automatic mounting of the disk was not set, or mounting was accidentally cancelled.

In this case, check the information in the system log and the server log, the disk access LED, and whether the disk is mounted correctly. If problems are detected, take appropriate action.

If mounting has been cancelled, it is possible that mounting was accidentally cancelled, or automatic mounting at the time of starting the operating system is not set. In this case, set the mounting to be performed automatically.

# 14.15 Anomaly Detection and Resolution

The following operations performed via the command line interface will result in an anomaly in WebAdmin:

- Changes to the port and backup\_destination parameters in postgresql.conf
- Changes to Mirroring Controller configuration of cluster replication added via WebAdmin

This section describes when WebAdmin checks for such anomalies, and what takes place when an anomaly is detected.

# 14.15.1 Port Number and Backup Storage Path Anomalies

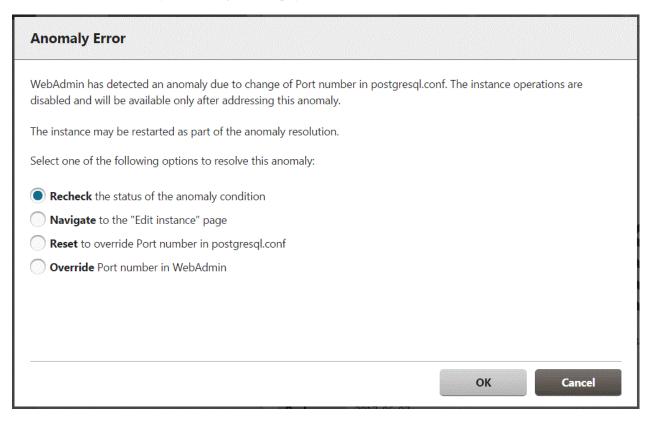
An anomaly occurs when the value of [Port number] and/or [Backup storage path] in WebAdmin is different from the value of its corresponding parameter in postgresql.conf - port and backup\_destination, respectively.

WebAdmin checks for anomalies when an instance is selected for viewing or any instance operation is performed. Anomalies will be identified for the selected instance only.

The following occurs when an anomaly is detected in port number and/or backup storage path:

- All instance operation buttons are disabled, except for "Edit instance", "Refresh instance", and "Delete Mirroring Controller"
- A red error status indicator is displayed on the instance icon
- For an anomaly specific to backup storage path, a red error status indicator is displayed on the [Backup storage] disk icon, and [Backup storage status] is set to "Error"
- The message, "WebAdmin has detected an anomaly with...", is displayed in the [Message] section along with an associated [Solution] button

Click [Solution]. The [Anomaly Error] dialog box is displayed.



Select the required option, click [OK], and then resolve the anomaly error.

Refer to "Editing instance information" in the Installation and Setup Guide for Server for information on the [Edit instance] page.



Critical errors encountered during anomaly resolution will be displayed, however, rollback of the instance to its previous state is not supported.

# 14.15.2 Mirroring Controller Anomalies

The following conditions will cause a Mirroring Controller anomaly:

- The Mirroring Controller management folder or configuration files have been deleted
- The permissions to the Mirroring Controller management folder or configuration files have been changed such that:
  - The instance administrator's access to Mirroring Controller configuration is denied
  - Users other than an instance administrator have access privileges to Mirroring Controller configuration files

WebAdmin checks for anomalies when Mirroring Controller status check is performed.

The following occurs when a Mirroring Controller anomaly is detected:

- All Mirroring Controller functionality is disabled for the replication cluster, except for "Delete Mirroring Controller"
- [Mirroring Controller status] is set to "Error"
- Either of the following messages is displayed in the [Message] section

"Failed to access the Mirroring Controller management folder or configuration files 'path'. Mirroring Controller functionality has been disabled. Consider deleting Mirroring Controller and adding it again."

"Failed to find the Mirroring Controller management folder or configuration files 'path'. Mirroring Controller functionality has been disabled. Consider deleting Mirroring Controller and adding it again."

# Appendix A Parameters

This appendix describes the parameters to be set in the postgresql.conf file of FUJITSU Enterprise Postgres.

The postgresql.conf file is located in the data storage destination.

- core\_directory (string)

This parameter specifies the directory where the corefile is to be output. If this parameter is omitted, the data storage destination is used by default. This parameter can only be set when specified on starting an instance. It cannot be changed dynamically, while an instance is active.

- core\_contents (string)

This parameter specifies the contents to be included in the corefile.

- full: Outputs all contents of the server process memory to the corefile.
- none: Does not output a corefile.
- minimum: Outputs only non-shared memory server processes to the corefile. This reduces the size of the corefile. However, in some cases, this file may not contain sufficient information for examining the factor that caused the corefile to be output.

If this parameter is omitted, "minimum" is used by default. This parameter can only be set when specified on starting an instance. It cannot be changed dynamically, while an instance is active.

- keystore\_location (string)

This parameter specifies the directory that stores the keystore file. Specify a different location from other database clusters. This parameter can only be set when specified on starting an instance. It cannot be changed dynamically, while an instance is active.

- tablespace\_encryption\_algorithm (string)

This parameter specifies the encryption algorithm for tablespaces that will be created. Valid values are "AES128", "AES256", and "none". If you specify "none", encryption is not performed. The default value is "none". To perform encryption, it is recommended that you specify "AES256". Only superusers can change this setting.

- backup\_destination (string)

This parameter specifies the absolute path of the directory where pgx\_dmpall will store the backup data. Specify a different location from other database clusters. This parameter can only be set when specified on starting an instance. It cannot be changed dynamically, while an instance is active.

Place this directory on a different disk from the data directory to be backed up and the tablespace directory. Ensure that users do not store arbitrary files in this directory, because the contents of this directory are managed by the database system.

- search\_path (string)

When using the SUBSTR function compatible with Oracle databases, set "oracle" and "pg\_catalog" in the search\_path parameter. You must specify "oracle" before "pg\_catalog".



search\_path = '"\$user", public, oracle, pg\_catalog'



- The search\_path feature specifies the priority of the schema search path. The SUBSTR function in Oracle database is defined in the oracle schema.

- Refer to "Statement Behavior" under "Server Administration" in the PostgreSQL Documentation for information on search\_path.
- track\_waits (string)

This parameter enables collection of statistics for pgx\_stat\_lwlock and pgx\_stat\_latch.

- on: Enables collection of statistics.
- off: Disables collection of statistics.

If this parameter is omitted, "on" is assumed.

Only superusers can change this setting.

- track\_sql (string)

This parameter enables collection of statistics for pgx\_stat\_sql.

- on: Enables collection of statistics.
- off: Disables collection of statistics.

If this parameter is omitted, "on" is assumed.

Only superusers can change this setting.

### Parameters for the in-memory feature

- reserve\_buffer\_ratio (numerical value)

This parameter specifies the proportion of shared memory to be used for a stable buffer table.

- Minimum value: 0

- Maximum value: 80

If this parameter is omitted, 0 will be used.

- vci.cost\_threshold (numerical value)

This parameter specifies the lowest cost that selects an execution plan that uses a VCI. If the cost of the best execution plan that does not use a VCI is lower than this value, that execution plan will be selected.

- Minimum value: 0
- Maximum value: Maximum value that can be expressed as a 4-byte signed integer

If this parameter is omitted or a value outside this range is specified, 18000 will be used.

- vci.control\_max\_workers (numerical value)

This parameter specifies the number of background workers that manage VCI. The number of workers for the entire instance is limited by max\_worker\_processes, so add the value specified here to max\_worker\_processes.

- Minimum value: 1
- Maximum value: 8388607

If this parameter is omitted or a value outside this range is specified, 8 will be used.

- vci.enable (string)

This parameter enables or disables VCI.

- on: Enables VCI.
- off: Disables VCI.

If this parameter is omitted, "on" will be used.

- vci.log\_query (string)

This parameter enables or disables log output when VCI is not used due to insufficient memory specified by vci.max local ros.

- on: Enables log output.
- off: Disables log output.

If this parameter is omitted, "off" will be used.

- vci.maintenance\_work\_mem (numerical value)

This parameter specifies the maximum memory size used for maintenance of VCI (when executing CREATE INDEX, for example).

- Minimum value: 1 MB
- Maximum value: Maximum value that can be expressed as a 4-byte signed integer

If this parameter is omitted or a value outside this range is specified, 256 MB will be used.

- vci.max\_local\_ros (numerical value)

This parameter specifies the maximum memory size used for VCI scan.

- Minimum value: 64 MB
- Maximum value: Maximum value that can be expressed as a 4-byte signed integer

If this parameter is omitted or a value outside this range is specified, 64 MB will be used.



The maximum value that can be expressed as a 4-byte signed integer changes according to the operating system. Follow the definition of the operating system in use.

vci.max\_parallel\_degree (numerical value)

This parameter specifies the maximum number of background workers used for parallel scan. The number of workers for the entire instance is limited by max\_worker\_processes, so add the value specified here to max\_worker\_processes.

A value from -8388607 to 8388607 can be specified.

- Integer (1 or greater): Parallel scan is performed using the specified degree of parallelism.
- 0: Stops the parallel scan process.
- Negative number: The specified value minus the maximum number of CPUs obtained from the environment is used as the degree of parallelism and parallel scan is performed.

If this parameter is omitted or a value outside this range is specified, 0 will be used.

- vci.shared\_work\_mem (numerical value)

This parameter specifies the maximum memory size used for VCI parallel scan.

- Minimum value: 32 MB
- Maximum value: Maximum value that can be expressed as a 4-byte signed integer

If this parameter is omitted or a value outside this range is specified, 1 GB will be used.

#### Parameters for the Global Meta Cache

- pgx\_global\_metacache (numerical value)

Specifies the memory size of the GMC area.

Specify a value calculated by the formula below.

A value lower than the calculated value will still work, but the meta cache may not be able to fit into the GMC area.

In this case, the system will discard the meta cache it thinks it is no longer needed, but if it is needed again, the meta cache will need to be expanded and will not perform well.

If the value is less than 10 MB and is set to a nonzero value that disables the feature, the database startup fails because the Global Meta Cache feature cannot operate.

A setting of 0 disables the Global Meta Cache feature. The default is 0.

Changing this setting requires restarting the database.

This value takes into account the case where both GMC before and after the change temporarily exist at the same time in shared memory when the table definition is changed or the row of the system catalog is changed.

- track\_gmc (string)

This parameter enables collection of statistics for pgx\_stat\_gmc.

- on: Enables collection of statistics.

\*1) Safety Factor (1.5)

- off: Disables collection of statistics.

If this parameter is omitted, "on" is used.

Only superusers can change this setting.



#### See

Refer to "Server Configuration" under "Server Administration" in the PostgreSQL Documentation for information on other postgresql.conf parameters.

.....

# Appendix B System Administration Functions

This appendix describes the system administration functions of FUJITSU Enterprise Postgres.



Refer to "System Administration Functions" under "The SQL Language" in the PostgreSQL Documentation for information on other system administration functions.

# **B.1 WAL Mirroring Control Functions**

The following table lists the functions that can be used for backup and recovery based on WAL mirroring.

Table B.1 WAL mirroring control functions

Name	Return type	Description
pgx_pause_wal_multiplexing()	void	Stops WAL multiplexing
pgx_resume_wal_multiplexing()	void	Resumes WAL multiplexing
pgx_is_wal_multiplexing_paused()	boolean	Returns true if WAL multiplexing has stopped

If WAL multiplexing has not been configured, these functions return an error. Setting the backup\_destination parameter in postgresql.conf configures WAL multiplexing.

Only superusers can execute these functions.

# **B.2 Transparent Data Encryption Control Functions**

The following table lists the functions that can be used for transparent data encryption.

Table B.2 Transparent data encryption control functions

Name	Return type	Description
pgx_open_keystore(passphrase)	void	Opens the keystore
pgx_set_master_key(passphrase)	void	Sets the master encryption key
pgx_set_keystore_passphrase(oldPassphrase, newPassphrase)	void	Changes the keystore passphrase

The pgx\_open\_keystore function uses the specified passphrase to open the keystore. When the keystore is opened, the master encryption key is loaded into the database server memory. In this way, you can access the encrypted data and create encrypted tablespaces. If the keystore is already open, this function returns an error.

Only superusers can execute this function. Also, this function cannot be executed within a transaction block.

The pgx\_set\_master\_key function generates a master encryption key and stores it in the keystore. If the keystore does not exist, this function creates a keystore. If the keystore already exists, this function modifies the master encryption key. If the keystore has not been opened, this function opens it.

The passphrase is a string of 8 to 200 bytes.

Only superusers can execute this function. Also, this function cannot be executed within a transaction block. Processing is not affected by whether the keystore is open.

The pgx\_set\_keystore\_passphrase function changes the keystore passphrase. Specify the current passphrase in *oldPassphrase*, and a new passphrase in *newPassphrase*.

The passphrase is a string of 8 to 200 bytes.

Only superusers can execute this function. Also, this function cannot be executed within a transaction block. Processing is not affected by whether the keystore is open.

# **B.3 Data Masking Control Functions**

The table below lists the functions that can be used for data masking.

Table B.3 Data masking control functions

Name	Return type	Description
pgx_alter_confidential_policy	boolean	Changes masking policies
pgx_create_confidential_policy	boolean	Creates masking policies
pgx_drop_confidential_policy	boolean	Deletes masking policies
pgx_enable_confidential_policy	boolean	Enables or disables masking policies
pgx_update_confidential_values	boolean	Changes replacement characters when full masking is specified for masking type

### B.3.1 pgx\_alter\_confidential\_policy

### Description

Changes masking policies

#### **Format**

The format varies depending on the content to be changed. The format is shown below.

- Common format

- Add a masking target to a masking policy

```
pgx_alter_confidential_policy(
commonArg,
[action
            := 'ADD_COLUMN', ]
column_name := 'colName'
[, function_type := 'FULL']
[, function_type := 'PARTIAL', partialOpt] |
[, function_type := 'REGEXP', regexpOpt]
partialOpt:
function_parameters := 'maskingFmt'
regexpOpt:
regexp_pattern
                   := 'regexpPattern',
regexp_replacement := 'regexpReplacementChar',
[, regexp_flags
                    := 'regexpFlags']
```

- Delete a masking target from a masking policy

- Change the masking condition

- Change the content of a masking policy set for a masking target

- Change the masking policy description

- Change the masking target description

```
pgx_alter_confidential_policy(
commonArg,
action := 'SET_COLUMN_DESCRIPTION',
column_name := 'colName',
column_description := 'colDesc'
)
```

### **Argument**

The argument varies depending on the content to be changed. Details are as follows.

- Common arguments

Masking type for which an argument can be specified	Argument	Data type	Description	Default value
All	schema_name	varchar(63)	Schema name of table for which a masking policy is applied	'public'
	table_name	varchar(63)	Name of table for which a masking policy is applied	Mandatory
	policy_name	varchar(63)	Masking policy name	Mandatory

### - Add a masking target to a masking policy

Masking type for which an argument can be specified	Argument	Data type	Description	Default value
All	action	varchar(63)	'ADD_COLUMN'	'ADD_COLU MN'
	column_name	varchar(63)	Masking target name	Mandatory
	function_type	varchar(63)	Masking type	'FULL'
			- 'FULL': Full masking	
			- 'PARTIAL': Partial masking	
			- 'REGEXP': Regular expression masking	
Partial masking	function_parameters	varchar(1024)	Masking format for partial masking	Mandatory
Regular expression	regexp_pattern	varchar(1024)	Search pattern for regular expression masking	Mandatory
masking	regexp_replacement	varchar(1024)	Replacement character/string for regular expression masking	Mandatory
	regexp_flags	varchar(20)	Regular expression flags	NULL

### - Delete a masking target from a masking policy

Masking type for which an argument can be specified	Argument	Data type	Description	Default value
All	action	varchar(63)	'DROP_COLUMN'	Mandatory
	column_name	varchar(63)	Masking target name	Mandatory

### - Change the masking condition

Masking type for which an argument can be specified	Argument	Data type	Description	Default value
All	action	varchar(63)	'MODIFY_EXPRESSION'	Mandatory
	expression	varchar(1024)	Masking condition to be changed	Mandatory

- Change the content of a masking policy set for a masking target

Masking type for which an argument can be specified	Argument	Data type	Description	Default value
All	action	varchar(63)	'MODIFY_COLUMN'	Mandatory
	column_name	varchar(63)	Masking target name	Mandatory
	function_type	varchar(63)	Masking type	'FULL'
			- 'FULL': Full masking	
			- 'PARTIAL': Partial masking	
			- 'REGEXP': Regular expression masking	
Partial masking	function_parameters	varchar(1024)	Masking format for partial masking	Mandatory
Regular expression	regexp_pattern	varchar(1024)	Search pattern for regular expression masking	Mandatory
masking	regexp_replacement	varchar(1024)	Replacement character/string for regular expression masking	Mandatory
	regexp_flags	varchar(20)	Regular expression flags	NULL

### - Change the masking policy description

Masking type for which an argument can be specified	Argument	Data type	Description	Default value
All	action	varchar(63)	'SET_POLICY_DESCRIPTION'	Mandatory
	policy_description	varchar(1024)	Masking policy description	Mandatory

### - Change the masking target description

Masking type for which an argument can be specified	Argument	Data type	Description	Default value
All	action	varchar(63)	'SET_COLUMN_DESCRIPTION'	Mandatory
	column_name	varchar(63)	Masking target name	Mandatory
	column_description	varchar(1024)	Masking target description	Mandatory

Details about whether arguments can be omitted are as follows.

Argument		Mandatory or optional								
	ADI	D_COL	UMN	DROP_C	MODIFY_COLUM			SET_POLICY	SET_COLUMN	
			OLUMN _EXF		_EXPRE		N			_DESCRIPTIO
	Full	Partia	Regul		SSION	Full	Partia	Regul	ON	N
	mas	- 1	ar			mas	- 1	ar		
	king	maski	expre			king	maski	expre		
		ng	ssion				ng	ssion		
			maski					maski		
			ng					ng		
schema_name	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
table_name	N	N	N	N	N	N	N	N	N	N

Argument		Mandatory or optional								
	ADD_COLUMN		DROP_C OLUMN	MODIFY _EXPRE	MOD	MODIFY_COLUM N		SET_POLICY _DESCRIPTI	SET_COLUMN _DESCRIPTIO	
	Full mas king	Partia I maski ng	Regul ar expre ssion maski ng		SSION	Full mas king	I	Regul ar expre ssion maski ng	ON	N
policy_name	N	N	N	N	N	N	N	N	N	N
action	Y	Y	Y	N	N	N	N	N	N	N
column_name	N	N	N	N	-	N	N	N	-	N
function_type	Y	N	N	-	-	Y	N	N	-	-
expression	-	-	-	-	N	-	-	-	-	-
policy_description	-	-	-	-	-	-	-	-	N	-
column_description	-	-	-	-	-	-	-	-	-	N
function_parameters	-	N	-	-	-	-	N	-	-	-
regexp_pattern	-	-	N	-	-	-	-	N	-	-
regexp_replacement	-	-	N	-	-	-	-	N	-	-
regexp_flags	-	-	Y	-	-	-	-	Y	-	-

Y: Can be omitted; N: Cannot be omitted; -: Ignored when specified

### Return value

Return value	Description
TRUE	Ended normally
FALSE	Ended abnormally

### **Execution example 1**

Adding masking policy p1 to masking target c2

### **Execution example 2**

Deleting masking target c1 from masking policy p1

```
postgres=# select pgx_alter_confidential_policy(table_name := 't1', policy_name := 'p1', action :=
'DROP_COLUMN', column_name := 'c1');
pgx_alter_confidential_policy
-----t
t
(1 row)
```

### **Execution example 3**

Changing the masking condition for masking policy p1

```
postgres=# select pgx_alter_confidential_policy(table_name := 't1', policy_name := 'p1', action :=
'MODIFY_EXPRESSION', expression := 'false');
pgx_alter_confidential_policy
-----t
(1 row)
```

### **Execution example 4**

Changing the content of masking policy p1 set for masking target c2

```
postgres=# select pgx_alter_confidential_policy(table_name := 't1', policy_name := 'p1', action :=
'MODIFY_COLUMN', column_name := 'c2', function_type := 'FULL');
pgx_alter_confidential_policy
------t
(1 row)
```

### **Execution example 5**

Changing the description of masking policy p1

### **Execution example 6**

Changing the description of masking target c2

```
postgres=# select pgx_alter_confidential_policy(table_name := 't1', policy_name := 'p1', action :=
'SET_COLUMN_DESCRIPTION', column_name := 'c2', column_description := 'c2 column is FULL.');
pgx_alter_confidential_policy
------t
t
(1 row)
```

### Description

- The arguments for the pgx\_alter\_confidential\_policy system management function can be specified in any order.
- The action parameters below can be specified. When action parameters are omitted, ADD\_COLUMN is applied.

Parameter	Description
ADD_COLUMN	Adds a masking target to a masking policy.
DROP_COLUMN	Deletes a masking target from a masking policy.
MODIFY_EXPRESSION	Changes expression.
MODIFY_COLUMN	Changes the content of a masking policy set for a masking target.
SET_POLICY_DESCRIPTION	Changes policy_description.
SET_COLUMN_DESCRIPTION	Changes column_description.

- The function\_parameters argument is enabled when the function\_type is PARTIAL. If the function\_type is other than PARTIAL, it will be ignored.

- The arguments below are enabled when the function\_type is REGEXP. If the function\_type is other than REGEXP, these arguments will be ignored.
  - regexp\_pattern
  - regexp\_replacement
  - regexp\_flags



### See

- Refer to "String Constants" in the PostgreSQL Documentation for information on the strings to specify for arguments.
- Refer to "POSIX Regular Expressions" in the PostgreSQL Documentation and check pattern, replacement, and flags for information on the values that can be specified for regexp\_pattern, regexp\_replacement, and regexp\_flags.

# B.3.2 pgx\_create\_confidential\_policy

### Description

Creates masking policies

#### **Format**

The format varies depending on the masking type. The format is shown below.

```
pgx_create_confidential_policy(
[schema_name := 'schemaName',]
                 := 'tableName',
table_name
table_name := 'table
policy_name := 'pol
expression := 'expr'
[. enable := 'po
                    := 'policyName',
[, enable
                   := 'policyStatus']
[, policy_description
                          := 'policyDesc']
[, column_name := 'colName'
                         := 'FULL']
     [, function_type
     [, function_type := 'PARTIAL', partialOpt] |
     [, function_type := 'REGEXP', regexpOpt]
                               := 'colDesc']
     [, column_description
])
partialOpt:
function_parameters := 'maskingFmt'
regexpOpt:
                := 'regexpPattern',
regexp_pattern
regexp_replacement := 'regexpRepl
[, regexp_flags := 'regexpFlags']
                           := 'regexpReplacementChar',
```

### **Argument**

Details are as follows.

Masking type for which an argument can be specified	Argument	Data type	Description	Default value
All	schema_name	varchar(63)	Schema name of table for which the masking policy is created	'public'
	table_name	varchar(63)	Name of table for which the masking policy is created	Mandatory

Masking type for which an argument can be specified	Argument	Data type	Description	Default value
	policy_name	varchar(63)	Masking policy name	Mandatory
	expression	varchar(1024)	Masking condition	Mandatory
	enable	boolean	Masking policy status	't'
			- 't': Enabled	
			- 'f': Disabled	
	policy_description	varchar(1024)	Masking policy description	NULL
	column_name	varchar(63)	Masking target name	NULL
	function_type	varchar(63)	Masking type	'FULL'
			- 'FULL': Full masking	
			- 'PARTIAL': Partial masking	
			- 'REGEXP': Regular expression masking	
	column_description	varchar(1024)	Masking target description	NULL
Partial masking	function_parameters	varchar(1024)	Masking format for partial masking	Mandatory
Regular expression masking	regexp_pattern	varchar(1024)	Search pattern for regular expression masking	Mandatory
	regexp_replacement	varchar(1024)	Replacement character/string for regular expression masking	Mandatory
	regexp_flags	varchar(20)	Regular expression flags	NULL

Details about whether arguments can be omitted are as follows.

Argument		Mandatory or op	tional
	Full masking	Partial masking	Regular expression masking
schema_name	Y	Y	Y
table_name	N	N	N
policy_name	N	N	N
expression	N	N	N
enable	Y	Y	Y
policy_description	Y	Y	Y
column_name	Y	Y	Y
function_type	Y	Y	Y
column_description	Y	Y	Y
function_parameters	-	N	-
regexp_pattern	-	-	N
regexp_replacement	-	-	N
regexp_flags	-	-	Y

Y: Can be omitted; N: Cannot be omitted; -: Ignored when specified

#### Return value

Return value	Description		
TRUE	Ended normally		
FALSE	Ended abnormally		

### **Execution example 1**

Creating masking policy p1 that does not contain a masking target

### **Execution example 2**

Creating masking policy p1 that contains masking target c1 of which the masking type is full masking

### **Execution example 3**

Creating masking policy p1 that contains masking target c2 of which the masking type is partial masking

### **Execution example 4**

Creating masking policy p1 that contains masking target c3 of which the masking type is regular expression masking

```
postgres=# select pgx_create_confidential_policy( table_name := 't1', policy_name := 'p1',
expression := '1=1', column_name := 'c3', function_type := 'REGEXP', regexp_pattern := '(.*)(@.*)',
regexp_replacement := 'xxx\2', regexp_flags := 'g');
pgx_create_confidential_policy
------t
(1 row)
```

### Description

- The arguments for the pgx\_create\_confidential\_policy system management function can be specified in any order.
- If column name is omitted, only masking policies that do not contain masking target will be created.
- One masking policy can be created for each table. Use the pgx\_alter\_confidential\_policy system management function to add a masking target to a masking policy.

- The function\_parameters argument is enabled when the function\_type is PARTIAL. If the function\_type is other than PARTIAL, it will be ignored.
- The arguments below are enabled when the function\_type is REGEXP. If the function\_type is other than REGEXP, these arguments will be ignored.

......

- regexp\_pattern
- regexp\_replacement
- regexp\_flags



If a table for which a masking policy is to be applied is deleted, delete the masking policy as well.



### See

- Refer to "String Constants" in the PostgreSQL Documentation for information on the strings to specify for arguments.
- Refer to "POSIX Regular Expressions" in the PostgreSQL Documentation and check pattern, replacement, and flags for information on the values that can be specified for regexp\_pattern, regexp\_replacement, and regexp\_flags.

### B.3.3 pgx\_drop\_confidential\_policy

### Description

Deletes masking policies

### **Format**

```
pgx_drop_confidential_policy(
[schema_name := 'schemaName', ]
table_name := 'tableName',
policy_name := 'policyName'
)
```

### **Argument**

Details are as follows.

Argument	Data type	Description	Default value
schema_name	varchar(63)	Schema name of table for which a masking policy is deleted	'public'
table_name	varchar(63)	Name of table for which a masking policy is deleted	Mandatory
policy_name	varchar(63)	Masking policy name	Mandatory

Details about whether arguments can be omitted are as follows.

Argument	Mandatory or optional
schema_name	Y
table_name	N
policy_name	N

Y: Can be omitted; N: Cannot be omitted

### Return value

Return value	Description
TRUE	Ended normally
FALSE	Ended abnormally

### **Execution example**

Deleting masking policy p1

```
postgres=# select pgx_drop_confidential_policy(table_name := 't1', policy_name := 'p1');
    pgx_drop_confidential_policy
    ------
    t
    (1 row)
```

### Description

The arguments for the pgx\_drop\_confidential\_policy system management function can be specified in any order.



If a table for which a masking policy is to be applied is deleted, delete the masking policy as well.



Refer to "String Constants" in the PostgreSQL Documentation for information on the strings to specify for arguments.

# B.3.4 pgx\_enable\_confidential\_policy

### Description

Enables or disables masking policies

#### **Format**

### Argument

Details are as follows.

Argument	Data type	Description	Default value
schema_name	varchar(63)	Schema name of table for which a masking policy is enabled or disabled	'public'
table_name	varchar(63)	Name of table for which a masking policy is enabled or disabled	Mandatory
policy_name	varchar(63)	Masking policy name	Mandatory
enable	boolean	Masking policy status	Mandatory
		- 't': Enabled	

Argument	Data type	Description	Default value
		- 'f': Disabled	

Details about whether arguments can be omitted are as follows.

Argument	Mandatory or optional
schema_name	Y
table_name	N
policy_name	N
enable	N

Y: Can be omitted; N: Cannot be omitted

#### Return value

Return value	Description
TRUE	Ended normally
FALSE	Ended abnormally

#### **Execution example**

Enabling masking policy p1

```
postgres=# select pgx_enable_confidential_policy(table_name := 't1', policy_name := 'p1', enable :=
't');
pgx_enable_confidential_policy
-----t
(1 row)
```

......

#### Description

The arguments for the pgx\_enable\_confidential\_policy system management function can be specified in any order.



#### See

Refer to "String Constants" in the PostgreSQL Documentation for information on the strings to specify for arguments.

## B.3.5 pgx\_update\_confidential\_values

#### Description

Changes replacement characters when full masking is specified for masking type

#### **Format**

```
pgx_update_confidential_values(
[number_value := 'numberValue']
[, char_value := 'charValue']
[, varchar_value := 'varcharValue']
[, date_value := 'dateValue']
[, ts_value := 'tsValue']
)
```

#### **Argument**

Details are as follows.

Argument	Data type	Description	
number_value	integer	Replacement character in numeric type	
char_value	varchar(1)	Replacement character in char type	
varchar_value	varchar(1)	Replacement character in varchar type	
date_value	date	Replacement character in date type	
ts_value	timestamp	Replacement character in timestamp type	

#### Return value

Return value	Description
TRUE	Ended normally
FALSE	Ended abnormally

#### **Execution example**

Using '\*' as a replacement character in char type and varchar type

```
postgres=# select pgx_update_confidential_values(char_value := '*', varchar_value := '*');
    pgx_update_confidential_values
    ------
    t
    (1 row)
```

#### Description

- The arguments for the pgx\_update\_confidential\_values system management function can be specified in any order.
- Specify one or more arguments for the pgx\_update\_confidential\_values system management function. A replacement character is not changed for an omitted argument.



Refer to "String Constants" in the PostgreSQL Documentation for information on the strings to specify for arguments.

#### **B.4 VCI Data Load Control Function**

The table below lists the function that loads VCI data to buffer cache.

Table B.4 VCI data load control function

Name	Return type	Description
pgx_prewarm_vci(vci_index regclass)	int8	Loads the VCI data to buffer cache.

pgx\_prewarm\_vci loads the specified VCI data to buffer cache and returns the number of blocks of the loaded VCI data.

The aggregation process using VCI may take time immediately after an instance is started, because the VCI data has not been loaded to buffer cache. Therefore, the first aggregation process can be sped up by executing pgx\_prewarm\_vci after an instance is started.

The amount of memory required for preloading is the number of blocks returned by pgx\_prewarm\_vci multiplied by the size of one block.

This function can only be executed if the user has reference privilege to the VCI index and execution privilege to the pg\_prewarm function.

## **B.5 High-Speed Data Load Control Functions**

The table below lists the functions that can be used for high-speed data load.

Table B.5 High-speed data load control functions

Name	Return type	Description
pgx_loader	bigint	Creates dynamic shared memory, starts parallel workers and loads data
pgx_loader_recovery	smallint	Resolves in-doubt transactions

The pgx\_loader command executes the above functions internally.

## Appendix C System Views

This appendix describes how to use the system views in FUJITSU Enterprise Postgres.



Refer to "System Views" under "Internals" in the PostgreSQL Documentation for information on other system views.

### C.1 pgx\_tablespaces

The pgx\_tablespaces catalog provides information related to the encryption of tablespaces.

Column	Туре	References	Description
spctablespace	oid	pg_tablespace.oid	Tablespace OID
spcencalgo	text		Tablespace encryption algorithm

The spcencalgo string displays one of the following values:

- none: Tablespace is not encrypted
- AES128: AES with key length of 128 bits
- AES256: AES with key length of 256 bits

## C.2 pgx\_stat\_lwlock

The pgx\_stat\_lwlock view displays statistics related to lightweight locks, with each type of content displayed on a separate line.

Table C.1 pgx\_stat\_lwlock view

Column	Туре	Description
lwlock_name	name	Name of the lightweight lock
total_waits	bigint	Number of waits caused by the lightweight lock
total_wait_time	double precision	Number of milliseconds spent in waits caused by the lightweight lock
stats_reset	timestamp with timezone	Last time at which this statistics was reset

## C.3 pgx\_stat\_latch

The pgx\_stat\_latch view displays statistics related to latches, with each type of wait information within FUJITSU Enterprise Postgres displayed on a separate line.

Table C.2 pgx stat latch view

Table C.2 pgx_stat_latch view		
Column	Туре	Description
latch_name	name	Name of the latch
total_waits	bigint	Number of waits caused a wait
total_wait_time	double precision	Number of milliseconds spent in waits caused by the latch
stats_reset	timestamp with timezone	Last time at which this statistic was reset

## C.4 pgx\_stat\_walwriter

The pgx\_stat\_walwriter view displays statistics related to WAL writing, in a single line.

Table C.3 pgx\_stat\_walwriter view

Column	Туре	Description
dirty_writes	bigint	Number of times old WAL buffers were written to the disk because the WAL buffer was full when WAL records were added
writes	bigint	Number of WAL writes
write_blocks	bigint	Number of WAL write blocks
total_write_time	double precision	Number of milliseconds spent on WAL writing
stats_reset	timestamp with timezone	Last time at which this statistic was reset

## C.5 pgx\_stat\_sql

 $The \ pgx\_stat\_sql\ view\ displays\ statistics\ related\ to\ SQL\ statement\ executions, with\ each\ type\ of\ SQL\ statement\ displayed\ on\ a\ separate\ line.$ 

Table C.4 pgx stat sql view

Column	Туре	Description
selects	bigint	Number of SELECT statements executed
		In database multiplexing mode, this number includes the SELECT statements executed in Mirroring Controller. Mirroring Controller executes the SELECT statement using the interval specified for the heartbeat_interval of the server definition file (milliseconds).
inserts	bigint	Number of INSERT statements executed
deletes	bigint	Number of DELETE statements executed
updates	bigint	Number of UPDATE statements executed
selects_with_parallelism	bigint	Number of times parallel scan was used in SELECT statements
inserts_with_parallelism	bigint	Not used
deletes_with_parallelism	bigint	Not used
updates_with_parallelism	bigint	Not used
copies_with_parallelism	bigint	Not used
declares	bigint	Number of DECLARE statements executed (number of cursor OPENs)
fetches	bigint	Number of FETCH statements executed
checkpoints	bigint	Number of CHECKPOINT statements executed
clusters	bigint	Number of CLUSTER statements executed
copies	bigint	Number of COPY statements executed
reindexes	bigint	Number of REINDEX statements executed
truncates	bigint	Number of TRUNCATE statements executed
locks	bigint	Number of times a lock occurred
stats_reset	timestamp with timezone	Last time at which this statistic was reset

## C.6 pgx\_stat\_gmc

The pgx\_stat\_gmc view provides information about the GMC areas.

Table C.5 pgx\_stat\_gmc view

Column	Туре	Description
searches	bigint	Number of times the cache table is searched.
hits	bigint	Number of times the cache table is hit.
size	bigint	The current amount of memory (bytes) used in the GMC area.
stats_reset	timestamp with timezone	Last time these statistics were reset.

## Appendix D Tables Used by Data Masking

This appendix explains tables used by the data masking feature.



These tables are updated by the data masking control function, so do not use SQL statements to directly update these tables.

## D.1 pgx\_confidential\_columns

This table provides information on masking target for which masking policies are set.

Column	Туре	Description
schema_name	varchar(63)	Schema name of table for which a masking policy is applied
table_name	varchar(63)	Name of table for which a masking policy is applied
policy_name	varchar(63)	Masking policy name
column_name	varchar(63)	Masking target name
function_type	varchar(63)	Masking type
		- 'FULL': Full masking
		- 'PARTIAL': Partial masking
		- 'REGEXP': Regular expression masking
function_parameters	varchar(1024)	Masking format for partial masking
regexp_pattern	varchar(1024)	Search pattern for regular expression masking
regexp_replacement	varchar(1024)	Replacement character/string for regular expression masking
regexp_flags	varchar(20)	Regular expression flags
column_description	varchar(1024)	Masking target description

#### **Execution example**

## D.2 pgx\_confidential\_policies

This table provides information on masking policies.

Column	Туре	Description
schema_name	varchar(63)	Schema name of table for which a masking policy is applied

Column	Туре	Description
table_name	varchar(63)	Name of table for which a masking policy is applied
policy_name	varchar(63)	Masking policy name
expression	varchar(1024)	Masking condition
enable	boolean	Masking policy status
		- 't': Enabled
		- 'f': Disabled
policy_description	varchar(1024)	Masking policy description

#### **Execution example**

## D.3 pgx\_confidential\_values

This table provides information on replacement characters when full masking is specified for masking type.

Column	Data type	Description	Default value
number_value	integer	Numeric	0
char_value	varchar(1)	char type	Spaces
varchar_value	varchar(1)	varchar type	Spaces
date_value	date	date type	'1970-01-01'
timestamp_value	timestamp	timestamp type	'1970-01-01 00:00:00'

#### **Execution example**

## Appendix E Tables Used by High-Speed Data Load

This appendix describes the tables used by high-speed data load.

## E.1 pgx\_loader\_state

 $The \ pgx\_loader\_state \ table \ provides \ information \ about \ transactions \ prepared \ by \ high-speed \ data \ load.$ 

Column	Туре	Description
id	serial	Unique identifier.
		This value is assigned from the pgx_loader_state_id_seq sequence.
gid	text	Global transaction identifier assigned to a transaction.
state	text	State of the transaction.
		The value can be one of the following:
		- commit: The prepared transaction has been committed.
		- rollback: The prepared transaction is in in-doubt state.
master_pid	integer	Process ID of the backend process (master process) that executed the pgx_loader control function.
role_oid	integer	Role identifier (OID). A prepared transaction can only be completed by the same user who executed the original transaction or by a superuser.
relation_oid	integer	Object identifier (OID).



The pgx\_loader\_state table and pgx\_loader\_state\_id\_seq sequence are updated by high-speed data load. Do not update these database objects directly using SQL.

## Appendix F Starting and Stopping the Web Server Feature of WebAdmin

To use WebAdmin for creating and managing a FUJITSU Enterprise Postgres instance on a server where FUJITSU Enterprise Postgres is installed, you must first start the Web server feature of WebAdmin.

- Using WebAdmin in a single-server configuration

You must start the Web server on the server on which FUJITSU Enterprise Postgres and WebAdmin are installed..

- Using WebAdmin in a multiserver configuration

You must start the Web server on all servers on which WebAdmin has been installed.

This appendix describes how to start and stop the Web server feature of WebAdmin.

Note that "<x>" in paths indicates the product version.



Refer to "Installing WebAdmin in a Multiserver Configuration" in the Installation and Setup Guide for Server for information on multiserver installation.

## F.1 Starting the Web Server Feature of WebAdmin

Follow the procedure below to start the Web server feature of WebAdmin.

1. Change to superuser

Acquire superuser privileges on the system.

#### Example

```
$ su -
Password:****
```

2. Start the Web server feature of WebAdmin

Execute the WebAdminStart command to start the Web server feature of WebAdmin.

#### Example

If WebAdmin is installed in /opt/fsepv<*x*>webadmin:

```
# cd /opt/fsepv<x>webadmin/sbin
# ./WebAdminStart
```

## F.2 Stopping the Web Server Feature of WebAdmin

This section describes how to stop the Web server feature of WebAdmin.

Follow the procedure below to stop the Web server feature of WebAdmin.

1. Change to superuser

Acquire superuser privileges on the system.

#### Example

```
$ su -
Password:*****
```

2. Stop the Web server feature of WebAdmin

Execute the WebAdminStop command to stop the Web server feature of WebAdmin.

#### Example

If WebAdmin is installed in /opt/fsepv<*x*>webadmin:

- # cd /opt/fsepv<x>webadmin/sbin
  # ./WebAdminStop

## Appendix G WebAdmin Wallet

This appendix describes how to use the Wallet feature of WebAdmin.

When a remote instance or a standby instance is created, it is necessary to provide user name and password for authentication with the remote machine or the database instance.

The Wallet feature in WebAdmin is a convenient way to create and store these credentials.

Once created, these credentials can be repeatedly used in one or more instances.

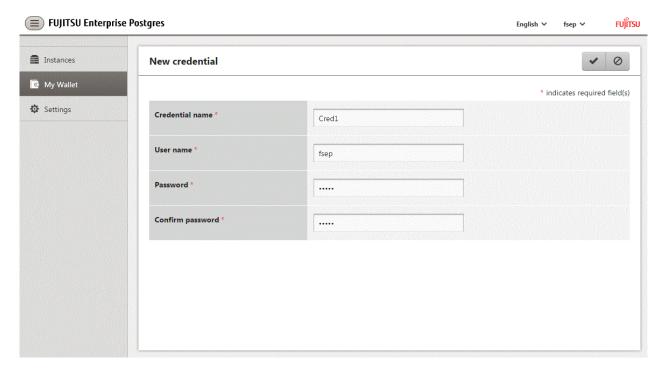


It is not mandatory to create a credential in the Wallet. It is possible to create a remote instance or a standby instance without creating any credential in the Wallet.

If no credential is created beforehand, a user name and password can be entered in the instance creation page. When creating a "Remote" instance, if operating system credentials are entered without using a credential stored in the Wallet, WebAdmin automatically creates a credential with the given user name and password, and stores it in the user's wallet for future use.

## G.1 Creating a Credential

- 1. In the [My Wallet] tab, click The [New credential] page will be displayed.
- 2. Enter the information for the credentials.



Enter the following items. Credential name, User name and Password should not contain hazardous characters. Refer to "Appendix H WebAdmin Disallow User Inputs Containing Hazardous Characters".

- [Credential name]: Name of the credential

The name must meet the conditions below:

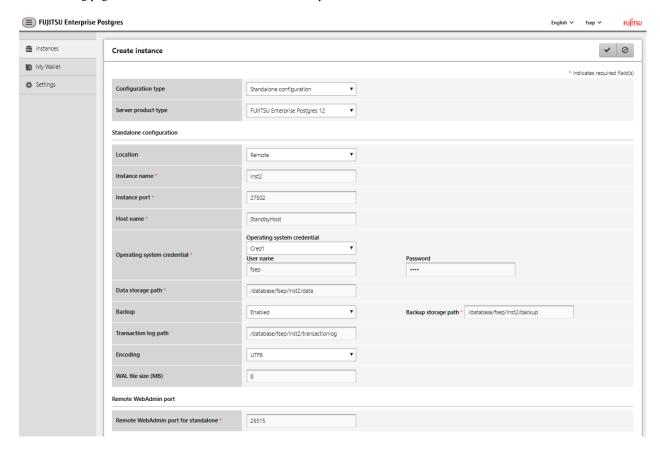
- Maximum of 16 characters
- The first character must be an ASCII alphabetic character
- The other characters must be ASCII alphanumeric characters

- [User name]: The operating system user name or database instance user name that will be used later
- [Password]: Password for the user
- [Confirm password]: Reenter the password.
- 3. Click **v** to store the credential.

## **G.2** Using a Credential

Once a credential is created in the Wallet, it can be used during remote instance creation or standby instance creation.

The following page uses the credential that was created in the previous section.



When "Cred1" is selected in [Operating system credential], the user name and password are automatically populated from the credential.

# Appendix H WebAdmin Disallow User Inputs Containing Hazardous Characters

WebAdmin considers the following as hazardous characters, which are not allowed in user inputs.

| (pipe sign)
& (ampersand sign)
; (semicolon sign)
\$ (dollar sign)
% (percent sign)
@ (at sign)
' (single apostrophe)
" (quotation mark)
\' (backslash-escaped apostrophe)
\" (backslash-escaped quotation mark)

⟨ (triangular parenthesis)
() (parenthesis)
+ (plus sign)
CR (Carriage return, ASCII 0x0d)
LF (Line feed, ASCII 0x0a)

, (comma sign)
\ (backslash)

- 153 -

# Appendix I Copy Command Samples that Use the Advanced Copy Feature of the ETERNUS Disk Array

Backup/recovery scripts that use OPC, an advanced copy feature of the FUJITSU Storage ETERNUS disk array (hereafter referred to as ETERNUS disk array), are supplied as copy command samples for use by the pgx\_dmpall and pgx\_revall commands.

Users can copy the samples to any file and make changes appropriate to their environment or operations.

The samples are stored in the directories below:

- Basic version

/installDir/share/copy\_command.esf\_acml.sh.sample

Advanced version

/installDir/share/copy\_command.esf\_acm2.sh.sample



#### See

The samples use the replication management command of the FUJITSU Storage ETERNUS SF AdvancedCopy Manager (hereafter referred to as ACM) for operating the advanced copy features of the ETERNUS disk array. Refer to the relevant manual for details.

#### Prerequisite for using the samples

To use the samples, it is necessary to configure the settings of the advanced copy features of the ETERNUS disk array in advance. The samples assume that the replication source/volume has been configured, its contents have been physically copied to the replication volume, and that the tracking process has been started.

#### Sample content (advanced version)

The sample uses two replication volumes as the replication source on which the database cluster and tablespace are located, alternating between them for each backup. It then registers to the backup information file the replication volume (group) used for the latest backup.

In addition, the backup information file is also used to determine the replication volume to restore during recovery.

The processing for each operation mode is described below:

#### prepare mode

- 1. The backup information file is read, and the replication volume to be used for the current backup is determined.
- 2. The completion status of physical copies to all replication volumes is checked using the swsrpstat command of ACM.
- 3. The replication volume determined in step 1 is written to a temporary file for later use by the backup mode.

#### backup mode

- 1. The temporary file is read, and the replication volume information is retrieved.
- 2. The file system buffer is written using the sync command of the operating system.
- 3. The file system targeted for backup is frozen using the fsfreeze command of the operating system.
- 4. Snapshot retrieval (implementation of a logical copy) is performed using the swsrpmake command of ACM.
- 5. The file system frozen in step 3 is unfrozen using the fsfreeze command of the operating system. Steps 3 to 5 are performed for all file systems targeted for backup.

#### finalize mode

1. The completion status of physical copies to all replication volumes is checked using the swsrpstat command of ACM.

2. Information about the replication volume used for the current backup is written to the backup information file.

#### restore mode

- 1. The backup information file is read, and the replication volume to be used for recovery is determined.
- 2. The file system targeted for recovery is unmounted using the umount command of the operating system.
- 3. Physical copy from the replication volume is performed using the swsrpmake command of ACM.
- 4. The file system targeted for recovery is mounted using the mount command of the operating system.
- 5. The completion status of physical copy from the replication volume is checked using the swsrpstat command of ACM. Steps 2 to 5 are performed for all file systems targeted for recovery.
- 6. Files and directories not needed for archive recovery are deleted.



- The samples cannot be used on SLES 12.
- The samples use the sudo command of the operating system so that operation of the commands and file systems of ACM is performed by the superuser of the operating system. Therefore, determine if these implementations satisfy the security standards on the database server, and if necessary, perform implementations using other means.
- The samples temporarily freeze the file system of the data storage destination (replication source volume) to protect the file system from copy processing by advanced copy features of the ETERNUS disk array. Therefore, consider the following:
  - Consider freeze time during the timeout period of an SQL statement.

When using a feature that links with the database, such as a cluster feature that accesses the data storage destination, consider freeze time in relation to the timeout periods below.

- When performing database multiplexing
  - Misdetection may occur during abnormality monitoring, so it is necessary to consider the monitoring interval, timeout period and number of retries for abnormality monitoring, and consider temporarily stopping only the Mirroring Controller process during backup.
- When performing failover operations using PRIMECLUSTER
  - If a failure occurs while the file system is frozen, switching is triggered when PRIMECLUSTER detects an issue, but will be performed only after the system is unfrozen. In addition, if it takes time for the system to unfreeze, the active node may trigger operating system panic, resulting in a switch.

## Appendix J Collecting Failure Investigation Data

If the cause of an error that occurs while building the environment or during operations is unclear, data must be collected for initial investigation.

This appendix describes how to collect data for initial investigation.

Use FJQSS (Information Collection Tool) to collect data for initial investigation.



Refer to the FJQSS manual for information on how to use FJQSS.



- When using FJQSS to collect data for initial investigation, you must set the following environment variables:
  - Environment variables required for using FUJITSU Enterprise Postgres

Refer to "Configure the environment variables" under the procedure for creating an instance in "Using the initdb Command" in the Installation and Setup Guide for Server for information on the values to be set in the environment variables.

.....

- PGDATA

Set the data storage destination.

- PGDATABASE

Set the database name from which you want to collect data for initial investigation.

- PGPORT

Set the instance port number. This does not need to be set if the default port number (27500) has not been changed.

- PGUSER

Set the database superuser.

Set the database superuser so that client authentication is possible.

FJQSS establishes a TCP/IP connection with the template1 database and collects data from the database.

- FSEP HOME

Set the FUJITSU Enterprise Postgres installation directory.

- Refer to "Collecting Failure Investigation Data" in the Cluster Operation Guide (Database Multiplexing) for information on how to collect failure investigation data when performing database multiplexing.

## <u>Index</u>

[/1]	F	20
[A]	Encrypting a Tablespace	
Actions in Response to Instance Startup Failure	Encrypting Existing Data	
All user data within the specified tablespace	Encryption mechanisms	
Approximate backup time	Errors in More Than One Storage Disk	110
Approximate recovery time	[F]	
Automatically opening the keystore	Faster encryption and decryption based on hardware	26
[B]	File system level backup and restore	
Backing Up and Recovering the Keystore31	The system is the custoff and restore in the control of the contro	
Backing Up and Restoring/Recovering the Database33	[H]	
Backup/Recovery Using the Copy Command79	High-Speed Data Load	<mark>7</mark> 1
Backup and recovery using the pgx_dmpall and pgx_rcvall		
commands	[1]	
backup cycle	If failure occurred in the data storage disk or the transaction	
Backup data	storage disk	
Backup operation	If failure occurred on the backup data storage disk	
Backup operation (file backup)	If failure occurred on the data storage disk or the transaction	
Backup status	storage directory	
Backup using the backup information file80	Importing and Exporting the Database	35
Backup Using the Copy Command	Installing and Operating the In-memory Feature	62
backup_destination (string)	11/2	
Building and starting a standby server	[K]	104
Building and starting a standby server	keystore_location (string)	124
[C]	[L]	
Changing a Masking Policy45	Logging in to WebAdmin	2
Changing the Keystore Passphrase30	log in	
Changing the Master Encryption Key30	10g III	
Changing the master encryption key and the passphrase 37	[M]	
Checking an Encrypted Tablespace29	Managing the Keystore	30
Checking backup status	Masking Condition	
Checking the operating status of an instance13,15	Masking Format	
Collecting Failure Investigation Data156	Masking Policy	
Configuration of the Copy Command	Masking Target	
Configuration of the copy command for backup81	Masking Type	
Configuration of the copy command for recovery	Monitoring Database Activity	
Confirming a Masking Policy45	·	
Continuous archiving and point-in-time recovery34	[O]	
Copy Command for Backup85	Opening the Keystore	28
Copy Command for Recovery87	Operating FUJITSU Enterprise Postgres	1
Copy Command Interface85	וחו	
Copy Command Samples that Use the Advanced Copy Feature of	[P]	70
the ETERNUS Disk Array	Parallel Query	
core_contents (string)124	Performing backup	
core_directory (string)124	Perform recovery	
Creating a Masking Policy44	Periodic Backup	
Cyclic usage of the backup area80	pgx_global_metacache (numerical value)	
-,g	pgx_stat_gmc view	
[D]	pgx_stat_latch view	
Data Masking39	pgx_stat_lwlock view	
Data Types for Masking47	pgx_stat_sql view	
Deleting a Masking Policy47	pgx_stat_walwriter view	
Determining the backup area of the latest backup84	pgx_tablespaces	
	Placement and automatic opening of the keystore file	
[E]	Placing the keystore file	
Enabling and Disabling a Masking Policy	Preparing for backup	83
Enabling Automatic Opening of the Keystore 30		

[R]	
Recovery Using the Copy Command	34
reserve_buffer_ratio (numerical value)12	25
[S]	
Scope of encryption	
search_path (string)	
Security-Related Notes	
Security Notes4	
Setting a restore point	
Setting the Master Encryption Key	
Starting and Stopping the Web Server Feature of WebAdmin14	
Starting an instance	
Starting pgAdmin	
Startup URL for WebAdmin	
Stopping an instance	
Streaming replication support	
Streaming Replication Using WebAdmin5	
Strong encryption algorithms	
System Administration Functions	
System Views	13
[T]	
tablespace_encryption_algorithm (string)12	0.1
Tables Used by Data Masking	
Tips for Installing Built Applications	
track_gmc (string)	
track_sql (string)	
track_waits (string)	
Transparent Data Encryption Control Functions	
Two-layer encryption key and the keystore	20
[U]	
User environment	3
Using Server Commands	4
[V]	
vci.control_max_workers (numerical value)12	
vci.cost_threshold (numeric)	
vci.enable (string)	
vci.log_query (string)	
vci.maintenance_work_mem (numerical value)12	
vci.max_local_ros (numerical value)	
vci.max_parallel_degree (numerical value) 12	
vci.shared_work_mem (numerical value)12	26



# FUJITSU Enterprise Postgres 12



Operation Guide (Windows)

## **Preface**

#### Purpose of this document

The FUJITSU Enterprise Postgres database system extends the PostgreSQL features and runs on the Windows platform.

This document is the FUJITSU Enterprise Postgres Operation Guide.

#### Intended readers

This document is intended for those who install and operate FUJITSU Enterprise Postgres.

Readers of this document are assumed to have general knowledge of:

- PostgreSQL
- SQL
- Windows

#### Structure of this document

This document is structured as follows:

#### Chapter 1 Operating FUJITSU Enterprise Postgres

Describes how to operate FUJITSU Enterprise Postgres.

#### Chapter 2 Starting an Instance and Creating a Database

Describes how to start a FUJITSU Enterprise Postgres instance, and how to create a database.

#### Chapter 3 Backing Up the Database

Describes how to back up the database.

#### Chapter 4 Configuring Secure Communication Using Secure Sockets Layer

Describes communication data encryption between the client and the server.

#### Chapter 5 Protecting Storage Data Using Transparent Data Encryption

Describes how to encrypt the data to be stored in the database.

#### Chapter 6 Data Masking

Describes the data masking feature.

#### Chapter 7 Periodic Operations

Describes the periodic database operations that must be performed on FUJITSU Enterprise Postgres.

#### Chapter 8 Streaming Replication Using WebAdmin

Describes how to create a streaming replication cluster using WebAdmin.

#### Chapter 9 Installing and Operating the In-memory Feature

Describes how to install and operate the in-memory feature.

#### Chapter 10 Parallel Query

Describes the factors taken into consideration by FUJITSU Enterprise Postgres when performing parallel

#### queries. Chapter 11 High-Speed Data Load

Describes how to install and operate high-speed data load.

#### Chapter 12 Global Meta Cache

Describes how to use Grobal Meta Cache feature.

#### Chapter 13 Backup/Recovery Using the Copy Command

Describes backup and recovery using the copy command created by the user.

#### Chapter 14 Actions when an Error Occurs

Describes how to perform recovery when disk failure or data corruption occurs.

#### Appendix A Parameters

Describes the FUJITSU Enterprise Postgres parameters.

#### Appendix B System Administration Functions

Describes the system administration functions of FUJITSU Enterprise Postgres.

#### Appendix C System Views

Describes how to use the system view in FUJITSU Enterprise Postgres.

#### Appendix D Tables Used by Data Masking

Describes the tables used by the data masking feature.

#### Appendix E Tables Used by High-Speed Data Load

Describes the tables used by high-speed data load.

#### Appendix F Starting and Stopping the Web Server Feature of WebAdmin

Describes how to start and stop WebAdmin (Web server feature).

#### Appendix G WebAdmin Wallet

Describes how to use the Wallet feature of WebAdmin.

#### Appendix H WebAdmin Disallow User Inputs Containing Hazardous Characters

Describes characters not allowed in WebAdmin.

#### Appendix I Collecting Failure Investigation Data

Describes how to collect information for initial investigation.

#### **Export restrictions**

Exportation/release of this document may require necessary procedures in accordance with the regulations of your resident country and/or US export control laws.

#### Issue date and version

Edition 1.0: July 2020

#### Copyright

Copyright 2015-2020 FUJITSU LIMITED

## **Contents**

Chapter 1 Operating FUJITSU Enterprise Postgres	
1.1 Operating Methods	1
1.2 Starting WebAdmin	2
1.2.1 Logging in to WebAdmin	3
1.3 Starting pgAdmin	
1.3.1 Starting pgAdmin	
1.3.2 Adding an Instance.	4
1.3.3 Connecting/Disconnecting an Instance	<i>6</i>
1.4 Operations Using Commands	<i>6</i>
1.5 Operating Environment of FUJITSU Enterprise Postgres	
1.5.1 Operating Environment	
1.5.2 File Composition	
1.6 Notes on Compatibility of Applications Used for Operations	10
1.7 Notes on Upgrading Database Instances	10
1.7.1 Additional Steps for upgrading to FUJITSU Enterprise Postgres with Vertical Clustered Index (VCI) Enabled	11
Chapter 2 Starting an Instance and Creating a Database	13
2.1 Starting and Stopping an Instance.	
2.1.1 Using WebAdmin.	
2.1.2 Using Commands.	
2.2 Creating a Database	
2.2.1 Using pgAdmin.	
2.2.2 Using Client Commands	
Chapter 3 Backing Up the Database	18
3.1 Periodic Backup	19
3.2 Backup Methods	19
3.2.1 Using WebAdmin.	19
3.2.2 Using Server Commands	20
Chanter 4 Configurity Converse Communication Union Converse Contrate Layer	2.
Chapter 4 Configuring Secure Communication Using Secure Sockets Layer	
4.1 Configuring Communication Data Encryption	
4.1.1 Issuing a Certificate.	
4.1.2 Deploying a Server Certificate File and a Server Private Key File	
4.1.3 Distributing a CA Certificate File to the Client	
4.1.4 Configuring the Operating Environment for the Database Server.	
4.1.5 Configuring the Operating Environment for the Client	
4.1.6 Performing Database Multiplexing	20
Chapter 5 Protecting Storage Data Using Transparent Data Encryption	27
5.1 Protecting Data Using Encryption.	
5.2 Setting the Master Encryption Key	
5.3 Opening the Keystore	
5.4 Encrypting a Tablespace	
5.5 Checking an Encrypted Tablespace	
5.6 Managing the Keystore	
5.6.1 Changing the Master Encryption Key	
5.6.2 Changing the Keystore Passphrase	
5.6.3 Enabling Automatic Opening of the Keystore	
5.6.4 Backing Up and Recovering the Keystore	
5.7 Backing Up and Restoring/Recovering the Database.	
5.8 Importing and Exporting the Database	
5.9 Encrypting Existing Data	
5.10 Operations in Cluster Systems.	
5.10 Operations in Cluster Systems.  5.10.1 HA Clusters that do not Use Database Multiplexing	
5.10.2 Database Multiplexing Mode	

5.11 Security-Related Notes	
5.12 Tips for Installing Built Applications	39
Chapter 6 Data Masking	40
6.1 Masking Policy	
6.1.1 Masking Target	
6.1.2 Masking Type	
6.1.3 Masking Condition.	
6.1.4 Masking Format	
6.2 Usage Method	
6.2.1 Creating a Masking Policy	45
6.2.2 Changing a Masking Policy	
6.2.3 Confirming a Masking Policy	46
6.2.4 Enabling and Disabling a Masking Policy	47
6.2.5 Deleting a Masking Policy	48
6.3 Data Types for Masking	48
6.4 Security Notes	49
Chapter 7 Periodic Operations	50
7.1 Configuring and Monitoring the Log	
7.2 Monitoring Disk Usage and Securing Free Space	
7.2.1 Monitoring Disk Usage	
7.2.2 Securing Free Disk Space	
7.3 Automatically Closing Connections	
7.4 Monitoring the Connection State of an Application	
7.4.1 Using the View (pg_stat_activity)	
7.4.2 Using pgAdmin.	
7.5 Reorganizing Indexes.	
7.6 Monitoring Database Activity	
7.6.1 Information that can be Collected	
7.6.2 Collection Configuration.	
7.6.3 Information Reset	
Chapter 8 Streaming Replication Using WebAdmin	
8.1 Creating a Standby Instance	
8.2 Promoting a Standby Instance.	
8.3 Converting an Asynchronous Replication to Synchronous	
8.4 Converting a Synchronous Replication to Asynchronous	
8.5 Joining a Replication Cluster	62
Chapter 9 Installing and Operating the In-memory Feature	
9.1 Installing Vertical Clustered Index (VCI)	
9.1.1 Evaluating whether to Install VCI	
9.1.2 Estimating Resources.	
9.1.3 Setting up	
9.1.3.1 Setting Parameters.	
9.1.3.2 Installing the Extensions.	
9.1.3.3 Creating a VCI	
9.1.3.4 Confirming that the VCI has been Created	
9.1.4 Data that can Use VCI	
9.1.4.1 Relation Types	
9.1.4.2 Data Types	
9.2.1 Commands that cannot be Used for VCI	
9.2.2 Data Preload Feature	
Chapter 10 Parallel Query	70
10.1 CPLLI and Calculation	77

10.2 Increase of Workers during Runtime	72
Chapter 11 High-Speed Data Load	73
11.1 Installing High-Speed Data Load	
11.1.1 Deciding whether to Install	73
11.1.2 Estimating Resources.	73
11.1.3 Setup	74
11.1.3.1 Setting Parameters	74
11.1.3.2 Installing the Extension	75
11.2 Using High-Speed Data Load	75
11.2.1 Loading Data	75
11.2.2 Recovering from a Data Load that Ended Abnormally	76
11.3 Removing High-Speed Data Load	77
11.3.1 Removing the Extension	77
Chapter 12 Global Meta Cache	79
12.1 Usage	
12.1.1 Deciding Whether to Enable the Global Meta Cache Feature	79
12.1.2 Estimating Memory for Global Meta Cache	79
12.1.3 How the GMC Memory Area Is Used	79
12.1.4 Enabling the Global Meta Cache Feature	79
12.1.5 Estimating Resources.	80
12.2 Statistics.	80
12.2.1 System View	80
Chapter 13 Backup/Recovery Using the Copy Command	81
13.1 Configuration of the Copy Command	81
13.2 Backup Using the Copy Command	
13.3 Recovery Using the Copy Command	85
13.4 Copy Command Interface	86
13.4.1 Copy Command for Backup	86
13.4.2 Copy Command for Recovery	88
Chapter 14 Actions when an Error Occurs	90
14.1 Recovering from Disk Failure (Hardware)	91
14.1.1 Using WebAdmin	91
14.1.2 Using Server Command	92
14.2 Recovering from Data Corruption	96
14.2.1 Using WebAdmin	96
14.2.2 Using the pgx_rcvall Command	97
14.3 Recovering from an Incorrect User Operation	98
14.3.1 Using WebAdmin	98
14.3.2 Using the pgx_rcvall Command	99
14.4 Actions in Response to an Application Error	
14.4.1 When using the view (pg_stat_activity)	
14.4.2 Using pgAdmin	101
14.5 Actions in Response to an Access Error	102
14.6 Actions in Response to Insufficient Space on the Data Storage Destination	
14.6.1 Using a Tablespace	
14.6.2 Replacing the Disk with a Larger Capacity Disk	
14.6.2.1 Using WebAdmin	
14.6.2.2 Using Server Commands	
14.7 Actions in Response to Insufficient Space on the Backup Data Storage Destination	
14.7.1 Temporarily Saving Backup Data	
14.7.1.1 Using WebAdmin	105
14.7.1.2 Using Server Commands	106
14.7.2 Replacing the Disk with a Larger Capacity Disk	
14.7.2.1 Using WebAdmin	109

14.7.2.2 Using Server Commands	110
14.8 Actions in Response to Insufficient Space on the Transaction Log Storage Destination	113
14.8.1 Replacing the Disk with a Larger Capacity Disk	113
14.8.1.1 Using WebAdmin	
14.8.1.2 Using Server Commands	
14.9 Errors in More Than One Storage Disk	116
14.10 Actions in Response to Instance Startup Failure	116
14.10.1 Errors in the Configuration File	116
14.10.2 Errors Caused by Power Failure or Mounting Issues	117
14.10.3 Other Errors	117
14.10.3.1 Using WebAdmin	117
14.10.3.2 Using Server Commands	117
14.11 Actions in Response to Failure to Stop an Instance	117
14.11.1 Using WebAdmin	118
14.11.2 Using Server Commands	118
14.11.2.1 Stopping the Instance Using the Fast Mode	118
14.11.2.2 Stopping the Instance Using the Immediate Mode	
14.11.2.3 Forcibly Stopping the Server Process.	
14.12 Actions in Response to Failure to Create a Streaming Replication Standby Instance	119
14.13 Actions in Response to Error in a Distributed Transaction	
14.14 I/O Errors Other than Disk Failure	
14.14.1 Network Error with an External Disk	121
14.14.2 Errors Caused by Power Failure or Mounting Issues	
14.15 Anomaly Detection and Resolution	121
14.15.1 Port Number and Backup Storage Path Anomalies.	121
14.15.2 Mirroring Controller Anomalies	
Appendix A Parameters	123
7 Apportuix 711 didinotors	
Appendix B System Administration Functions	127
B.1 WAL Mirroring Control Functions.	127
B.2 Transparent Data Encryption Control Functions.	
B.3 Data Masking Control Functions.	
B.3.1 pgx_alter_confidential_policy	128
B.3.2 pgx_create_confidential_policy	
B.3.3 pgx_drop_confidential_policy	
B.3.4 pgx_enable_confidential_policy	
B.3.5 pgx_update_confidential_values.	139
B.4 VCI Data Load Control Function.	
B.5 High-Speed Data Load Control Functions	141
Appendix C System Views	1/13
C.1 pgx_tablespaces	
C.2 pgx_stat_lwlock	
C.3 pgx_stat_latch	
C.4 pgx_stat_walwriter	
C.5 pgx_stat_sql	
C.6 pgx_stat_gmc.	
C.o pgx_smt_gnc	
Appendix D Tables Used by Data Masking	145
D.1 pgx_confidential_columns	145
D.2 pgx_confidential_policies	
D.3 pgx_confidential_values	
A	
Appendix E Tables Used by High-Speed Data Load	
E.1 pgx_loader_state	147
Appendix F Starting and Stopping the Web Server Feature of WebAdmin	148
F.1 Starting the Web Server Feature of WebAdmin.	

F.2 Stopping the Web Server Feature of WebAdmin	148
Appendix G WebAdmin Wallet	150
G.1 Creating a Credential	
G.2 Using a Credential.	151
Appendix H WebAdmin Disallow User Inputs Containing Hazardous Characters	152
Appendix I Collecting Failure Investigation Data	153
Index	154

## Chapter 1 Operating FUJITSU Enterprise Postgres

This chapter describes how to operate FUJITSU Enterprise Postgres.

## 1.1 Operating Methods

There are two methods of managing FUJITSU Enterprise Postgres operations:

- Operation management using GUI tools
- Operation management using commands



See

Before performing database multiplexing using database multiplexing, refer to "Database Multiplexing Mode" in the Cluster Operation Guide (Database Multiplexing).

#### Operation management using GUI tools

This involves managing operations using the WebAdmin and pgAdmin GUI tools.

- Management using WebAdmin

This removes the requirement for complex environment settings and operational design for backup and recovery that is usually required for running a database. It enables you to easily and reliably monitor the state of the database, create a streaming replication cluster, back up the database, and restore it even if you do not have expert knowledge of databases.

- Management using pgAdmin

When developing applications and maintaining the database, you can use pgAdmin to perform simple operations on database objects, such as:

- Rebuild indexes and update statistics
- Create, delete, and update database objects

In addition, from pgAdmin of FUJITSU Enterprise Postgres, you can use the expanded features provided by FUJITSU Enterprise Postgres on the PostgreSQL SQL commands.



See

Refer to pgAdmin Help for information on the expanded features of pgAdmin provided by FUJITSU Enterprise Postgres.

#### Operation management using commands

You can use commands for configuring and operating the database and managing operations.



You cannot combine WebAdmin and server commands to perform the following operations:

- Use commands to operate an instance created using WebAdmin.
- Use WebAdmin to recover a database backed up using commands.

For instances created with WebAdmin, however, backup can be obtained with the pgx\_dmpall command. Also, WebAdmin can perform recovery by using the backup obtained with the pgx\_dmpall command.

- To operate an instance created using the initdb command in WebAdmin, the instance needs to be imported using WebAdmin.

- You can perform backup and restoration in pgAdmin, but the backup data obtained with WebAdmin and pgx\_dmpall is not compatible with the backup data obtained with pgAdmin.
- Refer to pgAdmin Help for other notes on pgAdmin.

#### Features used in each phase

The following table lists the features used in each phase for GUI-based operations and command-based operations.

Operation		Operation with the GUI	Operation with commands	
Setup	Creating an instance	WebAdmin is used. The server machine capacity, and the optimum parameter for operations using WebAdmin, are set automatically.	The configuration file is edited directly using the initdb command.	
	Creating a standby instance	WebAdmin is used.  WebAdmin performs a base backup of the source instance and creates a standby instance.	A standby instance is created using the pg_basebackup command.	
	Changing the configuration files	WebAdmin is used.	The configuration file is edited directly.	
Starting and stopping an instance		WebAdmin is used.	The net command or sc command of the operating system is used.	
Creating a database		This is defined using pgAdmin of the GUI tool, or using the psql command or the application after specifying the DDL statement.		
Backing up the database		WebAdmin, or the pgx_dmpall command, is used.	It is recommended that the pgx_dmpall command be used. Recovery to the latest database can be performed.	
Database recovery		WebAdmin is used.	To use the backup that was performed using the pgx_dmpall command, the pgx_rcvall command is used.	
Monitoring	Database errors	The status in the WebAdmin window can be checked. (*1)	The messages that are output to the database server log are monitored (*1)	
	Disk space	The status in the WebAdmin window can be checked. A warning will be displayed if the free space falls below 20%. (*1)	This is monitored using the fsutil command (check free space), and the dir command (check used space), of the operating system, for example. (*1)	
	Connection status	This can be checked using pgAdmin of the GUI tool, or referencing pg_stat_activity of the standard statistics view from psql or the application		

<sup>\*1:</sup> This can be used together with system log monitoring using operations management middleware (Systemwalker Centric Manager, for example).

## 1.2 Starting WebAdmin

This section describes how to start and log in to WebAdmin.

## 1.2.1 Logging in to WebAdmin

This section describes how to log in to WebAdmin.

#### **User environment**

It is recommended to use the following browsers with WebAdmin:

- Internet Explorer 11
- Microsoft Edge (Build41 or later)

WebAdmin will work with other browsers, such as Firefox and Chrome, however, the look and feel may be slightly different.

#### Startup URL for WebAdmin

In the browser address bar, type the startup URL of the WebAdmin window in the following format:

http://hostNameOrIpAddress:portNumber/

- hostNameOrIpAddress: The host name or IP address of the server where WebAdmin is installed.
- portNumber. The port number of WebAdmin. The default port number is 27515.



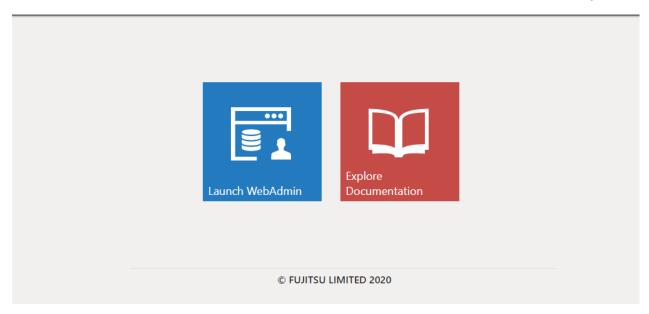
For a server with IP address "192.0.2.0" and port number "27515"

http://192.0.2.0:27515/

The startup URL window shown below is displayed. From this window you can log in to WebAdmin or access the product documentation.

## **FUJITSU Enterprise Postgres**







- You must start the Web server feature of WebAdmin before using WebAdmin.

- Refer to "Appendix F Starting and Stopping the Web Server Feature of WebAdmin" for information on how to start the Web server feature of WebAdmin.

#### Log in to WebAdmin

Click [Launch WebAdmin] in the startup URL window to start WebAdmin and display the login window.

To log in, specify the following values:

- [User name]: User name (OS user account) of the instance administrator
- [Password]: Password corresponding to the user name



Use the OS user account as the user name of the instance administrator. Refer to "Creating an Instance Administrator" in the Installation and Setup Guide for Server for details.

## 1.3 Starting pgAdmin

This section describes how to start pgAdmin, how to add an instance required for managing a database, and how to connect to and disconnect from the instance.

You can use pgAdmin on the Windows client.

#### 1.3.1 Starting pgAdmin

This section explains how to start pgAdmin if you are using it from the product "FUJITSU Enterprise Postgres Client (AAbit) x SPz" (where AA is "32" or "64", "x" is the product version, and "z" is the product level (x SPz)).

Windows Server(R) 2012

From the [Start] screen, start [pgAdmin 4 (AAbit) (x SPz)].

Windows(R) 8.1 or Windows Server(R) 2012 R2

From the [Apps] view, start [pgAdmin 4 (AAbit) (x SPz)].

Windows(R) 10, Windows Server(R) 2016, or Windows Server(R) 2019

Click [Start] >> [All apps] >> [FUJITSU Enterprise Postgres Client(AAbit) x SPz] and start [pgAdmin 4 (AAbit) (x SPz)].



- You must start the instance to be connected to before using pgAdmin.
- Refer to "2.1 Starting and Stopping an Instance" for information on how to start an instance.
- When using pgAdmin4 with Microsoft Edge, enable network access by loopback in Microsoft Edge. Also, add Microsoft Edge to the loopback exclusion list.

## 1.3.2 Adding an Instance

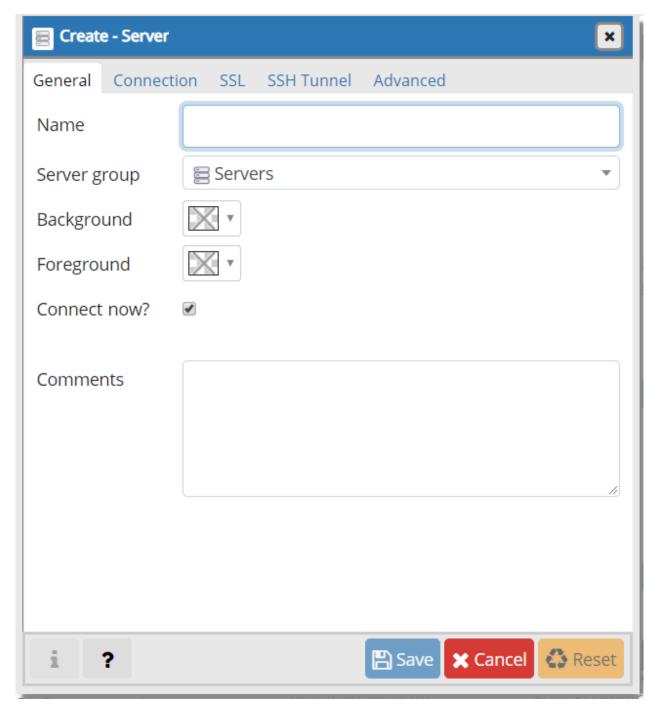
This section describes how to add an instance to be connected to.



If you use a link-local address with version 6 of the TCP/IP protocol, you may encounter the following error when registering the server with pgAdmin. Therefore, do not use link-local addresses.

unsupported format character '"' (0x22) at index 96

- 1. In the [Browser] pane, right-click [Servers], and then click [Create] >> [Server].
- 2. In the [Create Server] window, specify a value for each item.



#### [General] tab

- [Name]: Name of the instance to be managed

#### [Connection] tab

- [Host name/address]: Host name or IP address of the server where FUJITSU Enterprise Postgres is installed
- [Port]: Port number of the instance

- [Username]: User name of the instance administrator
- [Password]: Password for the user name specified in [Username]

When you add an instance using pgAdmin, the instance is automatically connected to immediately after the addition is completed.



If you select [Save password], the FUJITSU Enterprise Postgres connection password is stored in the following location. Set the appropriate access permissions for the password file to protect it from unauthorized access.

-  $APPDATA\%\Roaming\pgAdmin\pgadmin4.db$ 

#### 1.3.3 Connecting/Disconnecting an Instance

This section describes how to connect pgAdmin to an instance, and how to disconnect it.



To connect to an instance created using WebAdmin, you must first configure the settings in the [Client authentication] window of WebAdmin to permit connection from pgAdmin.

......



Refer to "Changing the settings" in the Installation and Setup Guide for Server for information on the [Client authentication] window of WebAdmin.

#### Connecting to an instance

Starting pgAdmin does not connect it to any instance.

To connect to an instance, in the [Browser] pane, right-click the instance, and then click [Connect Server].

If a password was not saved when the instance was added, enter a password in the password entry window that is displayed.

#### Disconnecting from an instance

To disconnect from an instance, in the [Browser] pane, right-click the server, and then click [Disconnect Server].

## 1.4 Operations Using Commands

You can operate and manage the database using the following commands:

- Server commands

This group of commands includes commands for creating a database cluster and controlling the database. You can run these commands on the server where the database is operating.

To use these commands, you must configure the environment variables.



#### See

 Refer to "PostgreSQL Server Applications" under "Reference" in the PostgreSQL Documentation, or "Reference" for information on server commands.

- Refer to "Configure the environment variables" in the procedure to create instances in "Using the initdb Command" in the Installation and Setup Guide for Server for information on configuring the environment variables.

#### - Client commands

This group of commands includes the psql command and commands for extracting the database cluster to a script file. These commands can be executed on the client that can connect to the database, or on the server on which the database is running.

To use these commands, you must configure the environment variables.



#### See

- Refer to "PostgreSQL Client Applications" under "Reference" in the PostgreSQL Documentation, or "Reference" for information on client commands.

......

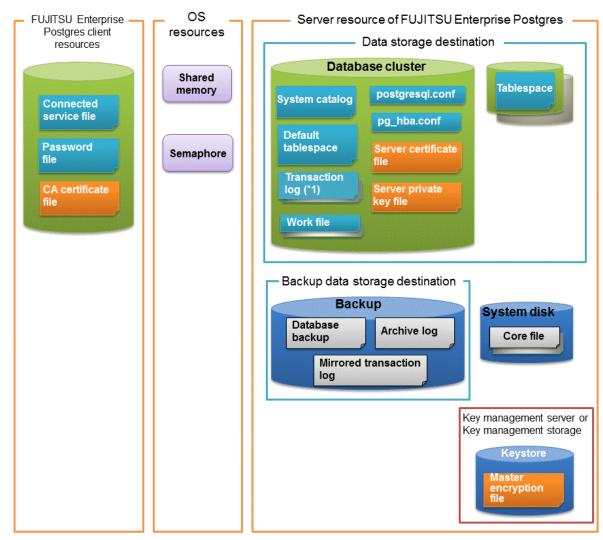
- Refer to "Configuring Environment Variables" in the Installation and Setup Guide for Client for information on the values to be set in the environment variables.

## 1.5 Operating Environment of FUJITSU Enterprise Postgres

This section describes the operating environment and the file composition of FUJITSU Enterprise Postgres.

#### 1.5.1 Operating Environment

The following figure shows the configuration of the FUJITSU Enterprise Postgres operating environment. The tables given below list the roles of the OS resources and FUJITSU Enterprise Postgres resources.



<sup>\*1:</sup> To distribute the I/O load, place the transaction log on a different disk from the data storage destination.

Table 1.1 OS resources

Туре	Role
Shared memory	Used when a database process exchanges information with an external
Semaphore	process.

Table 1.2 FUJITSU Enterprise Postgres client resources

Туре	Role
Connection service file	Specifies information, such as the host name, user name, and password, for connecting to FUJITSU Enterprise Postgres.
Password file	Securely manages the password for connecting to FUJITSU Enterprise Postgres.
CA certificate file	CA (certificate authority) certificate used for server authentication when encrypting communication data.

Table 1.3 Server resources of FUJITSU Enterprise Postgres

Туре	Role	
Database cluster	Database storage area on the database storage disk. It is a collection of databases managed by an instance.	
System catalog	Contains information required for the system to run, including the database definition information and the operation information created by the user.	
Default tablespace	Contains table files and index files stored by default.	
Transaction log	Contains log information in case of a crash recovery or rollback. This is the same as the WAL (Write Ahead Log).	
Work file	Work file used when executing applications or commands.	
postgresql.conf	Contains information that defines the operating environment of FUJITSU Enterprise Postgres.	
pg_hba.conf	FUJITSU Enterprise Postgres uses this file to authenticate individual client hosts.	
Server certificate file	Contains information about the server certificate to be used when encrypting communication data and authenticating a server.	
Server private key file	Contains information about the server private key to be used when encrypting communication data and authenticating a server	
Tablespace	Stores table files and index files in a separate area from the database cluster. Specify a space other than that under the database cluster.	
Backup	Stores the data required for recovering the database when an error, such as disk failure, occurs.	
Database backup	Contains the backup data for the database.	
Archive log	Contains the log information for recovery.	
Mirrored transaction log (mirrored WAL)	Enables a database cluster to be restored to the state immediately before an error even if both the database cluster and transaction log fail when performing backup/recovery operations using the pgx_dmpall command or WebAdmin.	
Core file	FUJITSU Enterprise Postgres process core file that is output when an error occurs during a FUJITSU Enterprise Postgres process.	
Key management server or key management storage	Server or storage where the master encryption key file is located.	
Master encryption key file	Contains the master encryption key to be used when encrypting storage data. The master encryption key file is managed on the key management server or key management storage.	

## 1.5.2 File Composition

FUJITSU Enterprise Postgres consists of the following files for controlling and storing the database. The table below shows the relationship between the number of such files and their location within a single instance.

Table 1.4 Number of files within a single instance and how to specify their location

File type	Required	Quantity	How to specify the location
			Note that "< <i>x</i> >" indicates the product version.
Program files	Y	Multiple	By default, the following location is used:
			%Program Files%\Fujitsu\fsepv< <i>x</i> >server64
Database cluster	Y	1	Specify using WebAdmin or server commands.

File type	Required	Quantity	How to specify the location
Tablespace	Y	Multiple	Specify a space other than that under the database cluster, using pgAdmin or the DDL statement.
Backup	Y	Multiple	Specify using WebAdmin or server commands.
Core file	Y	Multiple	Specify using WebAdmin, server commands, or postgresql.conf.
Server certificate file (*1)	N	1	Specify using postgresql.conf.
Server private key file (*1)	N	1	Specify using postgresql.conf.
Master encryption key file (*1)	N	1	Specify the directory created as the key store using postgresql.conf.
Connection service file (*1)	N	1	Specify using environment variables.
Password file (*1)	N	1	Specify using environment variables.
CA certificate file (*1)	N	1	Specify using environment variables.

Y: Mandatory N: Optional

<sup>\*1:</sup> Set manually when using the applicable feature.



If anti-virus software is used, set scan exception settings for directories so that none of the files that comprise FUJITSU Enterprise Postgres are scanned for viruses. Alternatively, if the files that comprise FUJITSU Enterprise Postgres are to be scanned for viruses, stop FUJITSU Enterprise Postgres and perform the scan when tasks that use FUJITSU Enterprise Postgres are not operating.

# 1.6 Notes on Compatibility of Applications Used for Operations

When you upgrade FUJITSU Enterprise Postgres to a newer version, there may be some effect on applications due to improvements or enhancements in functionality.

Take this into account when creating applications so that you can maintain compatibility after upgrading to a newer version of FUJITSU Enterprise Postgres.



Refer to "Notes on Application Compatibility" in the Application Development Guide for details.

# 1.7 Notes on Upgrading Database Instances

When upgrading FUJITSU Enterprise Postgres 9.4 or newer database instances to FUJITSU Enterprise Postgres 10 or later using pg\_upgrade, there are certain steps you need to follow.

Before using pg\_upgrade, remove the following extensions from all databases in the instance, except "template0":

- pg\_stat\_statements
- pgx\_io
- pgx\_paging

- pgx\_network
- pgx\_network\_err
- pgx\_cpu
- pgx\_memory
- pgx\_swap
- pgx\_disk
- pgx\_process
- pgx\_log
- oracle\_compatible
- pg\_dbms\_stats
- pg\_hint\_plan

For all databases except "template0", execute the following command to remove these extensions:

```
DROP EXTENSION extensionName;
```

Once the pg\_upgrade operation is complete, for all databases except "template0", execute the following command to re-create these extensions as required:

CREATE EXTENSION extensionName;



- It is strongly recommended to back up the database using pg\_dump before performing pg\_upgrade or using DROP EXTENSION.
- If there are any columns created in the user tables using a data type from these extensions, then DROP EXTENSION will also drop these columns. Therefore, it is essential that alternate upgrade mechanisms are considered instead of pg\_upgrade, in such scenarios. These may include pg\_dump/pg\_restore.

# 1.7.1 Additional Steps for upgrading to FUJITSU Enterprise Postgres with Vertical Clustered Index (VCI) Enabled

When upgrading FUJITSU Enterprise Postgres 11.0 or earlier instances that are using the VCI extension to FUJITSU Enterprise Postgres 12 using pg\_upgrade, additional steps must be performed because of the incompatibility of the VCI extension between FUJITSU Enterprise Postgres 12 and FUJITSU Enterprise Postgres 11 or earlier.

Follow the procedure below in all databases in the FUJITSU Enterprise Postgres 11 or earlier instance, except "template0".

### Before upgrading

1. Obtain the CREATE INDEX Definitions

Run the query below to list all the VCI indexes created in the database. Ensure that these indexes are re-created in the FUJITSU Enterprise Postgres 12 or later instance after pg\_upgrade has finished.

```
SELECT nspname || '.' || relname AS index_relname,* FROM pg_class, pg_namespace
WHERE relnamespace = pg_namespace.oid AND relam IN (SELECT oid FROM pg_am WHERE amname='vci');
```

For each index\_relname listed above, execute the commands below to obtain the CREATE INDEX definition (to use the same SQL syntax while re-creating the indexes on the FUJITSU Enterprise Postgres 12 or later instance).

```
SELECT pg_get_indexdef('indexName'::regclass);
```

2. Drop the VCI indexes and VCI extension along with all its dependencies.

To remove all the VCI indexes and VCI internal objects that are created in FUJITSU Enterprise Postgres, execute the commands below. VCI internal objects will be created in FUJITSU Enterprise Postgres 12 or later automatically when CREATE EXTENSION for VCI is executed.

DROP EXTENSION VCI CASCADE;



To restore the VCI extension in the FUJITSU Enterprise Postgres 11 or earlier instance, execute CREATE EXTENSION.

### After upgrading

Once the pg\_upgrade operation is complete, for all databases except "template0", execute CREATE EXTENSION to create the VCI extension, and then execute CREATE INDEX for all the VCI indexes as required.

# Chapter 2 Starting an Instance and Creating a Database

This chapter describes basic operations, from starting an instance to creating a database.

# 2.1 Starting and Stopping an Instance

This section describes how to start and stop an instance.

- 2.1.1 Using WebAdmin
- 2.1.2 Using Commands



To automatically start or stop an instance when the operating system on the database server is started or stopped, refer to "Configuring Automatic Start and Stop of an Instance" in the Installation and Setup Guide for Server and configure the settings.



The collected statistics are initialized if an instance is stopped in the "Immediate" mode or if it is abnormally terminated. To prepare for such initialization of statistics, consider regular collection of the statistics by using the SELECT statement. Refer to "The Statistics Collector" in "Server Administration" in the PostgreSQL Documentation for information on the statistics.

### 2.1.1 Using WebAdmin

WebAdmin enables you to start or stop an instance and check its operating status.

### Starting an instance

Start an instance by using the [Instances] tab in WebAdmin.

is displayed when an instance is stopped.

To start a stopped instance, click .

### Stopping an instance

Stop an instance by using the [Instances] tab in WebAdmin.

is displayed when an instance is active.

To stop an active instance, click .

#### Stop mode

Select the mode in which to stop the instance. The following describes the operations of the modes:

Stop mode	Connected clients	Backup being executed using the command
Smart mode (*1)	Waits for all connected clients to be disconnected.	Waits for backups being executed using the command to finish.
Fast mode	Rolls back all transactions being executed and forcibly disconnects clients.	Terminates backups being executed using the command.
Immediate mode	All server processes are terminated immediately. Crash recovery is executed the next time the instance is started.	

Stop mode	Connected clients	Backup being executed using the command
Kill process mode	Send SIGKILL to the process and abort all active transactions. This will lead to a crash-recovery run at the next restart.	

<sup>\*1:</sup> When the processing to stop the instance in the Smart mode has started and you want to stop immediately, use the following procedure:

- 1. Restart the Web server feature of WebAdmin.
- 2. In the [Instances] tab, click 2.
- 3. In the [Instances] tab, click , and select the Immediate mode to stop the instance.

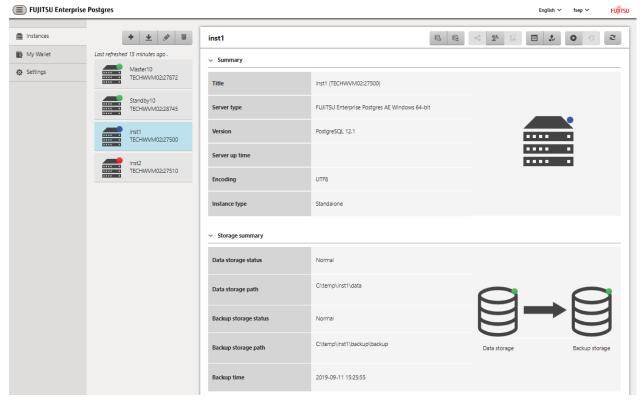
### Checking the operating status of an instance

You can check the operating status of an instance by using the [Instances] tab. The following indicators are used to show the status of a resource.

Status indicator	Explanation
•	The resource is operating normally.
•	The resource is stopped.
•	There is an error in the resource.
•	An operation is in progress on this resource or the status is being checked.
<u> </u>	The resource is not operating optimally and needs intervention.

If an instance stops abnormally, remove the cause of the stoppage and start the instance by using WebAdmin.

Figure 2.1 Example of operating status indicators





- When operating WebAdmin, click to update the status. WebAdmin will reflect the latest status of the operation or the instance resources from the server.

- If an error occurs while communicating with the server, there may be no response from WebAdmin. When this happens, close the browser and then log in again. If this does not resolve the issue, check the event log of the server and confirm whether a communication error has occurred.
- The following message is output during startup of an instance when the startup process is operating normally, therefore, the user does not need to be aware of this message:

FATAL: the database system is starting up

### 2.1.2 Using Commands

The Windows service-related commands enable you to start or stop an instance and to check its operating state.

If you are to use Windows services, you should register instances in Windows services.



Refer to "When an instance was created with WebAdmin" in "Configuring Automatic Start and Stop of an Instance" in the Installation Guide for Server for information on registering instances in Windows services.



While it is also possible for you to execute the pg\_ctl command to start and stop instances without having to register instances in Windows services, it is recommended that you use Windows services to start and stop instances for the following reason:

- If you use the pg\_ctl command to start an instance, the instance will be started as a user process. Therefore, when you close the [Command Prompt] window in which you executed the command, Windows forces the postgres process to stop.

### Starting an instance

You can start an instance by specifying the service name in the net start command or sc start command.

Also, you can use the following procedure to start an instance in the Windows services window:

- 1. Display the [Services] window
  - Windows Server(R) 2012 and Windows Server(R) 2012 R2:
     In the [Start] screen, select [Administrative Tools], and then click [Services].
  - All other operating systems:

    In the [Start] menu, select [Administrative Tools], and then click [Services].
- 2. Start a service

Select the instance name that you wish to start from the services list, and click [Start Service].

### Stopping an instance

You can stop an instance by specifying the service name in the net stop command or sc stop command.

Also, you can use the following procedure to stop an instance in the Windows services window:

- 1. Display the [Services] window
  - Windows Server(R) 2012 and Windows Server(R) 2012 R2:
     In the [Start] screen, select [Administrative Tools], and then click [Services].
  - All other operating systems:

In the [Start] menu, select [Administrative Tools], and then click [Services].

2. Stop the service

Select the instance name that you wish to stop from the services list, and click [Stop Service]. If you stop a service while applications and commands are running, FUJITSU Enterprise Postgres will force those applications and commands to close and will stop normally.

### Checking the operating state of an instance

Use the following procedure to check if an instance is operating correctly immediately after performing the operation to start an instance:

1. Display the [Services] window

In the [Start] menu, select [Administrative Tools], and then click [Services].

2. Check the state of the service

In the services list, check the state of the services for the applicable FUJITSU Enterprise Postgres.

To check the operating state of an instance during operation, use the pg\_ctl command.

Specify the following in the pg\_ctl command:

- Specify "status" as the mode.
- Specify the data storage destination directory in the -D option. If the -D option is omitted, the value of the PGDATA environment variable is used by default.



When the instance is active:

```
> pg_ctl status -D D:\database\inst1
pg_ctl: server is running (PID: 1234)
```

When the instance is inactive:

```
> pg_ctl status -D D:\database\inst1
pg_ctl: no server running
```



You can also use the net start command or sc query command to check the operating state of an instance.



Refer to "pg\_ctl" in "Reference" in the PostgreSQL Documentation for information on the pg\_ctl command.

# 2.2 Creating a Database

This section explains how to create a database.

- 2.2.1 Using pgAdmin

### 2.2.1 Using pgAdmin

Follow the procedure below to define a database using pgAdmin.

- 1. In the pgAdmin window, right-click [Databases] in the [Browser] pane, and then click [Create] >> [Database] to display a [Create Database] window.
- 2. Specify appropriate values for the following items in the [Create Database] window.
  - [General] tab
    - [Database]: Name of the database to be managed
- 3. Click [Save] to create the database.

### 2.2.2 Using Client Commands

Follow the procedure below to define a database using client commands.

An example of operations on the server is shown below.

Use psql command to connect to the postgres database.
 Execute psql postgres.

```
> psql postgres
psql (12.1)
Type "help" for help.
```

2. Create the database.

To create the database, execute the CREATE DATABASE databaseName; statement.

```
postgres=# CREATE DATABASE db01;
CREATE DATABASE
```

3. Confirm that the database is created.

Execute \l+, and confirm that the name of the database created in step 2 is displayed.

```
postgres=# \l+
```

4. Disconnect from the postgres database.

Execute  $\q$  to terminate the psql command.

```
postgres=# \q
```

You can create a database using the createdb command.



Refer to "Creating a Database" in "Tutorial" in the PostgreSQL Documentation for information on creating a database using the createdb command.

••••••

# Chapter 3 Backing Up the Database

This chapter describes how to back up the database.

#### Backup methods

The following backup methods enable you to recover data to a backup point or to the state immediately preceding disk physical breakdown or data logical failure.

- Backup using WebAdmin

This method enables you to back up data through intuitive window operations using the GUI.

WebAdmin is used for recovery.

- Backup using the pgx\_dmpall command

Execute the pgx\_dmpall command with a script to perform automatic backup.

To back up data automatically, you must register the process in the automation software of the operating system. Follow the procedure given in the documentation for your operating system.

The pgx\_rcvall command is used for recovery.



By using a copy command created by the user, the pgx\_dmpall command and the pgx\_rcvall command can back up database clusters and tablespaces to any destination and recover them from any destination using any copy method. Refer to "Chapter 13 Backup/Recovery Using the Copy Command" for details.

#### Approximate backup time

The formula for deriving the approximate backup time when you use WebAdmin or the pgx\_dmpall command is as follows:

backupTime = dataStorageDestinationUsage / diskWritePerformance x 1.5

- dataStorageDestinationUsage: Disk usage at the data storage destination
- diskWritePerformance: Maximum data volume (bytes/second) that can be written per second in the system environment where operation is performed
- 1.5: Coefficient to factor in tasks other than disk write (which is the most time-consuming step)

If using the copy command with the pgx\_dmpall command, the backup time will depend on the implementation of the copy command.



- Backup operation cannot be performed on an instance that is part of a streaming replication cluster in standby mode.
- Use the selected backup method continuously.

There are several differences, such as the data format, across the backup methods. For this reason, the following restrictions apply:

- It is not possible to use one method for backup and another for recovery.
- It is not possible to convert one type of backup data to a different type of backup data.
- Mirrored WALs can be used only for backup/recovery using the pgx\_dmpall command or WebAdmin.
- There are several considerations for the backup of the keystore and backup of the database in case the data stored in the database is encrypted. Refer to the following for details:
  - 5.6.4 Backing Up and Recovering the Keystore
  - 5.7 Backing Up and Restoring/Recovering the Database

- If you have defined a tablespace, back it up. If you do not back it up, directories for the tablespace are not created during recovery, which may cause the recovery to fail. If the recovery fails, refer to the event log, create the tablespace, and then perform the recovery process again.
- If performing backups with WebAdmin, the following password file is temporarily created during backup for WebAdmin to connect to the database:
  - userProfileFolder\localSettingsFolder\Fujitsu\fsep\_version\instanceName\pgpass.conf

Therefore, when you are backing up corefiles created in the core\_directory parameter of postgresql.conf, or log files created in the log\_directory parameter of postgresql.conf, ensure not to back up the password files located in the same directories at the same time.



The following methods can also be used to perform backup. Performing a backup using these methods allows you to restore to the point when the backup was performed.

- Backup using an SQL-based dump
  - Dump the data by using SQL. This backup method also enables data migration.
- File system level backup

This backup method requires you to stop the instance and use OS commands to backup database resources as files.

- Backup by continuous archiving

This is the standard backup method for PostgreSQL.

Refer to "Backup and Restore" in "Server Administration" in the PostgreSQL Documentation for information on these backup methods.

### 3.1 Periodic Backup

It is recommended that you perform backup periodically.

Backing up data periodically using WebAdmin or the pgx\_dmpall command has the following advantages:

- This method reduces disk usage, because obsolete archive logs (transaction logs copied to the backup data storage destination) are deleted. It also minimizes the recovery time when an error occurs.

### Backup cycle

The time interval when backup is performed periodically is called the backup cycle. For example, if backup is performed every morning, the backup cycle is 1 day.

The backup cycle depends on the jobs being run, but on FUJITSU Enterprise Postgres it is recommended that operations are run with a backup cycle of at least once per day.

# 3.2 Backup Methods

This section describes the methods for backing up the database.

- 3.2.1 Using WebAdmin
- 3.2.2 Using Server Commands

### 3.2.1 Using WebAdmin

You can use WebAdmin to perform backup and check the backup status.



- If backup is disabled for an instance, you will not be able to back up or restore the instance. Refer to "[Backup]" in "Creating an Instance" in the Installation and Setup Guide for Server for details.
- If the data to be stored in the database is to be encrypted, it is necessary to enable the automatic opening of the keystore before doing so. Refer to "5.6.3 Enabling Automatic Opening of the Keystore" for details.
- WebAdmin uses the labels "Data storage path", "Backup storage path" and "Transaction log path" to indicate "data storage destination", "backup data storage destination" and "transaction log storage destination" respectively. In this manual these terms are used interchangeably.

#### **Backup operation**

Follow the procedure below to back up the database.

1. Select the database to back up

In the [Instances] tab, select the instance to be backed up and click ...



2. Back up the database

The [Backup] dialog box is displayed. To perform backup, click [Yes]. An instance is automatically started when backup is performed.

### **Backup status**

If an error occurs and backup fails, [Error] is displayed adjacent to [Data storage status] or [Backup storage status] in the [Instances] tab. An error message is also displayed in the message list.

In this case, the backup data is not optimized. Ensure that you check the backup result whenever you perform backup. If backup fails, [Solution] appears to the right of the error message. Clicking this button displays information explaining how to resolve the cause of the error. Remove the cause of failure, and perform backup again.

### 3.2.2 Using Server Commands

Use the pgx\_dmpall command and pgx\_rcvall command to perform backup and check the backup result.

### Preparing for backup

You must prepare for backup before actually starting the backup process.

Follow the procedure below.



Refer to "Preparing Directories to Deploy Resources" in the Installation and Setup Guide for Server for information on the location of directories required for backup and for points to take into account.

1. Prepare the backup data storage disk

For backup, prepare a separate disk unit from the database storage disk and mount it using the operating system commands.

2. Create a directory where the backup data will be stored

Create an empty directory.

In [Properties] in Windows(R) Explorer, set appropriate permissions so that only the instance administrator can access the directory.



Refer to [Help and Support] in Windows(R) for information on [Properties].

3. Specify the settings required for backup

Stop the instance, and set the following parameters in the postgresql.conf file.

Start the instance after editing the postgresql.conf file.

Parameter name	Setting	Description
backup_destination	Name of the directory where the backup data will be stored	Specify the name of the directory where the backup data will be stored.
		Appropriate privileges that allow only the instance administrator to access the directory must already be set.
		Place the backup data storage destination directory outside the data storage destination directory, the tablespace directory, and the transaction log storage destination directory.
archive_mode	on	Specify the archive log mode.
		Specify [on] (execute).
archive_command	'cmd /c ""installationDirectory\\bin\ \pgx_walcopy.cmd" "%p" "backupDataStorageDestinationDirectory\ \archived_wal\\%f""	Specify the path name of the command that will save the transaction log and the storage destination.  Note the following when specifying the path:  - Specify \\ as the path delimiter.
		- Enclose the path in double quotes ("") if it contains spaces.

Refer to "Appendix A Parameters" and "Write Ahead Log" under "Server Administration" in the PostgreSQL Documentation for information on the parameters.

### **Backup operation (file backup)**

Use the pgx\_dmpall command to perform file backup. You can even embed the pgx\_dmpall command in OS automation software to perform backup.

The backup data is stored in the directory specified in the backup\_destination parameter of postgresql.conf.

Specify the data storage destination in the -D option. If the -D option is omitted, the value of the PGDATA environment variable is used by default.

......



> pgx\_dmpall -D D:\database\inst1



Backup stores the data obtained during the backup and the backup data of the data obtained during previous backup.

If the data to be stored in the database is encrypted, refer to the following and back up the keystore:

- 5.6.4 Backing Up and Recovering the Keystore

#### **Backup status**

Use the pgx\_rcvall command to check the backup status.

Specify the following values in the pgx\_rcvall command:

- The -l option indicates backup data information.
- Specify the data storage destination in the -D option. If the -D option is omitted, the value of the PGDATA environment variable is used by default.

If an error occurs and backup fails, a message is output to the event log.

In this case, the backup data is not optimized. Ensure that you check the backup result whenever you perform backup. If backup fails, remove the cause of failure and perform backup again.



#### See

Refer to "pgx\_dmpall" and "pgx\_rcvall" in the Reference for information on the pgx\_dmpall command and pgx\_rcvall command.

#### Setting a restore point

In case you want to recover your database to a certain point in time, you can name this particular point in time, which is referred to as the restore point, by using the psql command.

By setting a restore point before executing an application, it becomes easy to identify up to which point in time the data will be reverted.

A restore point can be set to any point in time after a backup is executed. However, if a restore point is set before a backup is executed, the database cannot be recovered to that point in time. This is because restore points are recorded in the archive logs, and the archive logs are discarded when backups are executed.



The following example uses the psql command to connect to the database and execute the SQL statement to set a restore point.

However, when considering continued compatibility of applications, do not use functions directly in SQL statements. Refer to "Notes on Application Compatibility" in the Application Development Guide for details.

.....

Refer to "14.3.2 Using the pgx\_rcvall Command" for information on using a restore point to recover the database.



- Name restore points so that they are unique within the database. Add the date and time of setting a restore point to distinguish it from other restore points, as shown below:
  - YYMMDD\_HHMMSS
    - YYMMDD: Indicates the date
    - HHMMSS: Indicates the time
- There is no way to check restore points you have set. Keep a record in, for example, a file.



### See

Refer to "System Administration Functions" under "Functions and Operators" in the PostgreSQL Documentation for information on pg\_create\_restore\_point.

......

# Chapter 4 Configuring Secure Communication Using Secure Sockets Layer

If communication data transferred between a client and a server contains confidential information, encrypting the communication data can protect it against threats, such as eavesdropping on the network.

### 4.1 Configuring Communication Data Encryption

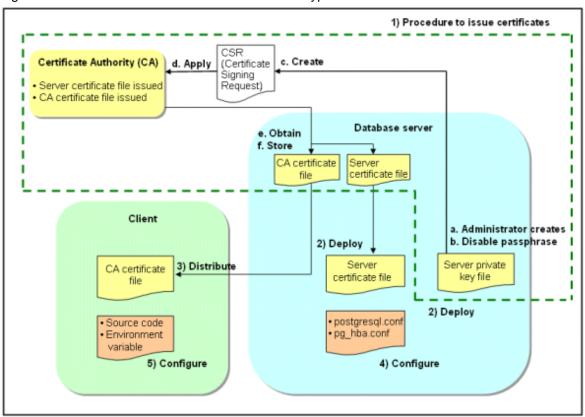
To encrypt communication data transferred between a client and a server, configure communication data encryption as described below. Communication data encryption not only protects the communication content, but it also guards against man-in-the-middle (MITM) attacks (for example, data and password theft through server impersonation).

Table 4.1 Configuration procedure

Configuration procedure
1) Issue a certificate
2) Deploy a server certificate file and a server private key file
3) Distribute a CA certificate file to the client
4) Configure the operating environment for the database server
5) Configure the operating environment for the client

The following figure illustrates the environment for communication data encryption.

Figure 4.1 Environment for communication data encryption



### 4.1.1 Issuing a Certificate

For authenticating servers, you must acquire a certificate issued by the certificate authority (CA).

FUJITSU Enterprise Postgres supports X.509 standard PEM format files. If the certificate authority issues a file in DER format, use a tool such as the openssl command to convert the DER format file to PEM format.

The following provides an overview of the procedure. Refer to the procedure published by the public or independent certificate authority (CA) that provides the certificate file for details.

- a. Create a server private key file
- b. Disable the passphrase for the server private key file
- c. Create a CSR (signing request for obtaining a server certificate) from the server private key file
- d. Apply to the certificate authority (CA) for a server certificate
- e. Obtain a server certificate file and a CA certificate file from the certificate authority (CA)
- f. Store the server certificate file and the CA certificate file Note: If you lose or destroy the certificates, you will need to have them re-issued.

The above procedure enables you to prepare the following files:

- Server private key file
- Server certificate file
- CA certificate file

### 4.1.2 Deploying a Server Certificate File and a Server Private Key File

Create a directory on the local disk of the database server and store the server certificate file and the server private key file in it. Use the operating system features to set access privileges for the server certificate file and the server private key file so that only the database administrator has load privileges.

Back up the server certificate file and the server private key file in the event that data corruption occurs and store them securely.

### 4.1.3 Distributing a CA Certificate File to the Client

Create a directory on the local disk of the client and place the distributed CA certificate file there. Use the operating system features to set load privileges to protect the CA certificate file against accidental deletion.

# 4.1.4 Configuring the Operating Environment for the Database Server



Refer to "Secure TCP/IP Connections with SSL" under "Server Administration" in the PostgreSQL Documentation for details.

# 4.1.5 Configuring the Operating Environment for the Client



Refer to the following sections in the Application Development Guide for details, depending on your application development environment:

- "Settings for Encrypting Communication Data" under "Setup" in "JDBC Driver"
- "Settings for Encrypting Communication Data" under "Setup" in "C Library (libpq)"
- "Settings for Encrypting Communication Data" under "Setup" in "Embedded SQL in C"

### 4.1.6 Performing Database Multiplexing

When you perform communication that uses database multiplexing and a Secure Socket Layer server certificate, certificates with the same "Common Name" must be used. To ensure this, take one of the following actions:

- Create one server certificate, replicate it, and place a copy on each server used for database multiplexing.
- Create a server certificate with the same "Common Name" for each server used for database multiplexing.



Refer to "Using the Application Connection Switch Feature" in the Application Development Guide for information on how to specify applications on the client.

......

# Chapter 5 Protecting Storage Data Using Transparent Data Encryption

This chapter describes how to encrypt data to be stored in the database.

### 5.1 Protecting Data Using Encryption

With PostgreSQL, data in a database is protected from access by unauthorized database users through the use of authentication and access controls. However, the OS file is not protected from attackers who bypass the database server's authentication and access controls.

With FUJITSU Enterprise Postgres, data inside the OS file is encrypted, so valuable information is protected even if the file or disk is stolen.

Data to be stored in a database is encrypted when it is written to the data file, and decrypted when it is read.

This is performed automatically by the instance, so the user and the application need not be aware of key management and encryption or decryption. This process is called TDE (Transparent Data Encryption).

The characteristics of TDE are described below.

#### **Encryption mechanisms**

#### Two-layer encryption key and the keystore

In each tablespace, there is a tablespace encryption key that encrypts and decrypts all the data within. The tablespace encryption key is encrypted by the master encryption key and saved.

Only one master encryption key exists in a database cluster. It is encrypted based on a passphrase specified by the user and stored in a keystore. FUJITSU Enterprise Postgres provides a file-based keystore. Attackers who do not know the passphrase cannot read the master encryption key from the keystore.

#### Strong encryption algorithms

TDE uses the Advanced Encryption Standard (AES) as its encryption algorithm. AES was adopted as a standard in 2002 by the United States Federal Government, and is used throughout the world.

#### Faster encryption and decryption based on hardware

TDE minimizes the overhead of encryption and decryption by using the AES-NI (Advanced Encryption Standard New Instructions) built into Intel(R) Xeon(R) processors since the 5600 series. This means that even in situations where previously the minimum encryption target was selected as a tradeoff between performance and security, it is now possible to encrypt all the data of an application.

Refer to the Intel Corporation's website for information on the list of processors equipped with AES-NI.

#### Zero overhead storage areas

Encryption does not change the size of data stored in tables, indexes, or WAL. There is, therefore, no need for additional estimates or disks

### Scope of encryption

#### All user data within the specified tablespace

The tablespace is the unit for specifying encryption. All tables, indexes, temporary tables, and temporary indexes created in the encrypted tablespace are encrypted. There is no need for the user to consider which tables and strings to encrypt.

Refer to "5.4 Encrypting a Tablespace" for details.

### Backup data

The pgx\_dmpall command and pg\_basebackup command create backup data by copying the OS file. Backups of the encrypted data are, therefore, also encrypted. Information is protected from leakage even if the backup medium is stolen.

#### WAL and temporary files

WAL, which is created by updating encrypted tables and indexes, is encrypted with the same security strength as the update target. When large merges and sorts are performed, the encrypted data is written to a temporary file in encrypted format.

#### Streaming replication support

You can combine streaming replication and transparent data encryption. The data and WAL encrypted on the primary server is transferred to the standby server in its encrypted format and stored.



The following are not encrypted:

- pg\_dump and pg\_dumpall output files
- Files output by the COPY command
- Notification event payloads that communicate using the LISTEN or NOTIFY command
- Checksum validation is not performed on encrypted tablespaces during backup and when using the pg\_checksum utility.

### 5.2 Setting the Master Encryption Key

To use transparent data encryption, you must create a keystore and set the master encryption key.

1. In the keystore\_location parameter of postgresql.conf, specify the directory to store the keystore.

Specify a different location for each database cluster.

```
keystore_location = 'C:\\key\\store\\location'
```

Refer to "Appendix A Parameters" for information on postgresql.conf.

After editing the postgresql.conf file, either start or restart the instance.

- Using WebAdmin

Refer to "2.1.1 Using WebAdmin", and restart the instance.

- Using commands

Refer to "2.1.2 Using Commands", and restart the instance.

2. Execute an SQL function, such as the one below, to set the master encryption key. This must be performed by the superuser. Execute it as the database superuser.

```
SELECT pgx_set_master_key('passphrase');
```

The value "passphrase" is the passphrase that will be used to open the keystore. The master encryption key is protected by this passphrase, so avoid specifying a short simple string that is easy to guess.

Refer to "B.2 Transparent Data Encryption Control Functions" for information on the pgx\_set\_master\_key function.



Note that if you forget the passphrase, you will not be able to access the encrypted data. There is no method to retrieve a forgotten passphrase and decrypt data. Do not, under any circumstances, forget the passphrase.

The pgx\_set\_master\_key function creates a file with the name keystore.ks in the keystore storage destination. It also creates a master encryption key from random bit strings, encrypts it with the specified passphrase, and stores it in keystore.ks. At this point, the keystore is open.

### 5.3 Opening the Keystore

To create encrypted tablespaces and access the encrypted data, you must first open the keystore. When you open the keystore, the master encryption key is loaded into the database server memory and becomes usable for encryption and decryption.

You need to open the keystore each time you start the instance. To open the keystore, the database superuser must execute the following SOL function.

```
SELECT pgx_open_keystore('passphrase');
```

The value "passphrase" is the passphrase specified during creation of the keystore.

Refer to "B.2 Transparent Data Encryption Control Functions" for information on the pgx\_open\_keystore function.

Note that, in the following cases, the passphrase must be entered when starting the instance, because the encrypted WAL must be decrypted for recovery. In this case, the above-mentioned pgx\_open\_keystore function cannot be executed.

- If performing crash recovery at the time of starting the instance
- If performing recovery using continuous archiving

For the above cases, select one of the following methods:

- Use an automatically opening keystore

Select this method if ease of operation has priority over enhanced security. When using an automatically opening keystore, the content of the keystore file is decrypted and a copy of the keystore file is generated. Although the content of this file is obfuscated, the level of security becomes slightly weaker.

Select this method if performing operations using WebAdmin.

- Enter a passphrase when starting an instance

Select this method if enhanced security has priority over ease of operation.

Specify the --keystore-passphrase in the pg\_ctl command and start the instance. This displays the prompt that asks for the passphrase to be entered.

```
> pg_ctl --keystore-passphrase start
Enter the passphrase:
The server is currently initiating
>
```

After performing the above operation, use the pg\_ctl command to stop the instance.

Then start the instance in Windows services. Refer to "2.1.2 Using Commands" for information on how to start an instance in Windows services.



When using an automatically opening keystore, you do not need to enter the passphrase and you can automatically open the keystore when the database server starts. Refer to "5.6.3 Enabling Automatic Opening of the Keystore" for details.

# 5.4 Encrypting a Tablespace

The keystore must be open before you can create an encrypted tablespace.

When creating a tablespace that will be encrypted, configure the encryption algorithm in the runtime parameters. For example, to create a tablespace with the name secure\_tablespace using AES with a key length of 256 bits as the encryption algorithm, configure as shown below.

```
-- Specify the encryption algorithm for the tablespace to be created below SET tablespace_encryption_algorithm = 'AES256';
CREATE TABLESPACE secure_tablespace LOCATION 'C:\My\Data\Dir';
-- Specify that the tablespace to be created below is not to be encrypted SET tablespace_encryption_algorithm = 'none';
```

Or

```
CREATE TABLESPACE secure_tablespace LOCATION '\My\Data\Dir' WITH (tablespace_encryption_algorithm = 'AES256' );
```

When the tablespace is empty, the encryption algorithm can be modified with the command below.

```
ALTER TABLESPACE secure_tablespace SET (tablespace_encryption_algorithm=AES256);
```

Trying to set the encryption algorithm for a non-empty tablespace causes an error.

You can use AES with a key length of 128 bits or 256 bits as the encryption algorithm. It is recommended that you use 256-bit AES. Refer to "Appendix A Parameters" for information on how to specify the runtime parameters.

If user provides both GUC and command line options while creating the tablespace, the preference is given to the command line option.

The pg\_default and pg\_global tablespaces cannot be encrypted.

Create tables and indexes in the encrypted tablespace that you created. Relations created in the encrypted tablespace are automatically encrypted.



#### Example 1: Specifying an encrypted tablespace when creating it

```
CREATE TABLE my_table (...)

TABLESPACE secure_tablespace;
```

#### Example 2: Not explicitly specifying a tablespace when creating it and instead using the default tablespace

```
SET default_tablespace = 'secure_tablespace';
CREATE TABLE my_table (...);
```

The process is the same for encrypting temporary tables and temporary indexes. In other words, either explicitly specify the TABLESPACE clause or list encrypted tablespaces in the temp\_tablespaces parameter, and then execute CREATE TEMPORARY TABLE or CREATE INDEX



If an encrypted tablespace is specified in the TABLESPACE clause of the CREATE DATABASE statement, relations created in the database without explicitly specifying a tablespace will be encrypted. Furthermore, the system catalog will also be encrypted, so the source code of user-defined functions is also protected.

Example: Specifying a tablespace in a database definition statement

```
CREATE DATABASE DB01 TABLESPACE=SP01 ...;
```

Part of the data is also stored in the system catalog - to encrypt this data as well, specify an encrypted tablespace as above and create a database.



An encrypted tablespace cannot be created from the window used for creating the pgAdmin tablespace, or from the query tool. To create an encrypted tablespace, click [PSQL Console] from the [Plugins] menu and create an encrypted tablespace in the psql console window.

# 5.5 Checking an Encrypted Tablespace

The pgx\_tablespaces system view displays information about whether each tablespace has been encrypted, and about the encryption algorithm. Refer to "C.1 pgx\_tablespaces" for information on strings.

You can discover which tablespaces have been encrypted by executing the following SQL statements.

However, when considering continued compatibility of applications, do not reference system catalogs (pg\_tablespace) directly in SQL statements

```
SELECT spcname, spcencalgo

FROM pg_tablespace ts, pgx_tablespaces tsx

WHERE ts.oid = tsx.spctablespace;
```





Refer to "Notes on Application Compatibility" in the Application Development Guide for information on how to maintain application compatibility.

### 5.6 Managing the Keystore

This section describes how to manage the keystore and the master encryption key to guard against the threat of theft.

### 5.6.1 Changing the Master Encryption Key

Using the same encryption key for an extended period gives attackers an opportunity to decipher the encrypted data. It is recommended that you change the key at regular intervals, or whenever the key is exposed to risk.

Adhere to the industry's best practices for encryption algorithms and key management when considering how often the key should be changed. For example, the NIST in the United States has published "NIST Special Publication 800-57". The PCI DSS also refers to this publication. This publication recommends changing the master encryption key once a year.

To change the master encryption key, execute the pgx\_set\_master\_key function, which is the same function used for configuring the key. Refer to "5.2 Setting the Master Encryption Key" for details.

After changing the master encryption key, you must immediately back up the keystore.

### 5.6.2 Changing the Keystore Passphrase

In security policies for organizations, it is usually a requirement that the passphrase be changed whenever a security administrator who knows the passphrase is removed from duties due to transfer or retirement. It is also recommended that the passphrase be changed if it is ever exposed to risks due to deception such as social engineering.

To change the keystore passphrase, execute the following SQL function as a superuser.

```
SELECT pgx_set_keystore_passphrase('oldPassphrase', 'newPassphrase');
```

After changing the passphrase, you must immediately back up the keystore.

Refer to "B.2 Transparent Data Encryption Control Functions" for information on the pgx\_set\_keystore\_passphrase function.

### 5.6.3 Enabling Automatic Opening of the Keystore

When using an automatically opening keystore, you do not need to enter the passphrase and you can automatically open the keystore when the instance starts. Execute the pgx\_keystore command to enable automatic opening of the keystore.

```
> pgx_keystore --enable-auto-open C:\key\store\location\keystore.ks
Enter the passphrase:
Automatic opening of the keystore is now enabled
>
```



Refer to "pgx\_keystore" in the Reference for information on pgx\_keystore command.

When automatic opening is enabled, an automatically opening keystore is created in the same directory as the original keystore. The file name of the automatically opening keystore is keystore.aks. The file keystore.aks is an obfuscated copy of the decrypted content of the keystore.ks file. As long as this file exists, there is no need to enter the passphrase to open the keystore when starting the instance.

Do not delete the original keystore file, keystore.ks. It is required for changing the master encryption key and the passphrase. When you change the master encryption key and the passphrase, keystore.aks is recreated from the original keystore file, keystore.ks.

Protect keystore.ks, keystore.aks, and the directory that stores the keystore so that only the user who starts the instance can access them.

Configure the permission of the files so that only the user who starts the instance can access the SQL functions and commands that create these files. Accordingly, manually configure the same permission mode if the files are restored.

Set the permission mode in [Properties] in Windows(R) Explorer.



Refer to [Help and Support] in Windows(R) for information on [Properties].

An automatically opening keystore will only open on the computer where it was created.

To disable automatic opening of the keystore, delete keystore.aks.



- To use WebAdmin for recovery, you must enable automatic opening of the keystore.
- Refer to "5.7 Backing Up and Restoring/Recovering the Database" after enabling or reconfiguring encryption to back up the database.
- Specify a different directory from those below as the keystore storage destination:
  - Data storage destination
  - Tablespace storage destination
  - Transaction log storage destination
  - Backup data storage destination

# 5.6.4 Backing Up and Recovering the Keystore

Back up the keystore at the following times in case it is corrupted or lost. Note that you must store the database and the keystore on separate data storage media. Storing both on the same data storage medium risks the danger of the encrypted data being deciphered if the medium is stolen. A passphrase is not required to open an automatically opening keystore, so store this type of keystore in a safe location.

- When the master encryption key is first configured
- When the master encryption key is changed
- When the database is backed up
- When the keystore passphrase is changed



Do not overwrite an old keystore when backing up a keystore. This is because during database recovery, you must restore the keystore to its state at the time of database backup. When the backup data of the database is no longer required, delete the corresponding keystore.



- Back up the database and the keystore on May 1, 2017.

```
> pgx_dmpall -D D:\database\inst1
> copy C:\key\store\location\keystore.ks C:\keybackup\keystore_20170501.ks
```

Specify the following in the pgx\_dmpall command:

- Specify the data storage destination in the -D option. If the -D option is omitted, the value of the PGDATA environment variable is used by default.
- Change the master encryption key, and back up the keystore on May 5, 2017.

```
> psql -c "SELECT pgx_set_master_key('passphrase')" postgres
> copy C:\key\store\location\keystore.ks C:\keybackup\keystore_20170505.ks
```

Specify the following in the psql command:

- Specify the SQL function that sets the master encryption key in the -c option.
- Specify the name of the database to be connected to as the argument.

If the keystore is corrupted or lost, restore the keystore containing the latest master encryption key. If there is no keystore containing the latest master encryption key, restore the keystore to its state at the time of database backup, and recover the database from the database backup. This action recovers the keystore to its latest state.



- Restore the keystore containing the latest master encryption key as of May 5, 2017.

```
> copy C:\keybackup\keystore_20170505.ks C:\key\store\location\keystore.ks
```

- If there is no backup of the keystore containing the latest master encryption key, recover the keystore by restoring the keystore that was backed up along with the database on 1 May 2017.

```
> copy C:\keybackup\keystore_20170501.ks C:\key\store\location\keystore.ks
> pgx_rcvall -B E:\backup\instl -D:\database\instl --keystore-passphrase
```

Specify the following in the pgx\_rcvall command:

- Specify the data storage directory in the -D option. If the -D option is omitted, the value of the PGDATA environment variable is used by default.
- Specify the backup data storage directory in the -B option.
- The --keystore-passphrase option prompts you to enter the passphrase to open the keystore.

If you have restored the keystore, repeat the process of enabling automatic opening of the keystore. This ensures that the contents of the automatically opening keystore (keystore.aks) are identical to the contents of the restored keystore.

It is recommended that you do not back up the automatically opening keystore file, keystore.aks. If the database backup medium and the backup medium storing the automatically opening keystore are both stolen, the attacker will be able to read the data even without knowing the passphrase.

If the automatically opening keystore is corrupted or lost, you must again enable automatic opening. The keystore.aks file will be recreated from keystore.ks at this time.



Refer to "pgx\_rcvall" and "pgx\_dmpall" in the Reference for information on the pgx\_rcvall and pgx\_dmpall commands.

Refer to "psql" under "Reference" in the PostgreSQL Documentation for information on the psql command.

Refer to "B.2 Transparent Data Encryption Control Functions" for information on the pgx\_set\_master\_key function.

Refer to "5.6.3 Enabling Automatic Opening of the Keystore" for information on how to enable automatic opening of the keystore.

# 5.7 Backing Up and Restoring/Recovering the Database

FUJITSU Enterprise Postgres enables you to use the five backup and recovery methods described below. Regardless of the method you use, you must back up the keystore at the same time.

Note that you must store the database and the keystore on separate data storage media. Storing both on the same data storage medium risks the danger of the encrypted data being deciphered if the medium is stolen.

### Backup and recovery using WebAdmin

- Backup

WebAdmin backs up encrypted data.

Back up the key store after backing up the database.

- Recovery

Restore the keystore to its state at the time of database backup. Refer to "5.6.4 Backing Up and Recovering the Keystore" for details.

Enable automatic opening of the keystore in accordance with the procedure described in "5.6.3 Enabling Automatic Opening of the Keystore". Then, use WebAdmin to recover the database.

#### Backup and recovery using the pgx\_dmpall and pgx\_rcvall commands

Backup

The pgx\_dmpall command backs up the encrypted data.

Back up the key store after backing up the database.

- Recovery

Restore the keystore to its state at the time of the database backup.

Configure automatic opening of the key store as necessary.

If automatic opening of the keystore is not enabled, execute the pgx\_rcvall command with the --keystore-passphrase option specified. This will display the prompt for the passphrase to be entered.



- Back up the database and the keystore on May 1, 2017.

> pgx\_dmpall -D D:\database\inst1
> copy C:\key\store\location\keystore.ks C:\keybackup\keystore\_20170501.ks

Specify the following in the pgx\_dmpall command:

- Specify the data storage destination in the -D option. If the -D option is omitted, the value of the PGDATA environment variable is used by default.

- Recover the database and the keystore from the backup taken on May 1, 2017.

```
> copy C:\keybackup\keystore_20170501.ks C:\key\store\location\keystore.ks
> pgx_keystore --enable-auto-open C:\key\store\location\keystore.ks (Execute only when enabling automatic opening)
> pgx_rcvall -B E:\backup\inst1 -D D:\database\inst1 --keystore-passphrase
```

Specify the following in the pgx\_rcvall command:

- Specify the data storage destination in the -D option. If the -D option is omitted, the value of the PGDATA environment variable is used by default.
- Specify the backup data storage directory in the -B option.
- The --keystore-passphrase option prompts you to enter the passphrase to open the keystore.

### **Dump and restore using SQL**

- Backup

The files output by the pg\_dump and pg\_dumpall commands are not encrypted. You should, therefore, encrypt the files using OpenSSL commands or other means before saving them, as described in "5.8 Importing and Exporting the Database" below.

Back up the key store after backing up the database.

- Restore

If the backup data has been encrypted using, for example Open SSL commands, decrypt that data.

The data generated by the pg\_dumpall command includes a specification to encrypt tablespaces by default. For this reason, the psql command encrypts tablespaces during restoration.

#### File system level backup and restore

- Backup

Stop the instance and backup the data directory and the tablespace directory using the file copy command of the operating system. The files of encrypted tablespaces are backed up in the encrypted state.

Back up the key store after performing the backup.

- Restore

Restore the keystore to its state at the time of the database backup.

Stop the instance and restore the data directory and the tablespace directory using the file copy command of the operating system.

### Continuous archiving and point-in-time recovery

- Backup

The pg\_basebackup command backs up the encrypted data as is.

Back up the key store after performing the backup.

- Recovery

Restore the keystore to its state at the time of the database backup.

Configure automatic opening of the key store as necessary.

If automatic opening of the keystore is not enabled, refer to "5.3 Opening the Keystore" for information on starting an instance by specifying pg\_ctl --keystore-passphrase start.



- Refer to "Reference" in the PostgreSQL Documentation for information on the following commands:
  - psql
  - pg\_dump
  - pg\_basebackup
- Refer to the Reference for information on the following commands:
  - pgx\_rcvall
  - pgx\_dmpall
  - pg\_dumpall

If you have restored the keystore, repeat the process of enabling automatic opening of the keystore This ensures that the contents of the automatically opening keystore (keystore.aks) are identical to the contents of the restored keystore.

Refer to "5.6.3 Enabling Automatic Opening of the Keystore" for information on how to enable automatic opening of the keystore.

### 5.8 Importing and Exporting the Database

The files output by the COPY TO command are not encrypted. Therefore, when transferring files to other systems, you should encrypt files using OpenSSL commands, or use file transfer software that performs encrypted communication for Windows, to encrypt the data being transferred

Use a safe method to delete obsolete plain text files.

You can use the following methods to safely delete files:

- fsutil command



```
# Export the contents of the table my_table to a CSV file.

> psql -c "COPY my_table TO 'C:\WINDOWS\Temp\my_table.csv' (FORMAT CSV)" postgres

# Encrypt the exported file.

> C:\OpenSSL-Win32\bin\openssl enc -e -aes256 -in C:\WINDOWS\Temp\my_table.csv -out my_table.csv.enc

(The user is prompted to enter the passphrase to be used for encryption)

# Check the size of plain text files, and delete them after zero padding

> dir C:\WINDOWS\Temp\my_table.csv

> fsutil file setzerodata offset=0 length=7 C:\WINDOWS\Temp\my_table.csv

> del C:\WINDOWS\Temp\my_table.csv

# Decrypt the encrypted files on other systems.

> C:\OpenSSL-Win32\bin\openssl enc -d -aes256 -in my_table.csv.enc -out my_table.csv

(The user is prompted to enter the passphrase to be used for decryption)
```

If you use COPY FROM to import data to tables and indexes in an encrypted tablespace, the imported data is automatically encrypted before being stored.

# 5.9 Encrypting Existing Data

You cannot encrypt existing unencrypted tablespaces. In addition, you cannot change encrypted tablespaces so that they do not encrypt.

As an alternative, transfer the tables and indexes to other tablespaces. You can use the following SQL commands for this.

ALTER TABLE table\_name SET TABLESPACE new\_tablespace;
ALTER INDEX index\_name SET TABLESPACE new\_tablespace;
ALTER DATABASE database\_name SET TABLESPACE new\_tablespace;



Refer to "SQL Commands" under "Reference" in the PostgreSQL Documentation for information on SQL commands.

# 5.10 Operations in Cluster Systems

This section describes how to use transparent data encryption on cluster systems such as high-availability systems, streaming replication, and database multiplexing.

### 5.10.1 HA Clusters that do not Use Database Multiplexing

Take the following points into account when using transparent data encryption in an HA cluster environment that does not use database multiplexing.

### Placement and automatic opening of the keystore file

There are two alternatives for placing the keystore file:

- Sharing the keystore file
- Placing a copy of the keystore file

#### Sharing the keystore file

This involves using the same keystore file on the primary server and the standby server.

As the standby server is not active while the primary server is running, this file would not be accessed simultaneously, and therefore, it can be shared.

To manage the keystore file in a more secure manner, place it on the key management server or the key management storage isolated in a secure location.

Enable the automatic opening of the keystore on both the primary and standby servers.

#### Placing a copy of the keystore file

This involves placing a copy of the primary server keystore file on the standby server.

You can do this if you cannot prepare a shared server or disk device that can be accessed from both the primary and standby servers.

However, if you change the master encryption key and the passphrase on the primary server, you must copy the keystore file to the standby server again.

To manage the keystore file in a more secure manner, prepare the key management server or the key management storage isolated in a secure location for both the primary and standby servers, and place the keystore files there.

Enable the automatic opening of the keystore on both the primary and standby servers. Note that copying the automatically opening keystore file (keystore.aks) to the standby server does not enable the automatic opening of the keystore.

### 5.10.2 Database Multiplexing Mode

Note the following when using transparent data encryption in environments that use streaming replication, or database multiplexing with streaming replication.

#### Placing the keystore file

Place a copy of the primary server keystore file on the standby server.

This is required as the keystore file cannot be shared, and both servers may need to access it simultaneously.



To manage the keystore file in a more secure manner, place it on the key management server or the key management storage isolated in a secure location. A keystore used by both the primary and standby servers can be managed on the same key management server or key management storage.

However, create different directories for the keystores to be used by the primary server and the standby server. Then copy the keystore for the primary server to the directory used on the standby server.

### Automatically opening the keystore

You must enable automatic opening of the keystore.

To do this, enable automatic opening of the keystore in all servers that make up database multiplexing. The settings for automatic opening of the keystore include information unique to each server, so simply copying the file does not enable it.

### Changing the passphrase

Changes to the passphrase are reflected in all servers that make up database multiplexing, so no special operation is required.

### Building and starting a standby server

Before using the pg\_basebackup command or pgx\_rcvall command to build a standby server, copy the keystore file from the primary server to the standby server. When using an automatically opening keystore, use the copied keystore file to enable automatic opening on the standby server.

Open the keystore each time you start the standby server. This step is necessary for decrypting and restoring encrypted WAL received from the primary server. To open the keystore, specify the --keystore-passphrase option in the pg\_ctl command or pgx\_rcvall command and enter the passphrase, or use an automatically opening keystore.

If specifying --keystore-passphrase in the pg\_ct command, refer to "5.3 Opening the Keystore" for details.

#### Changing the master encryption key and the passphrase

Change the master encryption key and the passphrase on the primary server. You need not copy the keystore from the primary server to the standby server. You need not even restart the standby server or reopen the keystore. Changes to the master encryption key and the passphrase are reflected in the keystore on the standby server.



Refer to "pgx\_rcvall" in the Reference for information on pgx\_rcvall command.

Refer to "pg\_basebackup" under "Reference" in the PostgreSQL Documentation for information on pg\_basebackup command.

Refer to "High Availability, Load Balancing, and Replication" under "Server Administration" in the PostgreSQL Documentation for information on how to set up streaming replication.

### 5.11 Security-Related Notes

- Decrypted data is cached in the database server memory (shared buffer). As a result, unencrypted data is stored in a minidump, which is the process memory dump. You should, therefore, safely delete the memory dump. You can safely delete files by using the following command:
  - fsutil command
- Unencrypted data may be written from the database server memory to the operating system's swap area. To prevent leakage of information from the swap area, consider either disabling the use of swap area or encrypting the swap area using a full-disk encryption product.
- The content of the server log file is not encrypted. Therefore, in some cases the value of a constant specified in a SQL statement is output to the server log file. To prevent this, consider setting a parameter such as log\_min\_error\_statement.

- When executing an SQL function that opens the keystore and modifies the master encryption key, ensure that the SQL statement containing the passphrase is not output to the server log file. To prevent this, consider setting a parameter such as log\_min\_error\_statement. If you are executing this type of SQL function on a different computer from the database server, encrypt the communication between the client and the database server with SSL.
- Starting with FEP 10, logical replication is available, which allows non-backed up clusters to subscribe to databases where transparent data encryption is enabled. Logical replication does not need to have the same encryption strategy between publisher and subscriber.

In this scenario, if the user wants to encrypt the subscribed copy of data as well, then it is the user's responsibility to create encryption policies to the subscribed databases. By default, published encrypted tablespace data will not be encrypted in the subscriber side.

### 5.12 Tips for Installing Built Applications

With transparent data encryption, you can easily encrypt all the data in an application without modifying the application. Database administrators install built applications in the following manner. However, this procedure stores data to the default tablespace, so take necessary action if processing differs from the original design.

1. (Normal procedure) Create an owner and a database for the built application.

```
CREATE USER crm_admin ...;
CREATE DATABASE crm_db ...;
```

2. (Procedure for encryption) Create an encrypted tablespace to store the data for the built application.

```
SET tablespace_encryption_algorithm = 'AES256';
CREATE TABLESPACE crm_tablespace LOCATION 'C:\crm\data';
```

3. (Procedure for encryption) Configure an encrypted tablespace as the default tablespace for the owner of the built application.

```
ALTER USER crm_admin SET default_tablespace = 'crm_tablespace';
ALTER USER crm_admin SET temp_tablespaces = 'crm_tablespace';
```

4. (Normal procedure) Install the built application. The application installer prompts you to enter the host name and the port number of the database server, the user name, and the database name. The installer uses the entered information to connect to the database server and execute the SQL script. For applications that do not have an installer, the database administrator must manually execute the SQL script.

Normally, the application's SQL script includes logic definition SQL statements, such as CREATE TABLE, CREATE INDEX, and GRANT or REVOKE, converted from the entity-relationship diagram. It does not include SQL statements that create databases, users, and tablespaces. Configuring the default tablespace of the users who will execute the SQL script deploys the objects generated by the SQL script to the tablespace.

# Chapter 6 Data Masking

Data masking is a feature that can change the returned data for queries generated by applications, so that it can be referenced by users. For example, for a query of employee data, digits except the last four digits of an eight-digit employee number can be changed to "\*" so that it can be used for reference.

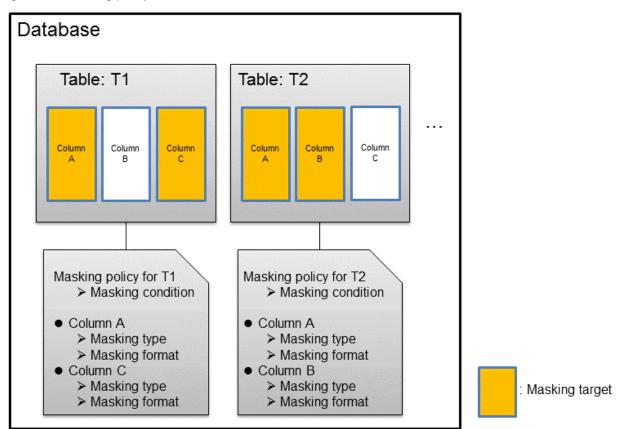


When using this feature, it is recommended that the changed data be transferred to another medium for users to reference. This is because, if users directly access the database to extract the masked data, there is a possibility that they can deduce the original data by analyzing the masking policy or query result to the masking target column.

# 6.1 Masking Policy

Masking policy is a method of changing data under specific conditions when it is returned for a query from an application. One masking policy can be created per table. You can configure masking target, masking type, masking condition and masking format in a masking policy.

Figure 6.1 Masking policy





When a masking policy is defined, the search performance for the corresponding table may deteriorate.

### 6.1.1 Masking Target

Masking target refers to a column to which a masking policy will be applied. When referring to a masking target or a function that includes a masking target, the execution result will be changed and obtained.

The following commands can change the execution result:

- SELECT
- COPY
- pg\_dump
- pg\_dumpall



- If a masking target is specified to INSERT...SELECT target columns, processing will be performed using data before change.
- If a masking target other than SELECT target columns is specified, processing will be performed using data before change.
- If a masking target is specified in a function where the data type will be converted, an error will occur.

### 6.1.2 Masking Type

Masking type is a method to change column data that is returned from queries. Specify the masking type in the function\_type parameter. The following masking types can be specified and selected depending on the masking target data type.

### **Full masking**

All the data in the specified column is changed. The changed value returned to the application that made the query varies depending on the column data type.

For example, 0 is used for a numeric type column and a space is used for a character type column.

### Partial masking

The data in the specified column is partially changed.

For example, digits except the last four digits of an employee number can be changed to "\*".

#### Regular expression masking

The data in the specified column is changed via a search that uses a regular expression.

For example, for strings such as email address that can have variable length, "\*" can be used to change characters preceding "@" by using a regular expression. Regular expression masking can only be used for character type data.



- If multiple valid masking targets are specified for a function, the masking type for the left-most masking target will be applied. For example, if "SELECT GREATEST(c1, c2) FROM t1" is executed for numeric type masking target c1 and c2, the masking type for c1 will be applied.

- When masking the data that includes multibyte characters, do not specify partial masking for masking type. The result may not be as expected.

### 6.1.3 Masking Condition

Masking condition refers to the conditions configured to perform masking. Specify the masking condition in the expression parameter. Changed or actual data can be displayed for different users by defining masking condition. An expression that returns a boolean type result needs to be specified in masking condition and masking is performed only when TRUE is returned. Refer to "Value Expressions" in the PostgreSQL Documentation for information on the expressions that can be specified. Note that expressions that include a column cannot

be specified.

For example, when masking data only for "postgres" users, specify 'current\_user = "postgres" in the masking condition.



Specify '1=1' so the masking condition is always evaluated to be TRUE and masking is performed all the time.

### 6.1.4 Masking Format

Masking format is a combination of change method and displayed characters when the masking condition is met. Masking format varies depending on the masking type. The following describes the masking format.

#### Full masking

With full masking, all characters are changed to values as determined by the database. Changed characters can be referenced in the pgx\_confidential\_values table. Also, replacement characters can be changed using the pgx\_update\_confidential\_values system management function.



Refer to "6.3 Data Types for Masking" for information on the data types for which data masking can be performed.

### Partial masking

With partial masking, data is changed according to the content in the function\_parameters parameter. The method of specifying function\_parameters varies depending on the data type.

Category	Method of specifying function_parameters
Numeric type	'replacementCharacter, startPosition, endPosition'
	- replacementCharacter. Specify the number to display. Specify a value from 0 to 9.
	- startPosition: Specify the start position of masking. Specify a positive integer.
	- <i>endPosition</i> : Specify the end position of masking. Specify a positive integer that is greater than <i>startPosition</i> .
	Example
	Specify as below to change the values from the 1st to 5th digits to 9.
	function_parameters := '9, 1, 5'
	In this example, if the original data is "123456789", it will be changed to "999996789".
Character type	'inputFormat, outputFormat, replacementCharacter, startPosition, endPosition'
	- <i>inputFormat</i> : Specify the current format of the data. Specify "V" for characters that will potentially be masked, and specify "F" for values such as spaces or hyphens that will not be masked.
	- <i>outputFormat</i> : Define the method to format the displayed data. Specify "V" for characters that will potentially be masked. Any character to be output can be specified for each character "F" in <i>inputFormat</i> . If you want to output a single quotation mark, specify two of them consecutively.
	- replacementCharacter. Specify any single character. If you want to output a single quotation mark, specify two of them consecutively.
	- <i>startPosition</i> : Specify the position of "V" as the start position of masking. For example, to specify the position of the 4th "V" from the left, specify 4. Specify a positive integer.

Category	Method of specifying function_parameters
	- <i>endPosition</i> : Specify the position of "V" as an end position of masking. When working out the end position, do not include positions of "F". For example, to specify the position of the 11th "V" from the left, specify 11. Specify a positive integer that is greater than <i>startPosition</i> .
	Example
	Specify as below to mask a telephone number other than the first three digits using *.
	function_parameters := 'VVVFVVVVVVVVVVVVVVVVVVVVVVVVVVVVVVVVV
	In this example, if the original data is "012-3156-7890", it will be changed to "012-****.
Date/timestamp type	'MDYHMS'
Date/timestamp type	- M: Masks month. To mask month, enter the month from 1 to 12 after a lowercase letter m. Specify an uppercase letter M to not mask month.
	- D: Masks date. To mask date, enter the date from 1 to 31 after a lowercase letter d. If a value bigger than the last day of the month is entered, the last day of the month will be displayed. Specify an uppercase letter D to not mask date.
	- Y: Masks year. To mask year, enter the year from 1 to 9999 after a lowercase letter y. Specify an uppercase letter Y to not mask year.
	- H: Masks hour. To mask hour, enter the hour from 0 to 23 after a lowercase letter h. Specify an uppercase letter H to not mask hour.
	- M: Masks minute. To mask minute, enter the minute from 0 to 59 after a lowercase letter m. Specify an uppercase letter M to not mask minute.
	- S: Masks second. To mask second, enter the second from 0 to 59 after a lowercase letter s. Specify an uppercase letter S to not mask second.
	Example
	Specify as below to mask hour, minute, and second and display 00:00:00.
	function_parameters := 'MDYh0m0s0'
	In this example, if the original data is "2010-10-10 10:10:10", it will be changed to "2010-10-10 00:00:00".



- Refer to "B.3.2 pgx\_create\_confidential\_policy" for information on function\_parameters.
- Refer to "6.3 Data Types for Masking" for information on the data types for which masking can be performed.

### Regular expression masking

With regular expression masking, data is changed according to the content of the regexp\_pattern, regexp\_replacement and regexp\_flags parameters. For regexp\_pattern, specify the search pattern using a regular expression. For regexp\_replacement, specify the replacement character to use when data matches the search pattern. For regexp\_flags, specify the regular expression flags.

.....

••••••



Specify as below to change all three characters starting from b to X.

regexp\_pattern := 'b..'

regexp\_replacement:= 'X'

regexp\_flags := 'g'

In this example, if the original data is "foobarbaz", it will be changed to "fooXX".



#### See

- Refer to "POSIX Regular Expressions" in the PostgreSQL Documentation and check pattern, replacement, and flags for information on the values that can be specified for regexp\_pattern, regexp\_replacement, and regexp\_flags.
- Refer to "6.3 Data Types for Masking" for information on the data types for which masking can be performed.



- When column data type is character(n) or char(n) and if the string length after change exceeds n, the extra characters will be truncated and only characters up to the nth character will be displayed.
- When column data type is character varying(*n*) or varchar(*n*) and if the string length after change exceeds the length before the change, the extra characters will be truncated and only characters up to the length before change will be displayed.

### 6.2 Usage Method

### Preparation

The following preparation is required to use this feature.

- 1. Set the postgresql.conf file parameters.
  - Prepend "pgx\_datamasking" to the shared\_preload\_libraries parameter.
- 2. Restart the instance.
- 3. Execute CREATE EXTENSION for the database that will use this feature.

The target database is described as "postgres" here.

Use the psql command to connect to the "postgres" database.



#### Example

postgres=# CREATE EXTENSION pgx\_datamasking;
CREATE EXTENSION



You must always prepend "pgx\_datamasking" to the "shared\_preload\_libraries" parameter.



- Specify "false" for pgx\_datamasking.enable to not use this feature. Data will not be masked even if a masking policy is configured. This feature becomes available again once "true" is specified for pgx\_datamasking.enable. This setting can be made

by specifying a SET statement or specifying a parameter in the postgresql.conf file. Example

```
postgres=# SET pgx_datamasking.enable=false;
```

- Hereafter, also perform this preparatory task for the "template1" database, so that this feature can be used by default when creating a new database.

#### Usage

To perform masking, a masking policy needs to be configured. The masking policy can be created, changed, confirmed, enabled, disabled or deleted during operation.

The procedures to perform these tasks are explained below with examples.

- 1. Creating a masking policy
- 2. Changing a masking policy
- 3. Confirming a masking policy
- 4. Enabling and disabling a masking policy
- 5. Deleting a masking policy



Only database superusers can configure masking policies.

### 6.2.1 Creating a Masking Policy

An example of the operation on the server is shown below.

1. Create a masking policy

Execute the pgx\_create\_confidential\_policy system management function to create a masking policy.

The following values are configured in this example.

- Masking target: Numeric type c1
- Masking type: FULL
- Masking condition: '1=1'

```
postgres=# select pgx_create_confidential_policy(table_name := 't1', policy_name := 'p1',
expression := '1=1', column_name := 'c1', function_type := 'FULL');
pgx_create_confidential_policy
-----t
t
(1 row)
```

2. Confirm the displayed data

Confirm that the masking target data (column c1) has been correctly changed.



- Refer to "B.3.2 pgx\_create\_confidential\_policy" for information on the pgx\_create\_confidential\_policy system management function.



- Only one masking policy can be created per table.
- All users can view the masking policy created, so do not grant the login privilege of the database where this feature is set to the users who refer to the changed data. Masking policies are defined in the "pgx\_confidential\_columns", "pgx\_confidential\_policies" and "pgx\_confidential\_values" tables.

### 6.2.2 Changing a Masking Policy

- 1. An example of the operation on the server is shown below.
- 2. Change a masking policy

Execute the pgx\_alter\_confidential\_policy system management function to change a masking policy.

The following values are changed in this example.

- Content of change: Add a masking target
- Masking target: Character type c2
- Masking type: PARTIAL
- Masking condition: 'VVVFVVVVFVVVV, VVV-VVVV-VVVV, \*, 4, 11'

3. Confirm the displayed data

Confirm that the masking target data has been correctly changed.

```
postgres=# select * from t1;
c1 | c2
----+-----
0 | 012-***-***
0 | 012-***-***
0 | 012-***-***
(3 row)
```



- Refer to "B.3.1 pgx\_alter\_confidential\_policy" for information on the pgx\_alter\_confidential\_policy system management function.

## 6.2.3 Confirming a Masking Policy

An example of the operation on the server is shown below.

Confirm information about a masking target where a masking policy is set
 Refer to the pgx\_confidential\_columns table to confirm the masking target where the masking policy is set.

2. Confirm information about the masking policy content

Refer to pgx\_confidential\_policies to confirm the masking policy content.



- Refer to "D.1 pgx\_confidential\_columns" for information on the pgx\_confidential\_columns table.
- Refer to "D.2 pgx\_confidential\_policies" for information on the pgx\_confidential\_policies table.

### 6.2.4 Enabling and Disabling a Masking Policy

An example of the operation on the server is shown below.

1. Disable a masking policy

Execute the pgx\_enable\_confidential\_policy system management function to disable a masking policy.

```
postgres=# select pgx_enable_confidential_policy(table_name := 't1', policy_name := 'p1',
enable := 'f');
pgx_enable_confidential_policy
-----t
t
(1 row)
```

2. Confirm the displayed data

Confirm that the original data is displayed by disabling the masking policy.

3. Enable a masking policy

Execute the pgx\_enable\_confidential\_policy system management function to enable a masking policy.

```
postgres=# select pgx_enable_confidential_policy(table_name := 't1', policy_name := 'p1',
enable := 't');
pgx_enable_confidential_policy
-----t
(1 row)
```

4. Confirm the displayed data

Confirm that the masking target data has been correctly changed.

```
postgres=# select * from t1;
c1 | c2
---+----
0 | 012-***-***
0 | 012-***-***
```

```
0 | 012-***-***
(3 row)
```



- Refer to "B.3.4 pgx\_enable\_confidential\_policy" for information on the pgx\_enable\_confidential\_policy system management function.

### 6.2.5 Deleting a Masking Policy

An example of the operation on the server is shown below.

1. Delete a masking policy

Execute the pgx\_drop\_confidential\_policy system management function to delete a masking policy.

```
postgres=# select pgx_drop_confidential_policy(table_name := 't1', policy_name := 'p1');
pgx_drop_confidential_policy
-----t
t
(1 row)
```

2. Confirm the displayed data

Confirm that the original data is displayed by deleting the masking policy.



- Refer to "B.3.3 pgx\_drop\_confidential\_policy" for information on the gx\_drop\_confidential\_policy function.

# 6.3 Data Types for Masking

The data types for which data masking can be performed are shown below.

Category	Data type			
		Full masking	Partial masking	Regular expression masking
Numeric type	smallint	Y	Y	N
	integer	Y	Y	N
	bigint	Y	Y	N
	decimal	Y	Y	N
	numeric	Y	Y	N
	float	Y	Y	N
	real	Y	Y	N
	double precision	Y	Y	N

Category	Data type	Masking type		
		Full masking	Partial masking	Regular expression masking
Character type	character varying(n)	Y	Y	Y
	varchar(n)	Y	Y	Y
	character(n)	Y	Y	Y
	char(n)	Y	Y	Y
Date/timestamp type	date	Y	Y	N
	timestamp	Y	Y	N

# **6.4 Security Notes**

- Starting with FEP 10, logical replication is available, which allows non-backed up clusters to subscribe to databases where data masking policies are enabled. Logical replication allows publisher and subscriber databases to have their own or the same data masking policies.
  - In this scenario, the user must disable data masking on the publisher database whenever a subscription is created. This ensures that subscribers are able to obtain the original data (initial copy) instead of the masked version. Then, it is the user's responsibility to set masking policies to each subscribed database.
- Take strong caution in publishing data masking's confidential tables (pgx\_confidential\_policies, pgx\_confidential\_columns, etc.) unless the user is publishing all tables of the database and wants to apply the same data masking's policies on the subscribed database for all of them.

Otherwise, as these confidential tables contain the masking policies for all tables of the database, confidential policies of unpublished tables may be unintentionally published. Additionally, it is not possible to apply different data masking policies on the subscriber database.

# Chapter 7 Periodic Operations

This chapter describes the operations that must be performed periodically when running daily database jobs.

## 7.1 Configuring and Monitoring the Log

FUJITSU Enterprise Postgres enables you to output database errors and warnings to a log file.

This information is useful for identifying if errors have occurred and the causes of those errors.

By default, this information is output to the event log. It is recommended that you configure FUJITSU Enterprise Postgres to collect logs from its log files (for example, log\_destination) before operating FUJITSU Enterprise Postgres.

Periodically monitor the log files to check if any errors have occurred.



- Refer to "Error Reporting and Logging" under "Server Administration" in the PostgreSQL Documentation for information on logs.

......

- Refer to "Configuring Parameters" in the Installation and Setup Guide for Server for information on log settings when operating with WebAdmin.

## 7.2 Monitoring Disk Usage and Securing Free Space

When a database is used for an extended period, free space on the disk is continuously consumed and in some cases the disk space runs out. When this happens, database jobs may stop and no longer run.

You should, therefore, periodically monitor the usage of disk space, and delete obsolete files located in the disk.

Monitor the disk usage of the disk where the following directories are located:

- Data storage destination directory
- Transaction log storage destination (if the transaction log is stored in a different directory from the data storage destination directory)
- Backup data storage destination directory
- Tablespace storage destination directory

## 7.2.1 Monitoring Disk Usage

To check the disk usage, use the following operating system commands:

- fsutil volume diskfree command

You can even use SQL statements to check tables and indexes individually.

Refer to "Determining Disk Usage" under "Server Administration" in the PostgreSQL Documentation for information on this method.



If you are using WebAdmin for operations, a warning is displayed when disk usage reaches 80%

## 7.2.2 Securing Free Disk Space

Secure free disk space by using the following operating system commands to delete unnecessary files, other than the database, from the same disk unit.

- del command

You can also secure disk space by performing the following tasks periodically:

- To secure space on the data storage destination disk:

Execute the REINDEX statement. Refer to "7.5 Reorganizing Indexes" for details.

- To secure space on the backup data storage destination disk:

Execute backup using WebAdmin or the pgx\_dmpall command.

## 7.3 Automatically Closing Connections

If an application stops responding and abnormally terminates for any reason, the connection from the application may remain active on the database server. If this situation continues for an extended period, other applications attempting to connect to the database server may encounter an error, or an error indicating that the tables are unavailable may occur.

It is, therefore, recommended that idle connections be closed automatically at regular intervals.

Set the following parameters in the postgresql.conf file to indicate the time permitted to elapse before a connection is closed.

Parameter name	Setting	Description
tcp_keepalives_idle	Time until keepalive is sent (seconds)  If 0, the default value of the system is used.	Sends keepalive to an idle connection at the specified interval in seconds  It is recommended to specify 30 seconds.
tcp_keepalives_interval	keepalive send interval (seconds)  If 0, the default value of the system is used.	Sends keepalive at the specified interval  It is recommended to specify 6 seconds.



The maximum number of connections allowed is 125, unless the desktop heap setting is changed.



Refer to "Connection Settings" under "Server Administration" in the PostgreSQL Documentation for information on the parameters.

.....

# 7.4 Monitoring the Connection State of an Application

FUJITSU Enterprise Postgres does not immediately delete the updated or deleted data. If the VACUUM determines there are no transactions that reference the database, FUJITSU Enterprise Postgres collects obsolete data.

However, obsolete data is not collected if there are connections that have remained active for an extended period or connections occupying resources. In this case the database may expand, causing performance degradation.



See

Refer to "Routine Vacuuming" under "Server Administration" in the PostgreSQL Documentation for information on the VACUUM command.

.......

In such cases, you can minimize performance degradation of the database by monitoring problematic connections.

The following methods are supported for monitoring connections that have been in the waiting status for an extended period:

- 7.4.1 Using the View (pg\_stat\_activity)
- 7.4.2 Using pgAdmin

### 7.4.1 Using the View (pg\_stat\_activity)

Use the view (pg\_stat\_activity) to identify and monitor connections where the client has been in the waiting status for an extended period.



The example below shows connections where the client has been in the waiting status for at least 60 minutes.

However, when considering continued compatibility of applications, do not reference system catalogs directly in the following SQL statements.

```
postgres=# select * from pg_stat_activity where state='idle in transaction' and current_timestamp >
cast(query_start + interval '60 minutes' as timestamp);
-[ RECORD 1 ]----+-----
             13003
datname
             | db01
pid
             4638
             | 10
usesysid
             | fsep
usename
application_name | apl01
           | 192.33.44.15
client_addr
client_hostname
             27500
client_port
backend_start | 2018-02-24 09:09:21.730641+09
            2018-02-24 09:09:23.858727+09
xact_start
            2018-02-24 09:09:23.858727+09
query_start
            2018-02-24 09:09:23.858834+09
state_change
wait_event_type | Client
wait_event
            ClientRead
state
             | idle in transaction
backend_xid
backend_xmin
query
             | begin;
            | client backend
backend_type
.....
```



- Refer to "Notes on Application Compatibility" in the Application Development Guide for information on maintaining application compatibility.

......

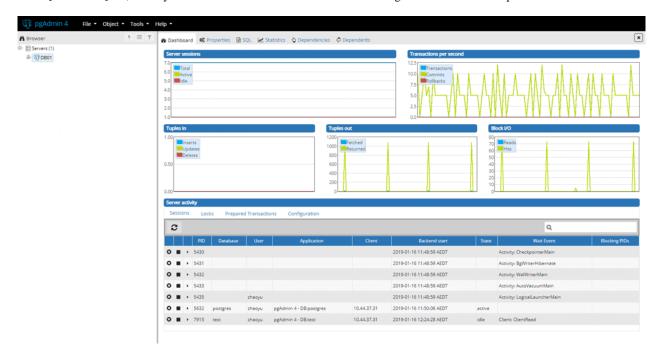
- Refer to "The Statistics Collector" under "Server Administration" in the PostgreSQL Documentation for information on pg\_stat\_activity.

## 7.4.2 Using pgAdmin

This section describes the procedure for monitoring connections using [Server Status] in pgAdmin.

1. In the [Browser] pane, click the database server for monitoring.

2. In the [Dashboard] tab, identify client connections that have been in the waiting state for an extended period.



## 7.5 Reorganizing Indexes

Normally, a database defines indexes in tables, but if data is frequently updated, indexes can no longer use free space in the disk efficiently. This situation can also cause a gradual decline in database access performance.

To rearrange used space on the disk and prevent the database access performance from declining, it is recommended that you periodically execute the REINDEX command to reorganize indexes.

Check the disk usage of the data storage destination using the method described in "7.2 Monitoring Disk Usage and Securing Free Space".



Because the REINDEX command retrieves the exclusive lock for an index being processed and locks writing of tables that are the source of the index, other processes that access these may stop while waiting to be locked.

Therefore, it is necessary to consider measures such as executing the command after the task is completed.



Refer to "Routine Reindexing" under "Server Administration" in the PostgreSQL Documentation for information on reorganizing indexes by periodically executing the REINDEX command.



Typically, reorganize indexes once a month at a suitable time such as when conducting database maintenance. Use SQL statements to check index usage. If this usage is increasing on a daily basis, adjust the frequency of recreating the index as compared to the free disk space.

The following example shows the SQL statements and the output.

However, when considering continued compatibility of applications, do not reference system catalogs and functions directly in the following SQL statements. Refer to "Notes on Application Compatibility" in the Application Development Guide for details.

#### [SQL statements]

```
SELECT
  nspname AS schema_name,
  relname AS index_name,
  round(100 * pg_relation_size(indexrelid) / pg_relation_size(indrelid)) / 100 AS index_ratio,
  pg_size_pretty(pg_relation_size(indexrelid)) AS index_size,
  pg_size_pretty(pg_relation_size(indrelid)) AS table_size
FROM pg_index I
  LEFT JOIN pg_class C ON (C.oid = I.indexrelid)
  LEFT JOIN pg_namespace N ON (N.oid = C.relnamespace)
WHERE
  C.relkind = 'i' AND
  pg_relation_size(indrelid) > 0
ORDER BY pg_relation_size(indexrelid) DESC, index_ratio DESC;
```

#### [Output]

schema_name	index_name	inc	dex_ratio	index_size	table_size
public pg_catalog pg_catalog	pgbench_accounts_pkey   pg_depend_depender_index   pg_depend_reference_index		0.6	2208 KB 224 KB 216 KB	13 MB   368 KB   368 KB



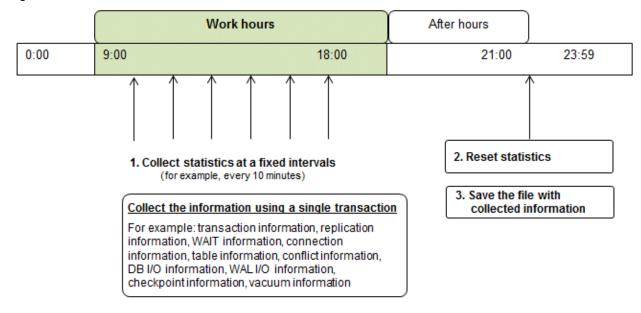
Refer to "Notes on Application Compatibility" in the Application Development Guide for information on maintaining application compatibility.

## 7.6 Monitoring Database Activity

FUJITSU Enterprise Postgres enables you to collect information related to database activity. By monitoring this information, you can check changes in the database status.

This information includes wait information for resources such as internal locks, and is useful for detecting performance bottlenecks. Furthermore, you should collect this information in case you need to request Fujitsu technical support for an investigation.

Figure 7.1 Overview of information collection



1. Collect statistics at fixed intervals during work hours.

Accumulate the collected information into a file.

Wherever possible, collect data from the various statistics views using a single transaction, because it enables you to take a snapshot of system performance at a given moment.

Refer to "7.6.1 Information that can be Collected" for information on the system views that can be collected.

2. Reset statistics after work hours, that is, after jobs have finished.

Refer to "7.6.3 Information Reset" for information on how to reset statistics.

3. Save the file with collected information.

Keep the file with collected information for at least two days, in order to check daily changes in performance and to ensure that the information is not deleted until you have sent a query to Fujitsu technical support.

Where jobs run 24 hours a day, reset statistics and save the file with collected information when the workload is low, for example, at night.



Statistics cumulatively add the daily database value, so if you do not reset them, the values will exceed the upper limit, and therefore will not provide accurate information.

The subsections below explain the following:

- Information that can be collected
- Collection configuration
- Information reset

#### 7.6.1 Information that can be Collected

Information that can be collected is categorized into the following types:

- Information common to PostgreSQL
- Information added by FUJITSU Enterprise Postgres

#### Information common to PostgreSQL



Refer to "Monitoring Database Activity" under "Server Administration" in the PostgreSQL Documentation for information common to PostgreSQL.

#### Information added by FUJITSU Enterprise Postgres

You can collect the following information added by FUJITSU Enterprise Postgres.

Table 7.1 Information added by FUJITSU Enterprise Postgres

View name	Description
pgx_stat_lwlock	Displays statistic related to lightweight lock, with each type of content displayed on a separate line. This information helps to detect bottlenecks.  Refer to "C.2 pgx_stat_lwlock" for details.
pgx_stat_latch	Displays statistics related latches, with each type of wait information within FUJITSU Enterprise Postgres displayed on a separate line. This information helps to detect bottlenecks.  Refer to "C.3 pgx_stat_latch" for details.

View name	Description
pgx_stat_walwriter	Displays statistics related to WAL writing, in a single line.
	Refer to "C.4 pgx_stat_walwriter" for details.
pgx_stat_sql	Displays statistics related to SQL statement executions, with each type of SQL statement displayed on a separate line.
	Refer to "C.5 pgx_stat_sql" for details.
pgx_stat_gmc	Displays statistics related to Global Meta Cache hit ration and used memory size.
	Refer to "C.6 pgx_stat_gmc" for detail. Also refer to Chapter 12 Global Meta Cache" for information on the Global Meta Cache.

# 7.6.2 Collection Configuration

The procedure for configuring collection depends on the information content.

- Information common to PostgreSQL
- Information added by FUJITSU Enterprise Postgres

#### Information common to PostgreSQL



Refer to "The Statistics Collector" in "Monitoring Database Activity" under "Server Administration" in the PostgreSQL Documentation for information on information common to PostgreSQL.

.....

#### Information added by FUJITSU Enterprise Postgres

Information added by FUJITSU Enterprise Postgres is collected by default.

To enable or disable information collection, change the configuration parameters in postgresql.conf. The following table lists the views for which you can enable or disable information collection, and the configuration parameters.

View name	Parameter
pgx_stat_lwlock	track_waits (*1)
pgx_stat_latch	
pgx_stat_sql	track_sql
pgx_stat_gmc	track_gmc

Remarks: You cannot change the collection status for  $pgx\_stat\_walwriter$ .

\*1: When executing the SQL statement with EXPLAIN ANALYZE, processing time may increase because of this information collection. It is recommended to set this parameter to "off" when executing EXPLAIN ANALYZE to check the processing time.

Refer to "Appendix A Parameters" for information on the parameters.

#### 7.6.3 Information Reset

This section describes how to reset information.

#### Information added by FUJITSU Enterprise Postgres

You can reset information added by FUJITSU Enterprise Postgres by using the pg\_stat\_reset\_shared function in the same way as for information common to PostgreSQL.

Configure the following parameters in the  $pg\_stat\_reset\_shared$  function:

Function	Type of return value	Description	
pg_stat_reset_shared(text)	void	Reset some cluster-wide statistics counters to zero, depending on the argument (requires superuser privileges).	
		Calling pg_stat_reset_shared('lwlock') will zero all counters shown in pgx_stat_lwlock.	
		Similarly, in the following cases, all values of the pertinent statistics counter are reset:	
		- If pg_stat_reset_shared('latch') is called:	
		All values displayed in pgx_stat_latch	
		- If pg_stat_reset_shared('walwriter') is called:	
		All values displayed in pgx_stat_walwriter	
		- If pg_stat_reset_shared('sql') is called:	
		All values displayed in pgx_stat_sql	
		- If pg_stat_reset_shared('gmc') is called:	
		All values except size column in pgx_stat_gmc	



Refer to "Statistics Functions" in "Monitoring Database Activity" under "Server Administration" in the PostgreSQL Documentation for information on other parameters of the  $pg\_stat\_reset\_shared$  function.

......

# Chapter 8 Streaming Replication Using WebAdmin

This chapter describes how to create a streaming replication cluster using WebAdmin.

Streaming replication allows the creation of one or more standby instances, which connect to the master instances and replicate the data using WAL records. The standby instance can be used for read-only operations.

WebAdmin can be used to create a streaming replication cluster. WebAdmin allows the creation of a cluster in the following configurations:

- Master-Standby Configuration: This configuration creates a master and standby instance together.
- Standby Only Configuration: This configuration creates a standby instance from an already existing instance.

## Point

- A standby instance can be created from a standalone instance, a master instance, or even from another standby instance.
- If a streaming replication cluster is created using WebAdmin, the network with the host name (or IP address) specified in [Host name] will be used across sessions of WebAdmin, and also used as the log transfer network.
- To use a network other than the job network as the log transfer network, specify the host name other than the job network one in [Host name].

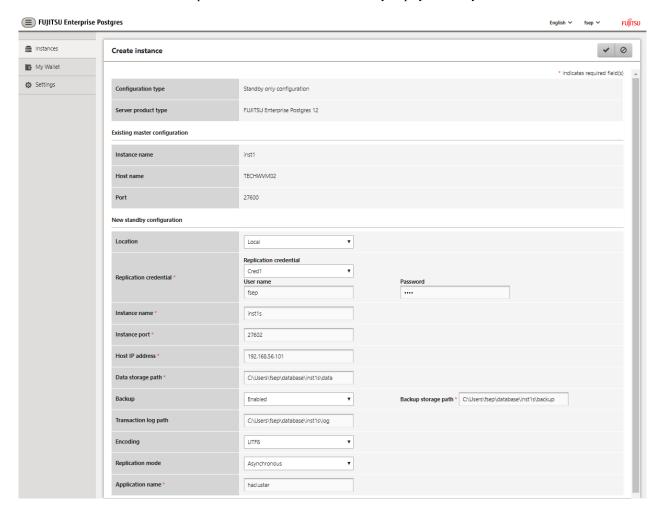
# 8.1 Creating a Standby Instance

Follow the procedure below to create a standby instance.

- 1. In the [Instances] tab, select the instance from which a standby instance is to be created.
- 2. Click

3. Enter the information for the standby instance to be created. In the example below, a standby instance is created from instance "inst1".

The instance name, host address and port of the selected instance are already displayed for easy reference.



#### Enter the following items:

- [Location]: Whether to create the instance in the server that the current user is logged in to, or in a remote server. The default is "Local", which will create the instance in the server machine where WebAdmin is currently running.
- [Replication credential]: The user name and password required for the standby instance to connect to the master instance. The user name and password can be entered or selected from the Wallet. Refer to "Appendix G WebAdmin Wallet" for information on creating wallet entries.
- [Instance name]: Name of the standby database instance to create.

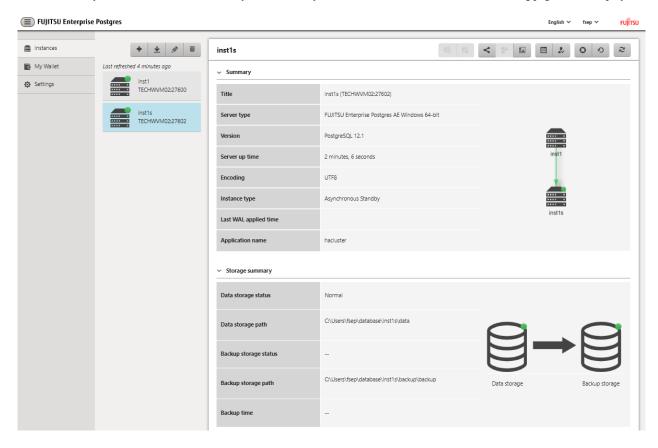
The name must meet the conditions below:

- Maximum of 16 characters
- The first character must be an ASCII alphabetic character
- The other characters must be ASCII alphanumeric characters
- [Instance port]: Port number of the standby database instance.
- [Host IP address]: The IP address of the server machine where the standby instance is to be created. This information is needed to configure the standby instance to be connected to the master.
- [Data storage path]: Directory where the database data will be stored
- [Backup storage path]: Directory where the database backup will be stored
- [Transaction log path]: Directory where the transaction log will be stored

- [Encoding]: Database encoding system
- [Replication mode]: Replication mode of the standby instance to be created ("Asynchronous" or "Synchronous")
- [Application name]: The reference name of the standby instance used to identify it to the master instance.

The name must meet the conditions below:

- Maximum of 16 characters
- The first character must be an ASCII alphabetic character
- The other characters must be ASCII alphanumeric characters
- 4. Click to create the standby instance.
- 5. Once the standby instance is created successfully, select standby instance in the [Instances] tab. The following page will be displayed:





- Backups are not possible for standby instances in WebAdmin. As a result, and are disabled and no value is shown for [Backup storage status] and [Backup time].
- If using WebAdmin to manage Mirroring Controller, the message below may be output to the server log or system log in the standby instance. No action is required, as the instance is running normally.

```
ERROR: pgx_rcvall failed (16491)
ERROR: pgx_rcvall: backup of the database has not yet been performed, or an incorrect backup storage directory was specified
```

- Replication credential (user name and password) should not contain hazardous characters. Refer to "Appendix H WebAdmin Disallow User Inputs Containing Hazardous Characters".

## 8.2 Promoting a Standby Instance

Streaming replication between a master and standby instance can be discontinued using WebAdmin.

Follow the procedure below to promote a standby instance to a standalone instance, thereby discontinuing the streaming replication.

- 1. In the [Instances] tab, select the standby instance that needs to be promoted.
- 2. Click Ta.
- 3. Click [Yes] from the confirmation dialog box.

The standby instance will be promoted and will become a standalone instance, which is not part of a streaming replication cluster.

Once the standby instance is promoted to become a standalone instance, the backup storage status will be "Error". This is because no backups are available when the instance is newly promoted to a standalone instance. The status will be reset if a new backup is performed by clicking [Solution] or ...

# 8.3 Converting an Asynchronous Replication to Synchronous

Streaming replication between a master and standby instance can be configured to be in Asynchronous or Synchronous mode. This mode can be changed even after the standby instance was successfully created.

Follow the procedure below to convert an Asynchronous standby instance to Synchronous.

- 1. In the [Instances] tab, select the master instance of the relevant cluster.
- 2. Click
- 3. In the [Streaming replication] section, edit the value for [Synchronous standby names].
  - Add the "Application name" of the standby instance you want to be in Synchronous mode.
- 4. Click .
- 5. Select the master instance and click
- 6. Select the standby instance. [Instance type] will now show the updated status.



- Converting an Asynchronous standby instance to Synchronous can cause the master instance to queue the incoming transactions until the standby instance is ready. For this reason, it is recommended that this operation be performed during a scheduled maintenance period.
- When adding a synchronous standby instance, FUJITSU Enterprise Postgres will only keep the first entry in [Synchronous standby names] in synchronous state.
- To learn more about the differences between synchronous and asynchronous standby modes and their behavior, refer to "Streaming Replication" in "High Availability, Load Balancing, and Replication" in the PostgreSQL Documentation.

# 8.4 Converting a Synchronous Replication to Asynchronous

Streaming replication between a master and standby instance can be configured to be in Asynchronous or Synchronous Mode. This mode can be changed even after the standby instance was successfully created.

Follow the procedure below to convert a Synchronous standby instance to Asynchronous.

- 1. In the [Instances] tab, select the master instance of the relevant cluster.
- 2. Click

- 3. In the [Streaming replication] section, edit the value for [Synchronous standby names].
  - Remove the "Application name" of the standby instance you want to be in Asynchronous mode.
- 4. Click .
- 5. Select the master instance and click .
- 6. Select the standby instance. [Instance type] will now show the updated status.



To learn more about the differences between synchronous and asynchronous standby modes and their behavior, refer to "Streaming Replication" in "High Availability, Load Balancing, and Replication" in the PostgreSQL Documentation.

# 8.5 Joining a Replication Cluster

WebAdmin facilitates the joining of an old master of the cluster as a standby node.

- 1. In the [Instances] tab, select the remote instance (from where the new cluster node will stream WAL entries), and then click 💂

- 2. Configure the node to accept streaming requests from the new node.
- 3. In the [Instances] tab, select the new standby instance (which needs to be connected to the cluster), and then click 🌠.



- 4. Set [Replication host name] to the remote instance.
- 5. Enter [Replication credential].

Specify the user name and password required for the standby instance to connect to the remote instance. The user name and password can be entered or selected from the Wallet. Refer to "Appendix G WebAdmin Wallet" for information on creating wallet entries. Replication credential (user name and password) should not contain hazardous characters. Refer to "Appendix H WebAdmin Disallow User Inputs Containing Hazardous Characters".

6. Enter [Host IP address].

Specify the IP address of the node where the standby instance was created.

- 7. Click v to open the [Join replication cluster] dialog box.
- 8. Select [Restart later] or [Restart now], and then click [Yes] to set up the standby instance.
- 9. Upon successful completion, the confirmation dialog box will be displayed.
- 10. Click [Close] to return to the instance details window.

The instance will become a standby instance, and will be part of the streaming replication cluster. The replication diagram will display the relationship between the standby instance and the remote instance. The user can change the replication relationship of the remote instance from asynchronous to synchronous (and vice versa) using the [Configuration] window.

# Chapter 9 Installing and Operating the In-memory Feature

The in-memory feature enables fast aggregation using Vertical Clustered Index (VCI) and memory-resident feature.

VCI has a data structure suitable for aggregation, and features parallel scan and disk compression, which enable faster aggregation through reduced disk I/O.

The memory-resident feature reduces disk I/O that occurs during aggregation. It consists of the preload feature that reads VCI data to memory in advance, and the stable buffer feature that suppresses VCI data eviction from memory. The stable buffer feature secures the proportion specified by parameter in the shared memory for VCI.

This chapter describes how to install and operate the in-memory feature.



This feature can only be used in Advanced Edition.

# 9.1 Installing Vertical Clustered Index (VCI)

This section describes the installation of VCI.

- 1. Evaluating whether to Install VCI
- 2. Estimating Resources
- 3. Setting up

### 9.1.1 Evaluating whether to Install VCI

VCI uses available resources within the server to increase scan performance.

It speeds up processing in many situations, and can be more effective in the following situations:

- Single table processing
- Processing that handles many rows in the table
- Processing that handles some columns in the table
- Processing that performs very heavy aggregation such as simultaneous sum and average aggregation

VCI will not be used in the following cases, so it is necessary to determine its effectiveness in advance:

- The data type of the target table or column contains VCI restrictions.
- The SQL statement does not meet the VCI operating conditions
- VCI is determined to be slower based on cost estimation



If performing operations that use VCI, the full\_page\_writes parameter setting in postgresql.conf must be enabled (on). For this reason, if this parameter is disabled (off), operations that use VCI return an error. In addition, to perform operations for tables that do not create a VCI when the full\_page\_writes parameter setting is temporarily disabled (off), do not create a VCI or perform operations to tables that created a VCI during that time.

......



- Refer to "9.1.4 Data that can Use VCI" for information on VCI restrictions.

- Refer to "Scan Using a Vertical Clustered Index (VCI)" - "Operating Conditions" in the Application Development Guide for information on VCI operating conditions.

### 9.1.2 Estimating Resources

Estimate resources before setting up VCI.

Select the aggregation that you want to speed up and identify the required column data. The additional resources below are required according to the number of columns.

- Memory

Secure additional capacity required for the disk space for the column for which VCI is to be created.

- Disk

Secure additional disks based on the disk space required for the column for which VCI is to be created, as VCI stores column data as well as existing table data on the disk. It is recommended to provide a separate disk in addition to the existing one, and specify it as the tablespace to avoid impact on any other jobs caused by I/O.



The operations on VCI can continue even if the memory configured for VCI is insufficient by using VCI data on the disk.



Refer to "Estimating Memory Requirements" and "Estimating Database Disk Space Requirements" in the Installation and Setup Guide for Server for information on how to estimate required memory and disk space.

### 9.1.3 Setting up

This section describes how to set up VCI.

#### Setup flow

- 1. Setting Parameters
- 2. Installing the Extensions
- 3. Creating VCI
- 4. Confirming that VCI has been Created

#### 9.1.3.1 Setting Parameters

Edit postgresql.conf to set the required parameters for VCI. After that, start or restart the instance.

The following table lists the parameters that need or are recommended to be configured in advance:

Parameter name	Setting value	Description	Required
shared_preload_libraries	Literal 'vci, pg_prewarm'	VCI and shared library to be preloaded at server start.	Y
session_preload_libraries	Literal 'vci, pg_prewarm'	VCI and shared library to be preloaded at connection start.	Y
reserve_buffer_ratio	Percentage of shared memory to be used for stable buffer table	Proportion of shared memory to be used for a stable buffer table.	N

Parameter name	Setting value	Description	Required
vci.control_max_workers	Number of background workers that manage VCI	Number of background workers that manage VCI.	N
		Add this value to max_worker_processes.	
vci.max_parallel_degree	Maximum number of background workers used for parallel scan	Maximum number of background workers used for parallel scan.  Add this value to max_worker_processes.	N
vci.smc_directory	Directory name in which a temporary file is created as the dynamic shared memory	Directory name in which a temporary file is created as the dynamic shared memory during a scan using a VCI.	N



```
shared_preload_libraries = 'vci, pg_prewarm'
                             'vci, pg_prewarm'
session_preload_libraries =
reserve_buffer_ratio = 20
vci.control_max_workers = 8
vci.max_parallel_degree = 4
max\_worker\_processes = 18 \# Example: If the initial value was 6, 6 + 8 + 4 = 18
vci.smc_directory = 'E:\\vci\\work'
```



If performing operations that use VCI, do not delete the shared library names specified to "shared\_preload\_libraries" and "session\_preload\_libraries". If operations that access VCI are performed after the VCI was defined, unexpected behavior may occur.



- Refer to "Appendix A Parameters" for information on all parameters for VCI. Refer also to default value for each parameter and details such as specification range in the same chapter. Refer to "Server Configuration" under "Server Administration" in the PostgreSQL documentation for information on shared\_preload\_libraries, session\_preload\_libraries, and max\_worker\_processes.

### 9.1.3.2 Installing the Extensions

Execute CREATE EXTENSION to install the VCI and pg\_prewarm extensions. Both extensions need to be installed for each database.

- Installing VCI

```
db01=# CREATE EXTENSION vci;
```

- Installing pg\_prewarm

```
db01=# CREATE EXTENSION pg_prewarm;
```



- Only superusers can install VCI extensions.
- VCI extensions can only be installed in public schema.
- Some operations cannot be performed for VCI extensions. Refer to "9.2.1 Commands that cannot be Used for VCI" for details.

#### 9.1.3.3 Creating a VCI

Execute the CREATE INDEX statement with the "USING vci" clause to create a VCI for the desired columns and the "WITH (stable buffer=true)" clause to enable the stable buffer feature.

To use a separate disk for the VCI, specify the TABLESPACE clause.

db01=# CREATE INDEX idx\_vci ON table01 USING vci (col01, col02) WITH (stable\_buffer=true);



- Some table types cannot be specified on the ON clause of CREATE INDEX. Refer to "9.1.4.1 Relation Types" for details.
- Some data types cannot be specified on the column specification of CREATE INDEX. Refer to "9.1.4.2 Data Types" for details.

- Some operations cannot be performed for VCI. Refer to "9.2.1 Commands that cannot be Used for VCI" for details.
- The same column cannot be specified more than once on the column specification of CREATE INDEX.
- VCI cannot be created for table columns that belong to the template database.
- CREATE INDEX creates multiple views named vci\_10digitRelOid\_5digitRelAttr\_1charRelType alongside VCI itself. These are called VCI internal relations. Do not update or delete them as they are used for VCI aggregation.
- All data for the specified column will be replaced in columnar format when VCI is created, so executing CREATE INDEX on an existing table with data inserted takes more time compared with a general index (B-tree). Jobs can continue while CREATE INDEX is running.
- When CREATE INDEX USING VCI is invoked on a partitioned table, the default behavior is to recurse to all partitions to ensure they all have matching indexes. Each partition is first checked to determine whether an equivalent index already exists, and if so, that index will become attached as a partition index to the index being created, which will become its parent index. If no matching index exists, a new index will be created and automatically attached; the name of the new index in each partition will be determined as if no index name had been specified in the command. If the ONLY option is specified, no recursion is done, and the index is marked invalid. (ALTER INDEX ... ATTACH PARTITION marks the index valid, once all partitions acquire matching indexes.) Note, however, that any partition that is created in the future using CREATE TABLE ... PARTITION OF will automatically have a matching index, regardless of whether ONLY is specified.
- Parallel index build is not supported on VCI indexes.

#### 9.1.3.4 Confirming that the VCI has been Created

Execute the SELECT statement to reference the pg\_indexes catalog, and confirm that the VCI was created for the target columns.



#### 9.1.4 Data that can Use VCI

This section describes on which relation types and for which data types VCIs can be created.

#### 9.1.4.1 Relation Types

VCIs cannot be created on some relation types.

The ON clause of CREATE INDEX described in "9.1.3.3 Creating a VCI" cannot specify relations on which VCIs cannot be created.

- Relations on which VCIs can be created
  - Normal tables
  - UNLOGGED TABLES
- Relations on which VCIs cannot be created
  - Materialized views
  - Temporary tables
  - Views
  - Temporary views
  - Foreign tables

# 9.1.4.2 Data Types

VCIs cannot be created for some data types.

The column specification of CREATE INDEX described in "9.1.3.3 Creating a VCI" cannot specify a column with data type on which VCIs cannot be created.

Category	Data type	Supported by VCI?
Numeric	smallint	Y
	integer	Y
	bigint	Y
	decimal	Y
	numeric	Y
	real	Y
	double precision	Y
	serial	Y
	bigserial	Y
Monetary	money	Y
Character	varchar(n)	Y
	char(n)	Y
	nchar	Y
	nvarchar(n)	Y
	text	Y
Binary	bytea	Y
Date/time	timestamp	Y
	timestamp with time zone	Y
	date	Y
	time	Y
	time with time zone	Y
	interval	Y
Boolean	boolean	Y
Geometric	point	N

Category	Data type	Supported by VCI?
	line	N
	lseg	N
	box	N
	path	N
	polygon	N
	circle	N
Network address	cidr	N
	inet	N
	macaddr	N
	macaddr8	N
Bit string	bit(n)	Y
	bit varying(n)	Y
Text search	tsvector	N
	tsquery	N
UUID	uuid	Y
XML	xml	N
JSON	json	N
	jbson	N
Range	int4range	N
	int8range	N
	numrange	N
	tsrange	N
	tstzrange	N
	daterange	N
Object identifier	oid	N
	regproc	N
	regprocedure	N
	regoper	N
	regoperator	N
	regclass	N
	regtype	N
	regconfig	N
	regdictionary	N
pg_lsn type	pg_lsn	N
Array type	-	N
User-defined type	-	N
(Basic type, enumerated type, composite type, and range type)		

# 9.2 Operating VCI

This section describes how to operate VCI.

### 9.2.1 Commands that cannot be Used for VCI

Some operations cannot be performed for VCI extensions and VCI itself.

This section describes SQL commands that cannot be executed for the VCI extensions and VCI itself, and client application commands.

#### **SQL** commands

- Operations that cannot be performed for the VCI extension

Command	Subcommand	Description
ALTER EXTENSION	UPDATE	The VCI extension cannot be specified.
	SET SCHEMA	This operation is not required for VCI.
	ADD	
	DROP	
CREATE EXTENSION	SCHEMA	The subcommands on the left cannot be performed if the VCI extension is specified.
		This operation is not required for VCI.

- Operations that cannot be performed on relations containing a VCI

Command	Subcommand	Description	
ALTER INDEX	SET	The subcommands on the left cannot be perform	
	SET TABLESPACE	if a VCI is specified.	
	ALL IN TABLESPACE	If the operation is required, delete the VCI using DROP INDEX, and re-create it using CREATE INDEX after completing the operation.	
ALTER OPERATOR	RENAME TO	The subcommands on the left cannot be performed	
CLASS	OWNER TO	if a VCI is specified.	
	SET SCHEMA	This operation is not supported in VCI.	
ALTER OPERATOR FAMILY	ADD		
	DROP		
	RENAME TO		
	OWNER TO		
	SET SCHEMA		
ALTER TABLE	ALL IN TABLESPACE name [OWNED BY roleName] SET TABLESPACE newTablespace	A tablespace that contains a VCI cannot be specified.  If the operation is required, delete the VCI using DROP INDEX, and re-create it using CREATE INDEX after completing the operation.	
	DROP [ COLUMN ] [ IF EXISTS ] colName [ RESTRICT   CASCADE ]	A column that contains a VCI cannot be specified.  If the operation is required, delete the VCI using	
	ALTER [ COLUMN ] colName [ SET DATA ] TYPE dataType [ COLLATE collation ] [ USING expr ]	DROP INDEX, and re-create it using CREATE INDEX after completing the operation.	

Command	Subcommand	Description
	CLUSTER ON indexName	A VCI cannot be specified.
	REPLICA IDENTITY {DEFAULT   USING INDEX indexName   FULL   NOTHING}	This operation is not supported in VCI.
CLUSTER	-	A table that contains a VCI and VCI cannot be specified.
		If the operation is required, delete the VCI using DROP INDEX, and re-create it using CREATE INDEX after completing the operation.
CREATE INDEX	UNIQUE	The subcommands on the left cannot be performed
	CONCURRENTLY	if a VCI is specified.
	[ ASC   DESC ]	This operation is not supported in VCI.
	[ NULLS { FIRST   LAST } ]	
	WITH	
	WHERE	
	INCLUDE	
CREATE OPERATOR CLASS	-	A VCI cannot be specified.  This operation is not supported in VCI.
CREATE OPERATOR FAMILY	-	This operation is not supported in ver.
CREATE TABLE	EXCLUDE	
DROP INDEX	CONCURRENTLY	The subcommands on the left cannot be performed if a VCI is specified.
		This operation is not supported in VCI.
REINDEX	-	A VCI cannot be specified.
		This command is not required as VCI uses daemon's automatic maintenance to prevent disk space from increasing.

#### Client application command

- Operations that cannot be performed on relations containing a VCI

Command	Description	
clusterdb	Clustering cannot be performed for tables that contain a VCI.	
reindexdb	VCIs cannot be specified on theindex option.	

### 9.2.2 Data Preload Feature

The first aggregation using VCI immediately after an instance is started may take time, because the VCI data has not been loaded to buffer. Therefore, use the preload feature to load the VCI data to buffer in advance when performing VCI aggregation after an instance is started. When using the preload feature, execute the function pgx\_prewarm\_vci to each VCI created with CREATE INDEX.

```
db01=# SELECT pgx_prewarm_vci('idx_vci');
```



Refer to "B.4 VCI Data Load Control Function" for information on pgx\_prewarm\_vci.

# Chapter 10 Parallel Query

FUJITSU Enterprise Postgres enhances parallel queries, by taking into consideration the aspects below:

- CPU load calculation
- Increase of workers during runtime

### 10.1 CPU Load Calculation

There may be a case when the user tries to execute a parallel query but there is not enough CPU available.

Adding dynamic workers at this stage will provide no benefits - instead, it may add overhead due to context switching.

FUJITSU Enterprise Postgres takes into consideration the current CPU load when deciding on the number of workers for parallel query.

# 10.2 Increase of Workers during Runtime

This FUJITSU Enterprise Postgres enhancement allows systems to allocate additional workers during query execution (if there are workers available at the time). This improves query performance, which could otherwise starve of CPU if there were fewer or no workers when the query started.



The ability to increase workers during runtime is available only with parallel sequence scan.

# Chapter 11 High-Speed Data Load

High-speed data load uses the pgx\_loader command to load data from files at high speed into FUJITSU Enterprise Postgres.



- This feature is available only in the Advanced Edition.
- This feature is not available in single-user mode. This is because in single-user mode instances run in a single process, and it cannot start parallel workers.

# 11.1 Installing High-Speed Data Load

This section describes how to install high-speed data load.

#### Installation flow

- 1. Deciding whether to Install
- 2. Estimating Resources
- 3. Setup

#### 11.1.1 Deciding whether to Install

The feature achieves high speed data load by executing the COPY FROM command in parallel. If the database system is unable to use sufficient resources due to the feature using more resources than the COPY FROM command of PostgreSQL, load performance may be inferior to that of the COPY FROM command of PostgreSQL. Therefore, determine if the feature will be effective by considering the factors below before deciding to install.

#### Database server memory

If the value of shared\_buffers in postgresql.conf is small, fewer data pages are cached to the shared memory of the database server. This will result in multiple parallel workers more often having to wait for write exclusive locks to the same data page. Moreover, the smaller the number of data pages, the more often the table expands. During table expansion, access to the table is exclusive (standby event name: extend), so write time increases. To cater for that, increase the value of shared\_buffers.



The standby event name is stored in the wait\_event column of the pg\_stat\_activity view. Refer to "wait\_event Description" in "The Statistics Collector" in the PostgreSQL Documentation for details.

#### Frequency of checkpoints

If checkpoints are issued at short intervals, write performance is reduced. If the messages below are output to the server log during data writes, increase the values of max\_wal\_size and checkpoint\_timeout in postgresql.conf to reduce the frequency of checkpoints.



LOG: checkpoints are occurring too frequently (19 seconds apart)
HINT: Consider increasing the configuration parameter "max\_wal\_size".

## 11.1.2 Estimating Resources

Estimate the memory requirements for high-speed data load.

Up to 128 parallel workers to perform data load can be specified for this feature. The additional resources below are required depending on the number of parallel workers.

- Dynamic shared memory created during data load

The feature creates shared memory and shared memory message queues during data load. These are used to send external data from the back end to the parallel workers, and for error notifications.



If the value of shared\_buffers in postgresql.conf is small, the system will often have to wait for write exclusive locks to the same data page (as described in "Database server memory" in "11.1.1 Deciding whether to Install"). Since input data cannot be loaded from the shared memory message queues during such waits, they will often be full. In these cases, it will not be possible to write to the shared memory message queues, resulting in degraded data load performance.



Refer to "High-Speed Data Load Memory Requirements" in the Installation and Setup Guide for Server for information on the formula for estimating memory requirements.

......

### 11.1.3 Setup

This section describes how to set up high-speed data load.

#### Setup flow

- 1. Setting Parameters
- 2. Installing the Extension

#### 11.1.3.1 Setting Parameters

Set the parameters required for high-speed data load in postgresql.conf. After that, start or restart the instance.

The table below lists the postgresql.conf parameters that must be changed, and the values that must be added to their current values. After editing postgresql.conf, start or restart the instance.

Setting	Required
Add the number of transactions that can be prepared by parallel workers during data load to the parameter's current value.	Mandatory
The resulting value must be equal to or greater than the maximum number of parallel workers used with this feature.	
Number of parallel workers to perform data load.	Mandatory
Add the maximum number of parallel workers to be used in a parallel query by this feature to the parameter's current value.  The resulting value must be equal to or greater than the number of parallel workers used with this feature.	Mandatory
	workers during data load to the parameter's current value.  The resulting value must be equal to or greater than the maximum number of parallel workers used with this feature.  Number of parallel workers to perform data load.  Add the maximum number of parallel workers to be used in a parallel query by this feature to the parameter's current value.



The example below shows how to configure 2 instances of high-speed data load being executed simultaneously using a degree of parallelism of 4.

```
max_prepared_transactions = 13 #Example if the initial value was 5: 5 + 2 x 4 = 13
max_worker_processes = 16 #Example if the initial value was 8: 8 + 2 x 4 = 16
max_parallel_workers = 12 #Example if the initial value was 4: 4 + 2 x 4 = 12
```



As shown in the example above, set the value of max\_prepared\_transactions, max\_worker\_processes and max\_parallel\_workers multiplied by the number of instances of this feature executed simultaneously.

The table below lists the postgresql.conf parameters that must also be checked.

Parameter	Setting	Required
dynamic_shared_memory_type	Implementation of dynamic shared memory to be used by the instance.	Mandatory
	The default value is recommended.	



Refer to "Resource Consumption" in the PostgreSQL Documentation for information on the parameters.

#### 11.1.3.2 Installing the Extension

Execute CREATE EXTENSION to install the high-speed data load extension. The extension needs to be installed on each database.



The example below installs the extension on the "postgres" database.

postgres=# CREATE EXTENSION pgx\_loader; CREATE EXTENSION



- Only superusers can install the high-speed data load extension.
- The high-speed data load extension can only be installed on the public schema.

## 11.2 Using High-Speed Data Load

This section describes how to use high-speed data load.

## 11.2.1 Loading Data

To load data from a file into a FUJITSU Enterprise Postgres table, execute the pgx\_loader command in load mode.



The example below loads the file E:\path\to\data.csv (2000 records) into table tbl using a degree of parallelism of 3.

```
> pgx_loader load -j 3 -c "COPY tbl FROM 'E:\path\to\data.csv' WITH CSV"
LOAD 2000
```



If an external file contains data that violates the format or constraints, the data load may fail partway through, resulting in delays for routine tasks such as nightly batch processing. Therefore, it is recommended to remove the invalid data before executing the data load.



The data inserted using this feature is dumped as a COPY command by the pg\_dump command and the pg\_dumpall command.



- Refer to "pgx\_loader" in the Reference for information on the command.
- Refer to "COPY" in the PostgreSQL Documentation for information on the deployment destination and access privileges for external files.

#### 11.2.2 Recovering from a Data Load that Ended Abnormally

If a system interruption such as a server failure occurs while high-speed data load is being performed, transactions prepared using this feature may be changed to the in-doubt state. At that point, resources occupied by the transaction will be locked, and access to the relevant resources from other transactions will be blocked, rendering them unusable.

In such cases, check transactions that are in an in-doubt state, and resolve them.

#### Checking for in-doubt transactions

This section describes how to check for in-doubt transactions.

1. Refer to the pgx\_loader\_state table in the pgx\_loader schema.

Retrieve the global transaction identifier (gid column) of in-doubt transactions. In-doubt transactions will contain "rollback" in the column "state".



The example below retrieves the global transaction identifier (gid) of in-doubt transactions performed by the database role myrole and that used table tbl. The retrieved global transaction identifiers pgx\_loader:9589 and pgx\_loader:9590 identify in-doubt transactions.

2. Refer to the pg\_prepared\_xacts system view.

Check if the in-doubt transactions retrieved above exist.



The example below checks if in-doubt transactions with the global transaction identifiers pgx\_loader:9589 and pgx\_loader:9590 exist.



Refer to "E.1 pgx\_loader\_state" for information on the pgx\_loader\_state table.

#### Resolving in-doubt transactions

Execute the pgx\_loader command in recovery mode to resolve in-doubt transactions.

After executing the pgx\_loader command in recovery mode, perform the procedure described in "Checking for in-doubt transactions" to check if the in-doubt transactions have been resolved.



The example below completes the in-doubt transactions prepared for table tbl.

```
$ pgx_loader recovery -t tbl
```



The recovery mode of the pgx\_loader command only resolves transactions prepared by high-speed data load. For transactions prepared by an application using distributed transactions other than this feature, follow the procedure described in "14.13 Actions in Response to Error in a Distributed Transaction".

## 11.3 Removing High-Speed Data Load

This section describes how to remove high-speed data load.

## 11.3.1 Removing the Extension

Execute DROP EXTENSION to remove the high-speed data load extension. The extension needs to be removed on each database.



The example below removes the extension on the "postgres" database.

postgres=# DROP EXTENSION pgx\_loader;
DROP EXTENSION



- The information required for operation of high-speed data load is stored in the pgx\_loader\_state table of the pgx\_loader schema. Do not remove the high-speed data load extension if the pgx\_loader\_state table is not empty.
- Only superusers can remove the high-speed data load extension.
- The high-speed data load extension can only be removed on the public schema.

# Chapter 12 Global Meta Cache

The Global Meta Cache (GMC) feature loads a meta cache into shared memory using the pgx\_global\_metacache parameter. This reduces the amount of memory required throughout the system.



This feature can only be used in Advanced Edition.

# **12.1 Usage**

Describes how to use the Global Meta Cache feature.

### 12.1.1 Deciding Whether to Enable the Global Meta Cache Feature

Global Meta Cache is a mechanism for sharing meta caches between processes, so it works well on systems with a high number of resources accessed and SQL connections. The number of resources is primarily the number of tables accessed by a process, the number of indexes, or the total number of all columns in all tables accessed.

In particular, consider using Global Meta Cache if the total size of the meta cache for each process exceeds the amount of installed memory, or takes up a large portion of that memory, thereby squeezing memory allocations to the database cache or the Operating system file cache. Using Global Meta Cache may slightly increase the time it takes to execute a single SQL to reference a meta cache on shared memory (For example, it will not exceed 5% of the execution time in a situation where all the data is located in the database cache), but you can expect a greater benefit from being able to allocate more memory, such as for the database cache.

If performance degradation using Global Meta Cache is not acceptable, you may want to limit the number of tables accessed by a process.

### 12.1.2 Estimating Memory for Global Meta Cache

To enable the Global Meta Cache feature, the pgx\_global\_metacache parameter must specify an upper limit on the size of the shared memory (Hereinafter, the GMC area) dedicated to Global Meta Cahche. Ideally, this upper limit should be the size estimated in "Appendix A Parameters". Values lower than this can still work, but refer to "12.1.3 How the GMC Memory Area Is Used" on using the GMC area to understand the disadvantages.

# 12.1.3 How the GMC Memory Area Is Used

At startup, the memory for the GMC area is not used much, but the GMC area grows as new meta caches are placed in the GMC area. If it does, it discards any meta caches that the system determines are not heavily used and places a new one in the GMC area.

Therefore, the GMC area will work even if it is smaller than the estimate, but the meta cache will be regenerated if the discarded meta cache needs to be reused. Note that if this happens frequently, it will degrade overall performance.

With this in mind, it may not be a problem if, for example, the tables to be accessed are different depending on the time zone, and the degradation of the time zone immediately after the change is acceptable.

In any case, be sure to test and tune the system thoroughly before running it.

# 12.1.4 Enabling the Global Meta Cache Feature

To Enable the Global Meta Cache feature edit the postgresql.conf file and set the pgx\_global\_metacahe parameter. Restarting the instance after editing the postgresql.conf file is required. Refer to "Appendix A Parameters" for information on the parameters.

Parameter Name	Description
pgx_global_metacache	Specify the maximum amount of memory for the GMC area on shared memory. When it's set to 0 (default value), the Global Meta Cache feature is disabled.  When enabled, the minimum value allowed is 10MB.

When the cache is created, if the total amount of meta caches on shared memory exceeds the value specified by pgx\_global\_metacache, the inactive, unreferenced meta caches are removed from the GMC area. Note that if all GMC are in use and the cache cannot be created in the GMC area, the cache is temporarily created in the local memory of the backend process.



Here is an example postgresql.conf configuration:

pgx\_global\_metacache = 800 MB

#### Wait Events

The Global Meta Cache feature may cause wait events. Wait events are identified in the wait\_event column of the pg\_stat\_activity view. GMC specific wait events are described below.

[GMC Feature Wait Events]

Wait Event Type	Wait Event Name	Description
LWLock	global_metacache	Waiting to find, add, and remove meta caches in the GMC area.
	global_metacache_sweep	Waiting to select a meta cache that can be deleted when GMC space is low.
		If the GMC is fully referencing and there is no deletable meta cache, it is waiting for the reference to be removed and a deletable meta cache to be selected.



If global\_metacache\_sweep is happened frequently, increase the pgx\_global\_metacache setting.



See

Refer to "Viewing Statistics" in the PostgreSQL Documentation for information on the pg\_stat\_activity view.

### 12.1.5 Estimating Resources

Refer to "Global Meta Cache Memory Requirements" in the Installation and Setup Guide for Server for formulas to estimate the amount of memory used by the Global Meta Cache feature.

#### 12.2 Statistics

Describes the statistics for the Global Meta Cache feature.

## 12.2.1 System View

You can check the cache hit ratio and size of the GMC area in the system view pgx\_stat\_gmc. Refer to "C.6 pgx\_stat\_gmc" for information on the columns.

If the cache hit ratio is low and the current memory usage is close to pgx\_global\_metacache, increase the pgx\_global\_metacache setting because performance may be degraded.

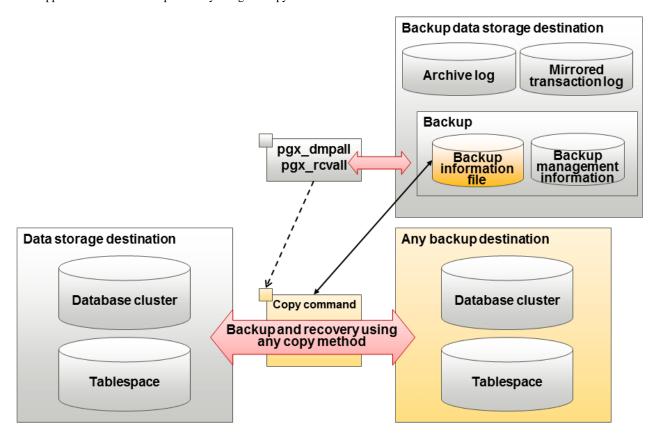
Refer to "7.6 Monitoring Database Activity" in the Operations Guide for information on the statistics.

# Chapter 13 Backup/Recovery Using the Copy Command

By using a copy command created by the user, the pgx\_dmpall command and the pgx\_rcvall command can perform backup to any destination and can perform recovery from any destination using any copy method.

Copy commands must be created in advance as executable scripts for the user to implement the copy process on database clusters and tablespaces, and are called when executing the pgx\_dmpall and pgx\_revall commands.

This appendix describes backup/recovery using the copy command.





It is also possible to back up only some tablespaces using the copy command. However, database resources not backed up using the copy command are still backed up to the backup data storage destination.



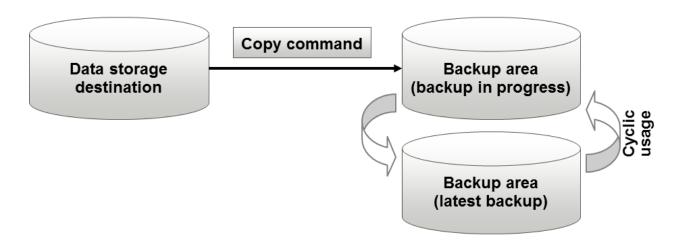
Both the backup data storage destination and the optional backup destination are necessary for recovery - if they are located in secondary media, combined management of these is necessary.

# 13.1 Configuration of the Copy Command

This section describes the configuration of the copy command for backup and recovery.

#### Cyclic usage of the backup area

Prepare two backup areas for the copy command in case an issue affects the data storage destination during backup. The copy command performs backup while cyclically using these backup areas.

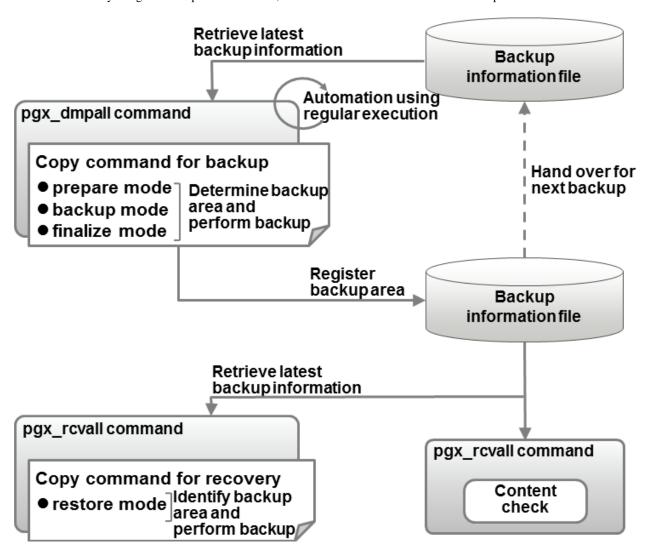




The backup data storage destination cannot be used as these backup areas used by the copy command.

# Backup using the backup information file

The copy command must determine the backup destination on each backup, as it is necessary to cycle through the backup areas. Backup can be automated by using the backup information file, which contains information about the backup destination.





The backup information file is prepared in the backup data storage destination by the pgx\_dmpall command, and contains information that can be read or updated by the copy command. This file is managed by associating it with the latest backup successfully completed by the pgx\_dmpall command, so the latest backup information relating to the copy command registered by the user can be retrieved. Additionally, the content of the backup information file can be displayed using the pgx\_revall command.

# Configuration of the copy command for backup

The pgx\_dmpall command calls the copy command for backup after execution for the three modes below. It is therefore necessary for the copy command for backup to implement the required processing for each of the modes.

- prepare mode

Determines which of the two backup areas will be used for the current backup.

The backup area to be used for the current backup is determined by reading the information relating to the latest backup destination where the backup information file was written to during the previous backup.

- backup mode

Performs backup on the backup area determined by prepare mode, using any copy method.

- finalize mode

Writes information relating to the destination of the current backup to the backup information file.

This enables the prepare mode to check the destination of the previous backup during the next backup.



The user can use any method to hand over backup information between modes within the copy command, such as creating temporary files.

## Configuration of the copy command for recovery

The pgx\_rcvall command calls the copy command for recovery for the mode below. It is therefore necessary for the copy command for recovery to implement the required processing for the mode.

- restore mode

Any copy method can be used to implement restore from the backup destination retrieved using the copy command for backup.



By referring to the mode assigned to the copy command as an argument, backup and recovery can be implemented using a single copy command.



# Using a batch file

```
if "%1" == "prepare" goto prepare
if "%1" == "backup" goto backup
if "%1" == "finalize" goto finalize
if "%1" == "restore" goto restore
:prepare
    processingRequiredForPrepareMode
    exit /b
:backup
    processingRequiredForBackupMode
```

```
exit /b
:finalize
   processingRequiredForFinalizeMode
   exit /b
:restore
   processingRequiredForRestoreMode
   exit /b
```



A sample batch file that backs up the database cluster and tablespace directory to a specific directory is supplied to demonstrate how to write a copy command.

The sample is stored in the directory below:

installDir\share\copy\_command.bat.sample

# 13.2 Backup Using the Copy Command

To perform backup using the copy command, in addition to performing the standard backup procedure, it is also necessary to create a copy command, and then execute the pgx\_dmpall command specifying it. This section describes the procedure specific to using the copy command.

## Preparing for backup

You must prepare for backup before actually starting the backup process.

Perform the following procedure:

1. Determine the database resources to be backed up

Determine the database resources to be backed up using the copy command. The copy command can back up the following resources:

- Database cluster
- Tablespace

To back up only some tablespaces, create a file listing them. This file is not necessary to back up all tablespaces.

## Example

To back up only tablespaces tblspc1 and tblspc2

tblspc1 tblspc2

2. Prepare a backup area

Prepare a backup area to save the database resources to be backed up, as determined in step 1.

3. Create the copy command

Create the copy commands for backup and recovery. Refer to "13.4 Copy Command Interface" for details.

# Performing backup

Execute the pgx\_dmpall command with the -Y option specifying the full path of the copy command for backup created in step 3 of preparation for backup.

The example below backs up only some tablespaces, but not the database cluster, using the copy command.



> pgx\_dmpall -D D:\database\inst1 -Y D:\database\command\backup.bat
--exclude-copy-cluster -P D:\database\command\tablespace\_list.txt



- To exclude up the database cluster from backup using the copy command, specify the --exclude-copy-cluster option.
- To back up only some tablespaces using the copy command, use the -P option specifying the full path of the file created in step 1 of preparation for backup.



- Refer to "pgx\_dmpall" in the Reference for information on the command.

# Checking backup status

Use the pgx\_rcvall command to check the backup status.

Execute the pgx\_rcvall command with the -l option specified to output backup data information. If backup was performed using the copy command, the resources backed up using the copy command will also be output.



# 13.3 Recovery Using the Copy Command

To perform recovery using the copy command, in addition to performing the standard recovery procedure, it is also necessary to create a copy command, and then execute the pgx\_rcvall command specifying it. This section describes the procedure specific to using the copy command.

# Determining the backup area of the latest backup

Check the backup information file to determine the backup area used for the latest backup, and confirm that it is in a recoverable state.

Execute the pgx\_rcvall command with the --view-results-of-copying option to output the content of the backup information file.



> pgx\_rcvall -D D:\database\inst1 --view-results-of-copying

# **Perform recovery**

Execute the pgx\_rcvall command with the -Y option specifying the full path of the copy command for recovery created in step 3 of the preparation for backup described in "13.2 Backup Using the Copy Command".

The example below recover only some tablespaces, but not the database cluster, using the copy command.



> pgx\_rcvall -D D:\database\inst1 -B E:\backup\inst1 -Y D:\database\command\recovery.bat



If the latest backup was performed using the copy command, the pgx\_rcvall command automatically recognizes which database resources were backed up using the copy command, or whether resources were backed up to the backup data storage destination. Therefore, recovery can be performed by simply executing the pgx\_rcvall command specifying the copy command for recovery.



Refer to "pgx\_rcvall" in the Reference for information on the command.

# 13.4 Copy Command Interface

The following types of copy command are available:

- Copy command for backup
- Copy command for recovery

This appendix describes the interface of each copy command.

# 13.4.1 Copy Command for Backup

# **Feature**

User exit (for the copy command) called from the pgx\_dmpall command.

# **Format**

The syntax for calling the copy command from the pgx\_dmpall command is described below.

If the operation mode is "prepare"

```
copyCommandName prepare "pathOfBackupInfoFile" "pathOfBackupTargetListFile"
```

If the operation mode is "backup"

copyCommandName backup

If the operation mode is "finalize"

copyCommandName finalize "pathOfBackupInfoFile"

# **Argument**

- Operation mode

Mode	Description		
prepare	Implements the preparation process for backing up using the copy command.		
	Called before the PostgreSQL online backup mode is started.		
backup	Implements the backup process.		
	Called during the PostgreSQL online backup mode.		

Mode	Description	
finalize	Implements the backup completion process.	
	Called after the PostgreSQL online backup mode is completed.	

- Full path of the backup information file

Full path of the backup information file of the latest backup, enclosed in double quotation marks. If a backup has not been performed, specify '-'.

- Full path of the backup target list file

Full path of the file containing the resources to be backed up using the copy command, enclosed in double quotation marks. One of the following is described in each resource name.

Resource	Description	
Database cluster	pg_data	
Tablespace	Tablespace name	



To back up the database cluster and the tablespaces dbspace and indexspace using the copy command, the file should contain the following:

pg\_data
dbspace
indexspace

# **Information**

The encoding of resource names output to the backup target list file by the pgx\_dmpall command is the encoding used when this command connects to the database with auto specified for the client\_encoding parameter, and is dependent on the locale at the time of command execution.

The number of arguments vary depending on operation mode. The argument of each operation mode is as follows.

Operation mode	First argument	Second argument	Third argument	
prepare	Operation mode	Backup information file path name	Backup target list file path name	
backup		None	None	
finalize		Backup information file path name		

Additionally, the access permissions for the backup information file and backup target list file are different depending on the operation mode. The access permissions of each operation mode are as follows:

Operation mode	Backup information file	Backup target list file
prepare	Can be viewed	Can be viewed
backup	-	-
finalize	Can be viewed and updated	-

## Return value

Return value	Description
0	Normal end

Return value	Description	
	The pgx_dmpall command continues processing.	
Other than 0	Abnormal end	
	The pgx_dmpall command terminates in error.	

# Description

- The copy command operates with the privileges of the operating system user who executed the pgx\_dmpall command. Therefore, grant copy command execution privileges to users who will execute the pgx\_dmpall command. Additionally, have the copy command change users as necessary.
- To write to the backup information file, use a method such as redirection from the copy command.
- Because the copy command is called for each mode, implement all processing for each one.
- To copy multiple resources simultaneously, have the copy command copy them in parallel.



- The backup information file and backup target list file cannot be deleted. Additionally, the privileges cannot be changed.
- Standard output and standard error output of the copy command are output to the terminal where the pgx\_dmpall command was executed.
- If the copy command becomes unresponsive, the pgx\_dmpall command will also become unresponsive. If the copy command is deemed to be unresponsive by the operating system, use an operating system command to forcibly stop it.
- Output the copy command execution trace and the result to a temporary file, so that if it terminates in error, the cause can be investigated at a later time.
- For prepare mode only, it is possible to use the PostgreSQL client application to access the database using the copy command. For all other modes, do not execute FUJITSU Enterprise Postgres commands or PostgreSQL applications.
- Enable the fsync parameter in postgresql.conf, because data on the shared memory buffer needs to have been already written to disk when backup starts.

# 13.4.2 Copy Command for Recovery

### **Feature**

User exit (for the copy command) called from the pgx\_rcvall command.

### **Format**

The syntax for calling the copy command from the pgx\_rcvall command is described below.

copyCommandName restore "pathOfBackupInfoFile" "pathOfBackupTargetListFile"

# Argument

- Operation mode

Mode	Description
restore	Performs restore.

- Full path of the backup information file

Full path of the backup information file, enclosed in double quotation marks.

- Full path of the backup target list file

Full path of the file containing the resources to be restored using the copy command, enclosed in double quotation marks.

The access permissions for the backup information file and backup target list file are as below.

Backup information file	Backup target list file
Can be viewed	Can be viewed

### Return value

Return value	Description
0	Normal end
	The pgx_rcvall command continues processing.
Other than 0	Abnormal end
	The pgx_rcvall command terminates in error.

# **Description**

- The copy command operates with the privileges of the operating system user who executed the pgx\_rcvall command. Therefore, grant copy command execution privileges to users who will execute the pgx\_rcvall command. Additionally, have the copy command change users as necessary.
- The copy command is called once only in restore mode.
- To copy multiple resources simultaneously, have the copy command copy them in parallel.



- The backup information file and backup target list file cannot be deleted. Additionally, the privileges cannot be changed.
- Standard output and standard error output of the copy command are output to the terminal where the pgx\_rcvall command was executed.

......

- If the copy command becomes unresponsive, the pgx\_rcvall command will also become unresponsive. If the status of the copy command is deemed to be unresponsive by the operating system, use an operating system command to forcibly stop it.
- Output the copy command execution trace and the result to a temporary file, so that if it terminates in error, the cause can be investigated at a later time.
- Do not execute FUJITSU Enterprise Postgres commands or PostgreSQL applications in the copy command.
- There may be files and directories not required for recovery using the archive log included in the backup, such as postmaster.pid, pg\_wal/subdirectory and pg\_replslot in the database cluster. If such unnecessary files and directories exist, have the copy command delete them after the restore.

# Chapter 14 Actions when an Error Occurs

This chapter describes the actions to take when an error occurs in the database or an application, while FUJITSU Enterprise Postgres is operating.

Depending on the type of error, it may be necessary to recover the database cluster. The recovery process recovers the following resources:

- Data storage destination
- Transaction log storage destination (if the transaction log is stored in a separate disk from the data storage destination)
- Backup data storage destination



Even if a disk is not defective, the same input-output error messages, as those generated when the disk is defective, may be output. The recovery actions differ for these error messages.

.....

Check the status of the disk, and select one of the following actions:

- If the disk is defective

Refer to "14.1 Recovering from Disk Failure (Hardware)", and take actions accordingly.

- If the disk is not defective

Refer to "14.14 I/O Errors Other than Disk Failure", and take actions accordingly.

A few examples of errors generated even if the disk is not defective include:

- Network error with an external disk
- Errors caused by power failure or mounting issues

## Determining the cause of an error

If an error occurs, refer to the WebAdmin message and the event log, and determine the cause of the error.



Refer to "Configuring Parameters" in the Installation and Setup Guide for Server for information on server logs.

# Approximate recovery time

The formulas for deriving the approximate recovery time of resources in each directory are given below.

If using the copy command with the pgx\_rcvall command, the recovery time will depend on the implementation of the copy command.

- Data storage destination or transaction log storage destination

 $\label{eq:recovery_time} \textit{Recovery time} = (\textit{usageByTheDataStorageDestination} + \textit{usageByTheTransactionLogStorageDestination}) \ / \ diskWritePerformance x 1.5$ 

- usageByTheDataStorageDestination: Disk space used by the database cluster
- usageByTheTransactionLogStorageDestination. Disk space used by the transaction log stored outside the database cluster
- *diskWritePerformance*: Measured maximum data volume (bytes/second) that can be written per second in the system environment where the operation is performed
- 1.5: Coefficient assuming the time excluding disk write, which is the most time-consuming step
- Backup data storage destination

Recovery time = usageByTheBackupDataStorageDestination / diskWritePerformance x 1.5

- usageByTheBackupDataStorageDestination: Disk space used by the backup data
- *diskWritePerformance*: Measured maximum data volume (bytes/second) that can be written per second in the system environment where the operation is performed
- 1.5: Coefficient assuming the time excluding disk write, which is the most time-consuming step

# 14.1 Recovering from Disk Failure (Hardware)

This section describes how to recover database clusters to a point immediately before failure, if a hardware failure occurs in the data storage disk or the backup data storage disk.

There are two methods of recovery:

- 14.1.1 Using WebAdmin
- 14.1.2 Using Server Command



Back up the database cluster after recovering it. Backup deletes obsolete archive logs (transaction logs copied to the backup data storage destination), freeing up disk space and reducing the recovery time.

# 14.1.1 Using WebAdmin

Recover the database cluster by following the appropriate recovery procedure below for the disk where the failure occurred.



Recovery operation cannot be performed on an instance that is part of a streaming replication cluster in standby mode.

If disk failure occurs on a standby instance, it may be necessary to delete and re-create the instance.

Recovery operation can be performed on an instance that is part of a streaming replication cluster in "Master" mode. If a recovery operation is performed on a master instance, it will break the replication cluster and streaming replication will stop between the master instance and all its standby instances. In such an event, the standby instances can be promoted to standalone instances or can be deleted and re-created.

# If failure occurred in the data storage disk or the transaction log storage disk

Follow the procedure below to recover the data storage disk or the transaction log storage disk.

1. Stop applications

Stop applications that are using the database.

2. Stop the instance

Stop the instance. Refer to "2.1.1 Using WebAdmin" for information on how to stop an instance. WebAdmin automatically stops instances if recovery of the database cluster is performed without stopping the instance.

3. Recover the failed disk

Replace the disk, and then recover the volume configuration information.

4. Create a tablespace directory

If a tablespace was defined after backup, create a directory for it.

5. Recover the keystore, and enable automatic opening of the keystore

Do the following if the data in the database has been encrypted:

- Restore the keystore to its state at the time of the database backup.
- Enable automatic opening of the keystore.

### 6. Recover the database cluster

Log in to WebAdmin, and in the [Instances] tab, click [Solution] for the error message in the lower-right corner.

### 7. Run recovery

In the [Restore Instance] dialog box, click [Yes].

Instance restore is performed. An instance is automatically started when recovery is successful.



WebAdmin does not support recovery of hash index. If you are using a hash index, then after recovery, execute the REINDEX command to rebuild it. Use of hash indexes is not recommended.

### 8. Resume applications

Resume applications that are using the database.



WebAdmin may be unable to detect disk errors, depending on how the error occurred.

If this happens, refer to "14.10.3 Other Errors" to perform recovery.

## If failure occurred on the backup data storage disk

Follow the procedure below to recover the backup data storage disk.

1. Recover the failed disk

Replace the disk, and then recover the volume configuration information.

2. Recover the backup data

Log in to WebAdmin, and in the [Instances] tab, click [Solution] for the error message.

3. Run backup

 $Perform\ backup\ to\ enable\ recovery\ of\ the\ backup\ data.\ In\ the\ [Backup]\ dialog\ box,\ click\ [Yes].\ The\ backup\ is\ performed.\ An\ instance$  is automatically started when backup is performed.

.....



If you click [Recheck the status], the resources in the data storage destination and the backup data storage destination are reconfirmed. As a result, the following occurs:

- If an error is not detected

The status of the data storage destination and the backup data storage destination returns to normal, and it is possible to perform operations as usual.

- If an error is detected

An error message is displayed in the message list again. Click [Solution], and resolve the problem by following the resolution for the cause of the error displayed in the dialog box.

# 14.1.2 Using Server Command

Recover the database cluster by following the appropriate recovery procedure below for the disk where the failure occurred.

## If failure occurred on the data storage disk or the transaction log storage directory

Follow the procedure below to recover the data storage disk or the transaction log storage directory.

### 1. Stop applications

Stop applications that are using the database.

### 2. Stop the instance

Stop the instance, refer to "2.1.2 Using Commands" for details.

If the instance fails to stop, refer to "14.11 Actions in Response to Failure to Stop an Instance".

### 3. Recover the failed disk

Replace the disk, and then recover the volume configuration information.

- 4. Create a storage destination directory
  - If failure occurred on the data storage disk
     Create a data storage destination directory. If a tablespace was defined, also create a directory for it.
  - If failure occurred on the translation log storage disk
     Create a transaction log storage destination directory.

In [Properties] in Windows(R) Explorer, set appropriate permissions so that only the instance administrator can access the storage destination directory. (Refer to [Help and Support] in Windows(R) for information on [Properties].)



Refer to "Preparing Directories to Deploy Resources" under "Setup" in the Installation and Setup Guide for Server for information on how to create a storage directory.

5. Recover the keystore, and enable automatic opening of the keystore

When the data in the database has been encrypted, restore the keystore to its state at the time of the database backup. Configure automatic opening of the keystore as necessary.

6. Recover the database cluster

Recover the database cluster using the backup data.

Specify the following in the pgx\_rcvall command:

- Specify the data storage location in the -D option. If the -D option is omitted, the value of the PGDATA environment variable is used by default.
- Specify the backup data storage location in the -B option.

## Example

> pgx\_rcvall -D D:\database\inst1 -B E:\backup\inst1



If recovery fails, remove the cause of the error in accordance with the displayed error message and then re-execute the pgx\_rcvall command.

If the message "pgx\_rcvall: an error occurred during recovery" is displayed, then the log recorded when recovery was executed is output after this message. The cause of the error is output in around the last fifteen lines of the log, so remove the cause of the error in accordance with the message and then re-execute the pgx\_rcvall command.

The following message displayed during recovery is output as part of normal operation of pgx\_rcvall command (therefore the user does not need not be concerned).

FATAL: The database system is starting

# 7. Start the instance

Refer to "2.1.2 Using Commands" for information on how to start an instance.

# 8. Resume applications

Resume applications that are using the database.

# If failure occurred on the backup data storage disk

The procedure for recovering the backup data storage disk is described below.

There are two methods of taking action:

- Performing recovery while the instance is active
- Stopping the instance before performing recovery

The following table shows the different steps to be performed depending on whether you stop the instance.

No	Step	Instance stopped	
		No	Yes
1	Confirm that transaction log mirroring has stopped	Y	N
2	Stop output of archive logs	Y	N
3	Stop applications	N	Y
4	Stop the instance	N	Y
5	Recover the failed disk	Y	Y
6	Create a backup data storage destination directory	Y	Y
7	Resume output of archive logs	Y	N
8	Resume transaction log mirroring	Y	N
9	Start the instance	N	Y
10	Run backup	Y	Y
11	Resume applications	N	Y

Y: Required

N: Not required

The procedure is as follows:

# If an instance has not been stopped

1. Confirm that transaction log mirroring has stopped

Use the following SQL function to confirm that transaction log mirroring has stopped.

```
postgres=# SELECT pgx_is_wal_multiplexing_paused();
pgx_is_wal_multiplexing_paused
------t
t
(1 row)
```

If transaction log mirroring has not stopped, then stop it using the following SQL function.

## 2. Stop output of archive logs

Transaction logs may accumulate during replacement of backup storage disk, and if the data storage disk or the transaction log storage disk becomes full, there is a risk that operations may not be able to continue.

To prevent this, use the following methods to stop output of archive logs.

- Changing archive\_command

Specify a command that will surely complete normally, so that archive logs will be regarded as having been output.

If you specify echo, a message is output to the server log, so it may be used as a reference when you conduct investigations.

- Reload the configuration file

Execute the pg\_ctl reload command or the pg\_reload\_conf SQL function to reload the configuration file.

If you simply want to stop output of errors without the risk that operations will not be able to continue, specify an empty string (") in archive\_command and reload the configuration file.

### 3. Recover the failed disk

Replace the disk, and then recover the volume configuration information.

4. Create a backup data storage destination

Create a backup data storage destination.

In [Properties] in Windows(R) Explorer, set appropriate permissions so that only the instance administrator can access the backup data storage destination directory. (Refer to [Help and Support] in Windows(R) for information on [Properties].)

Refer to "3.2.2 Using Server Commands" for information on how to create a backup data storage destination.

5. Resume output of archive logs

Return the archive\_command setting to its original value, and reload the configuration file.

6. Resume transaction log mirroring

Execute the pgx\_resume\_wal\_multiplexing SQL function.

### Example

```
SELECT pgx_resume_wal_multiplexing()
```

## 7. Run backup

Use the  $pgx\_dmpall$  command to back up the database cluster.

Specify the following value in the pgx\_dmpall command:

- Specify the data storage destination in the -D option. If the -D option is omitted, the value of the PGDATA environment variable is used by default.

## Example

```
> pgx_dmpall -D D:\database\inst1
```

# If an instance has been stopped

1. Stop applications

Stop applications that are using the database.

2. Stop the instance

Stop the instance. Refer to "2.1.2 Using Commands" for details.

If the instance fails to stop, refer to "14.11 Actions in Response to Failure to Stop an Instance".

3. Recover the failed disk

Replace the disk, and then recover the volume configuration information.

4. Create a backup data storage destination

Create a backup data storage destination.

In [Properties] in Windows(R) Explorer, set appropriate permissions so that only the instance administrator can access the backup data storage destination directory. (Refer to [Help and Support] in Windows(R) for information on [Properties].)

Refer to "3.2.2 Using Server Commands" for details.

5. Start the instance

Start the instance. Refer to "2.1.2 Using Commands" for information on how to start an instance.

6. Run backup

Use the pgx\_dmpall command to back up the database cluster.

Specify the following value in the pgx\_dmpall command:

- Specify the data storage destination in the -D option. If the -D option is omitted, the value of the PGDATA environment variable is used by default.

## Example

> pgx\_dmpall -D D:\database\inst1

7. Resume applications

Resume applications that are using the database.



- Refer to "pgx\_rcvall" and "pgx\_dmpall" in the Reference for information on the pgx\_rcvall command and pgx\_dmpall command.

......

- Refer to "Write Ahead Log" under "Server Administration" in the PostgreSQL Documentation for information on archive\_command.
- Refer to "B.1 WAL Mirroring Control Functions" for information on pgx\_resume\_wal\_multiplexing.

# 14.2 Recovering from Data Corruption

If data in a disk is logically corrupted and the database does not operate properly, you can recover the database cluster to its state at the time of backup.

There are two methods of recovery:

- 14.2.1 Using WebAdmin
- 14.2.2 Using the pgx\_rcvall Command



- Back up the database cluster after recovering it. Backup deletes obsolete archive logs (transaction logs copied to the backup data storage destination), freeing up disk space and reducing the recovery time.
- If you recover data to a point in the past, a new time series (database update history) will start from that recovery point. When recovery is complete, the recovery point is the latest point in the new time series. When you subsequently recover data to the latest state, the database update is re-executed on the new time series.

# 14.2.1 Using WebAdmin

If using WebAdmin, recover the data to the point immediately prior to data corruption by using the backup data.

Refer to "14.1.1 Using WebAdmin" for details.

# 14.2.2 Using the pgx\_rcvall Command

Recover the database cluster by specifying in the pgx\_rcvall command the date and time of the backup you want to read from. Then reexecute the transaction as required to recover the data.

Follow the procedure below to recover the data storage disk.

1. Stop applications

Stop applications that are using the database.

2. Stop the instance

Stop the instance. Refer to "2.1.2 Using Commands" for information on how to stop an instance.

If the instance fails to stop, refer to "14.11 Actions in Response to Failure to Stop an Instance".

3. Confirm the backup date and time

Pinpoint a date and time prior to the data corruption based on the content of the job log or event log.

4. Recover the keystore, and enable automatic opening of the keystore

When the data in the database has been encrypted, restore the keystore to its state at the time of the database backup. Configure automatic opening of the keystore as necessary.

5. Recover the database cluster

Use the pgx\_rcvall command to recover the database cluster.

Specify the following values in the pgx\_rcvall command:

- Specify the data storage destination in the -D option. If the -D option is omitted, the value of the PGDATA environment variable is used by default.
- Specify the backup storage directory in the -B option.
- Specify the recovery date and time in the -e option.

## Example

In the following examples, "May 20, 2017 10:00:00" is specified as the recovery time.

> pgx\_rcvall -D D:\database\inst1 -B E:\backup\inst1 -e "2017-05-20 10:00:00"



If recovery fails, remove the cause of the error in accordance with the displayed error message and then re-execute the pgx\_rcvall command.

If the message "pgx\_rcvall: an error occurred during recovery" is displayed, then the log recorded when recovery was executed is output after this message. The cause of the error is output in around the last fifteen lines of the log, so remove the cause of the error in accordance with the message and then re-execute the pgx\_rcvall command.

The following message displayed during recovery is output as part of normal operation of pgx\_rcvall command (therefore the user does not need not be concerned).

FATAL: The database system is starting

6. Start the instance

Start the instance. Refer to "2.1.2 Using Commands" for information on how to start an instance.

If necessary, re-execute transaction processing from the specified recovery time, and then resume database operations.



The pgx\_rcvall command cannot accurately recover a hash index. If you are using a hash index, wait for the instance to start and then execute the REINDEX command for the appropriate index.

7. Resume applications

Resume applications that are using the database.



Refer to "pgx\_rcvall" in the Reference for information on the pgx\_rcvall command.

# 14.3 Recovering from an Incorrect User Operation

This section describes how to recover database clusters when data has been corrupted due to erroneous user operations.

There are two methods of recovery:

- 14.3.1 Using WebAdmin
- 14.3.2 Using the pgx\_rcvall Command



- Back up the database cluster after recovering it. Backup deletes obsolete archive logs (transaction logs copied to the backup data storage destination), freeing up disk space and reducing the recovery time.
- If you recover data to a point in the past, a new time series (database update history) will start from that recovery point. When recovery is complete, the recovery point is the latest point in the new time series. When you subsequently recover data to the latest state, the database update is re-executed on the new time series.
- An effective restore point is one created on a time series for which you have made a backup. That is, if you recover data to a point in the past, you cannot use any restore points set after that recovery point. Therefore, once you manage to recover your target past data, make a backup.

# 14.3.1 Using WebAdmin

You can use WebAdmin to recover data to a backup point.



Recovery operation cannot be performed on an instance that is part of a streaming replication cluster in standby mode.

If disk failure occurs on a standby instance, it may be necessary to delete and re-create the instance.

Recovery operation can be performed on an instance that is part of a streaming replication cluster in "Master" mode. If a recovery operation is performed on a master instance, it will break the replication cluster and streaming replication will stop between the master instance and all its standby instances. In such an event, the standby instances can be promoted to standalone instances or can be deleted and re-created.

Follow the procedure below to recover the data in the data storage disk.

1. Stop applications

Stop applications that are using the database.

2. Stop the instance

Stop the instance. Refer to "2.1.1 Using WebAdmin" for information on how to stop an instance.

3. Recover the keystore, and enable automatic opening of the keystore

Do the following if the data in the database has been encrypted:

- Restore the keystore to its state at the time of the database backup.
- Enable automatic opening of the keystore.
- 4. Recover the database cluster

Log in to WebAdmin, and in the [Instances] tab, select the instance to be recovered and click 🚮.



5. Recover to the backup point

In the [Restore Instance] dialog box, click [Yes].

Recovery is performed. An instance is automatically started when recovery is successful.



WebAdmin cannot accurately recover a hash index. If you are using a hash index, then after recovery, execute the REINDEX command for the appropriate index.

6. Resume database operations

If necessary, re-execute transaction processing from the backup point to when an erroneous operation was performed, and then resume database operations.

# 14.3.2 Using the pgx\_rcvall Command

The pgx\_rcvall command recovers database clusters to the restore point created with the server command. Refer to "Setting a restore point" in "3.2.2 Using Server Commands" for information on how to create a restore point.

Follow the procedure below to recover the data in the data storage disk.

1. Stop applications

Stop applications that are using the database.

2. Stop the instance

Stop the instance. Refer to "2.1.2 Using Commands" for information on how to stop an instance.

If the instance fails to stop, refer to "14.11 Actions in Response to Failure to Stop an Instance".

3. Confirm the restore point

Use a restore point recorded in an arbitrary file, as explained in ""3.2.2 Using Server Commands", to determine a restore point prior to the erroneous operation.

4. Recover the keystore, and enable automatic opening of the keystore

When the data in the database has been encrypted, restore the keystore to its state at the time of the database backup. Configure automatic opening of the keystore as necessary.

5. Recover the database cluster

Use the pgx\_rcvall command to recover the database cluster.

Specify the following values in the pgx\_rcvall command:

- Specify the data storage destination in the -D option. If the -D option is omitted, the value of the PGDATA environment variable is used by default.
- Specify the backup data storage destination in the -B option.
- The -n option recovers the data to the specified restore point.

# Example

The following example executes the pgx\_rcvall command with the restore point "batch\_20170503\_1".

> pgx\_rcvall -D D:\database\inst1 -B E:\backup\inst1 -n batch\_20170503\_1



If recovery fails, remove the cause of the error in accordance with the displayed error message and then re-execute the pgx\_rcvall command.

If the message "pgx\_rcvall: an error occurred during recovery" is displayed, then the log recorded when recovery was executed is output after this message. The cause of the error is output in around the last fifteen lines of the log, so remove the cause of the error in accordance with the message and then re-execute the pgx\_rcvall command.

The following message displayed during recovery is output as part of normal operation of pgx\_rcvall (therefore the user does not need not be concerned).

FATAL: The database system is starting

6. Start the instance

Start the instance.

Refer to "2.1.2 Using Commands" for information on how to start an instance.



The pgx\_rcvall command cannot accurately recover a hash index. If you are using a hash index, wait for the instance to start and then execute the REINDEX command for the appropriate index.

7. Restart operation of the database

If necessary, re-execute transaction processing from the specified recovery time to the point when an erroneous operation was performed, and then resume database operations.



Refer to "pgx\_rcvall" in the Reference for information on the pgx\_rcvall command.

# 14.4 Actions in Response to an Application Error

If there is a connection from a client that has been in the waiting state for an extended period, you can minimize performance degradation of the database by closing the problematic connection.

The following methods are available for identifying a connection to be closed:

- view(pg\_stat\_activity) (refer to "14.4.1 When using the view (pg\_stat\_activity)")
- pgAdmin (refer to "14.4.2 Using pgAdmin")

Use the system management function (pg\_terminate\_backend) to disconnect connections.

# 14.4.1 When using the view (pg\_stat\_activity)

When using the view (pg\_stat\_activity), follow the procedure below to close a connection.

1. Use psql command to connect to the postgres database.

```
> psql postgres
psql (12.1)
Type "help" for help.
```

2. Close connections from clients that have been in the waiting state for an extended period.

Use pg\_terminate\_backend() to close connections that have been trying to connect for an extended period.

However, when considering continued compatibility of applications, do not reference or use system catalogs and functions directly in SQL statements. Refer to "Notes on Application Compatibility" in the Application Development Guide for details.

### Example

The following example closes connections where the client has been in the waiting state for at least 60 minutes.



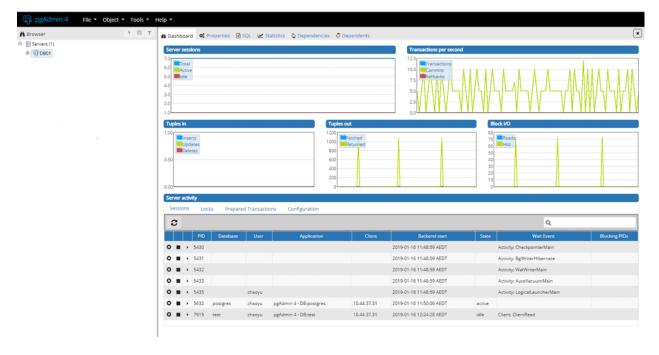
- Refer to "System Administration Functions" under "The SQL Language" in the PostgreSQL Documentation for information on pg\_terminate\_backend.
- Refer to "Notes on Application Compatibility" in the Application Development Guide for information on how to maintain application compatibility.

# 14.4.2 Using pgAdmin

If using pgAdmin, follow the procedure below to close connections.

- 1. In the [Browser] pane, click the relevant database server.
- 2. Close the client connections that have been in a wait state for an extended period.

Click the [Dashboard] tab. In the [Server activity] section, select the connections that have been in an "idle" or "idle in transaction" state for an extended period. For each of these connections, click 3 to close the session.



# 14.5 Actions in Response to an Access Error

If access is denied, grant privileges allowing the instance administrator to operate the following directories, and then re-execute the operation. Also, refer to the event log and the server log, and confirm that the file system has not been mounted as read-only due to a disk error. If the file system has been mounted as read-only, mount it properly and then re-execute the operation.

- Data storage destination
- Tablespace storage destination
- Transaction log storage destination
- Backup data storage destination



Refer to "Preparing Directories to Deploy Resources" under "Setup" in the Installation and Setup Guide for Server for information on the privileges required for the directory.

# 14.6 Actions in Response to Insufficient Space on the Data Storage Destination

If the data storage destination runs out of space, check if the disk contains any unnecessary files and delete them so that operations can continue.

If deleting unnecessary files does not solve the problem, you must migrate data to a disk with larger capacity.

There are two methods of migrating data:

- 14.6.1 Using a Tablespace
- 14.6.2 Replacing the Disk with a Larger Capacity Disk

# 14.6.1 Using a Tablespace

FUJITSU Enterprise Postgres enables you to use a tablespace to change the storage destination of database objects, such as tables and indexes, to a different disk.

The procedure is as follows:

1. Create a tablespace

Use the CREATE TABLESPACE command to create a new tablespace in the prepared disk.

2. Modify the tablespace

Use the ALTER TABLE command to modify tables for the newly defined tablespace.



Refer to "SQL Commands" under "Reference" in the PostgreSQL Documentation for information on the CREATE TABLESPACE command and ALTER TABLE command.

# 14.6.2 Replacing the Disk with a Larger Capacity Disk

Before replacing the disk with a larger capacity disk, migrate resources at the data storage destination using the backup and recovery features.

There are two methods of performing backup and recovery:

- 14.6.2.1 Using WebAdmin

## - 14.6.2.2 Using Server Commands

The following sections describe procedures that use each of these methods to replace the disk and migrate resources at the data storage destination.



- Before replacing the disk, stop applications and instances that are using the database.
- It is recommended that you back up the database cluster following recovery. Backup deletes obsolete archive logs (transaction logs copied to the backup data storage destination), freeing up disk space and reducing the recovery time.

# 14.6.2.1 Using WebAdmin

Follow the procedure below to replace the disk and migrate resources at the data storage destination by using WebAdmin.

1. Back up files

If the disk at the data storage destination contains any required files, back up the files. It is not necessary to back up the data storage destination.

2. Stop applications

Stop applications that are using the database.

3. Back up the database cluster

Back up the latest resources at the data storage destination. Refer to "3.2.1 Using WebAdmin" for details.

4. Stop the instance

Stop the instance. Refer to "2.1.1 Using WebAdmin" for information on how to stop an instance.

5. Replace with a larger capacity disk

Replace the disk. Then, recover the volume configuration information.

6. Recover the database cluster

Log in to WebAdmin, and perform recovery operations. Refer to steps 4 ("Create a tablespace directory") to 7 ("Run recovery") under "If failure occurred in the data storage disk or the transaction log storage disk" in "14.1.1 Using WebAdmin" for information on the procedure. An instance is automatically started when recovery is successful.

7. Resume applications

Resume applications that are using the database.

8. Restore the files

Restore the files backed up in step 1.

# 14.6.2.2 Using Server Commands

Follow the procedure below to replace the disk and migrate resources at the data storage destination by using server commands.

1. Back up files

If the disk at the data storage destination contains any required files, back up the files. It is not necessary to back up the data storage destination.

2. Stop applications

Stop applications that are using the database.

3. Back up the database cluster

Back up the latest resources at the data storage destination. Refer to "3.2.2 Using Server Commands" for details.

## 4. Stop the instance

After backup is complete, stop the instance. Refer to "2.1.2 Using Commands" for information on how to stop an instance.

If the instance fails to stop, refer to "14.11 Actions in Response to Failure to Stop an Instance".

5. Replace with a larger capacity disk

Replace the disk. Then, recover the volume configuration information.

6. Create a data storage destination

Create a data storage destination. If a tablespace was defined, also create a directory for it.

In [Properties] in Windows(R) Explorer, set appropriate permissions so that only the instance administrator can access the data storage destination directory. (Refer to [Help and Support] in Windows(R) for information on [Properties].)

7. Recover the keystore, and enable automatic opening of the keystore

When the data in the database has been encrypted, restore the keystore to its state at the time of the database backup. Configure automatic opening of the keystore as necessary.

8. Recover the database cluster

Use the pgx\_rcvall command to recover the database cluster.

- Specify the data storage destination in the -D option. If the -D option is omitted, the value of the PGDATA environment variable is used by default.
- Specify the backup storage directory in the -B option.

## Example

> pgx\_rcvall -D D:\database\inst1 -B E:\backup\inst1



If recovery fails, remove the cause of the error in accordance with the displayed error message and then re-execute the pgx\_rcvall command.

If the message "pgx\_rcvall: an error occurred during recovery" is displayed, then the log recorded when recovery was executed is output after this message. The cause of the error is output in around the last fifteen lines of the log, so remove the cause of the error in accordance with the message and then re-execute the pgx\_rcvall command.

The following message displayed during recovery is output as part of normal operation of pgx\_rcvall (therefore the user does not need not be concerned).

.....

FATAL: The database system is starting



Refer to "pgx\_rcvall" in the Reference for information on the pgx\_rcvall command.

9. Start the instance

Start the instance.

Refer to "2.1.2 Using Commands" for information on how to start an instance.



The pgx\_rcvall command cannot accurately recover a hash index. If you are using a hash index, wait for the pgx\_rcvall command to end and then execute the REINDEX command for the appropriate index.

### 10. Resume applications

Resume applications that are using the database.

11. Restore files

Restore the files backed up in step 1.

# 14.7 Actions in Response to Insufficient Space on the Backup Data Storage Destination

If space runs out on the backup data storage destination, check if the disk contains any unnecessary files and delete them, and then make a backup as required.

If deleting unnecessary files does not solve the problem, take the following action:

- 14.7.1 Temporarily Saving Backup Data
- 14.7.2 Replacing the Disk with a Larger Capacity Disk

# 14.7.1 Temporarily Saving Backup Data

This method involves temporarily moving backup data to a different directory, saving it there, and securing disk space on the backup data storage destination so that a backup can be made normally.

Use this method if you need time to prepare a larger capacity disk.

If space runs out on the backup data storage destination, archive logs can no longer be stored in the backup data storage destination. As a result, transaction logs continue to accumulate in the data storage destination or the transaction log storage destination.

If action is not taken soon, the transaction log storage destination will become full, and operations may not be able to continue.

To prevent this, secure space in the backup data storage destination, so that archive logs can be stored.

There are two methods of taking action:

- 14.7.1.1 Using WebAdmin
- 14.7.1.2 Using Server Commands

# 14.7.1.1 Using WebAdmin

Follow the procedure below to recover the backup data storage disk.

1. Temporarily save backup data

Move backup data to a different directory and temporarily save it, and secure space in the backup data storage destination directory.

The reason for saving the backup data is so that the data in the data storage destination can be recovered even if it is corrupted before you perform recovery. If there is no disk at the save destination and you consider that there is no risk of corruption at the data storage destination, delete the backup data.

The following example saves backup data from the backup data storage destination directory (E:\backup\inst1) under F:\mnt\usb\backup.

### Example

```
> mkdir F:\mnt\usb\backup
> move E:\backup\inst1\* F:\mnt\usb\backup
```

Note: Place the temporary backup destination directory in a location where it will not impact on operating system resources or FUJITSU Enterprise Postgres resources.

2. Back up the database cluster

Back up the latest resources at the data storage destination. Refer to "3.2.1 Using WebAdmin" for details.

3. Delete temporarily saved backup data

If backup completes normally, the temporarily saved backup data becomes unnecessary and is deleted.

The following example deletes backup data that was temporarily saved in F:\mnt\usb.

### Example

```
> rmdir /S /Q F:\mnt\usb\backup
```

# 14.7.1.2 Using Server Commands

The following describes the procedure for recovering the backup storage disk.

There are two methods of taking action:

- Performing recovery while the instance is active
- Stopping the instance before performing recovery

The following table shows the different steps to be performed depending on whether you stop the instance.

No	Step	Instance stopped	
		No	Yes
1	Stop transaction log mirroring	Y	N
2	Stop output of archive logs	Y	N
3	Stop applications	N	Y
4	Stop the instance	N	Y
5	Temporarily save backup data	Y	Y
6	Resume output of archive logs	Y	N
7	Resume transaction log mirroring	Y	N
8	Start an instance	N	Y
9	Run backup	Y	Y
10	Resume applications	N	Y
11	Delete temporarily saved backup data	Y	Y

Y: Required

N: Not required

The procedure is as follows:

# Performing recovery while the instance is active

1. Stop transaction log mirroring

Stop transaction log mirroring.

2. Stop output of archive logs

Transaction logs may accumulate during replacement of backup storage disk, and if the data storage disk or the transaction log storage disk becomes full, there is a risk that operations may not be able to continue.

To prevent this, use the following methods to stop output of archive logs.

- Changing the archive\_command parameter

Specify a command that will surely complete normally, so that archive logs will be regarded as having been output.

If you specify echo, a message is output to the server log, so it may be used as a reference when you conduct investigations.

- Reloading the configuration file

Run the pg\_ctl reload command or the pg\_reload\_conf SQL function.

If you simply want to stop output of errors without the risk that operations will not be able to continue, specify an empty string (") in archive\_command and reload the configuration file.

## 3. Temporarily save backup data

Move backup data to a different directory and temporarily save it, and secure space in the backup data storage destination directory.

The reason for saving the backup data is so that the data in the data storage destination can be recovered even if it is corrupted before you perform the next step. If there is no disk at the save destination and you consider that there is no risk of corruption at the data storage destination, delete the backup data.

The following example saves backup data from the backup data storage destination directory (E:\backup\inst1) under F:\mnt\usb\backup.

## Example

```
> mkdir F:\mnt\usb\backup
> move E:\backup\inst1\* F:\mnt\usb\backup
```

Note: Place the temporary backup destination directory in a location where it will not impact on operating system resources or FUJITSU Enterprise Postgres resources.

### 4. Resume output of archive logs

Return the archive\_command setting to its original value, and reload the configuration file.

5. Resume transaction log mirroring

Execute the pgx\_resume\_wal\_multiplexing SQL function.

### Example

```
SELECT pgx_resume_wal_multiplexing()
```

## 6. Run backup

Use the pgx\_dmpall command to back up the database cluster.

Specify the following option in the pgx\_dmpall command:

- Specify the directory of the data storage destination in the -D option. If the -D option is omitted, the value of the PGDATA environment variable is used by default.

### Example

```
> pgx_dmpall -D D:\database\inst1
```

# 7. Delete temporarily saved backup data

If backup completes normally, the temporarily saved backup data becomes unnecessary and is deleted.

The following example deletes backup data that was temporarily saved in F:\mnt\usb.

### Example

```
> rmdir /S /Q F:\mnt\usb\backup
```

## If an instance has been stopped

1. Stop applications

Stop applications that are using the database.

2. Stop the instance

Stop the instance. Refer to "2.1.2 Using Commands" for details.

If the instance fails to stop, refer to "14.11 Actions in Response to Failure to Stop an Instance".

3. Temporarily save backup data

Move backup data to a different directory and temporarily save it, and secure space in the backup data storage destination directory.

The reason for saving the backup data is so that the data in the data storage destination can be recovered even if it is corrupted before you perform recovery. If there is no disk at the save destination and you consider that there is no risk of corruption at the data storage destination, delete the backup data.

The following example saves backup data from the backup data storage destination directory (E:\backup\inst1) under F:\mnt\usb\backup.

### Example

```
> mkdir F:\mnt\usb\backup
> move E:\backup\inst1\* F:\mnt\usb\backup
```

Note: Place the temporary backup destination directory in a location where it will not impact on operating system resources or FUJITSU Enterprise Postgres resources.

4. Start the instance

Start the instance. Refer to "2.1.2 Using Commands" for information on how to start an instance.

5. Run backup

Use the pgx\_dmpall command to back up the database cluster.

Specify the following value in the pgx\_dmpall command:

- Specify the data storage destination in the -D option. If the -D option is omitted, the value of the PGDATA environment variable is used by default.

## Example

```
> pgx_dmpall -D D:\database\inst1
```

6. Resume applications

Resume applications that are using the database.

7. Delete temporarily saved backup data

If backup completes normally, the temporarily saved backup data becomes unnecessary and is deleted.

The following example deletes backup data that was temporarily saved in  $F:\mbox{\sc mnt}\sc s$ .

## Example

```
> rmdir /S /Q F:\mnt\usb\backup
```



 $- \ Refer to \ "pgx\_rcvall" \ and \ "pgx\_dmpall" \ in \ the \ Reference \ for \ information \ on \ the \ pgx\_rcvall \ command \ and \ pgx\_dmpall \ command.$ 

......

- Refer to "Write Ahead Log" under "Server Administration" in the PostgreSQL Documentation for information on archive\_command.

- Refer to "B.1 WAL Mirroring Control Functions" for information on the pgx\_is\_wal\_multiplexing\_paused and pgx\_resume\_wal\_multiplexing.

# 14.7.2 Replacing the Disk with a Larger Capacity Disk

This method involves replacing the disk at the backup data storage destination with a larger capacity disk, so that it does not run out of free space again. After replacing the disk, back up data to obtain a proper backup.

There are two methods of performing backup:

- 14.7.2.1 Using WebAdmin
- 14.7.2.2 Using Server Commands



Before replacing the disk, stop applications that are using the database.

# 14.7.2.1 Using WebAdmin

Follow the procedure below to recover the backup storage disk.

1. Back up files

If the disk at the backup data storage destination contains any required files, back up the files. It is not necessary to back up the backup data storage destination.

2. Temporarily save backup data

Save the backup data to a different directory.

The reason for saving the backup data is so that the data in the data storage destination can be recovered even if it is corrupted before you perform the next step. If there is no disk at the save destination and you consider that there is no risk of corruption at the data storage destination, delete the backup data.

The following example saves backup data from the backup data storage destination directory (E:\backup\inst1) under F:\mnt\usb \backup.

## Example

```
> mkdir F:\mnt\usb\backup
> move E:\backup\inst1\* F:\mnt\usb\backup
```

Note: Place the temporary backup destination directory in a location where it will not impact on operating system resources or FUJITSU Enterprise Postgres resources.

3. Replace with a larger capacity disk

Replace the disk. Then, recover the volume configuration information.

4. Run backup

Log in to WebAdmin, and perform recovery operations. Refer to steps 2 ("Recover the backup data") and 3 ("Run backup") under "If failure occurred on the backup storage disk" in "14.1.1 Using WebAdmin".

5. Restore files

Restore the files backed up in step 1.

6. Delete temporarily saved backup data

If backup completes normally, the temporarily saved backup data becomes unnecessary and is deleted.

The following example deletes backup data that was temporarily saved in F:\mnt\usb.

> rmdir /S /Q F:\mnt\usb\backup

# 14.7.2.2 Using Server Commands

The procedure for recovering the backup data storage disk is described below.

There are two methods of taking action:

- Performing recovery while the instance is active
- Stopping the instance before performing recovery

The following table shows the different steps to be performed depending on whether you stop the instance.

No	Step	Instance stopped	
		No	Yes
1	Back up files	Y	Y
2	Temporarily save backup data	Y	Y
3	Confirm that transaction log mirroring has stopped	Y	N
4	Stop output of archive logs	Y	N
5	Stop applications	N	Y
6	Stop the instance	N	Y
7	Replace with a larger capacity disk	Y	Y
8	Create a backup storage directory	Y	Y
9	Resume output of archive logs	Y	N
10	Resume transaction log mirroring	Y	N
11	Start the instance	N	Y
12	Run backup	Y	Y
13	Resume applications	N	Y
14	Restore files	Y	Y
15	Delete temporarily saved backup data	Y	Y

Y: Required

N: Not required

The procedure is as follows:

If an instance has not been stopped

1. Back up files

If the disk at the backup data storage destination contains any required files, back up the files. It is not necessary to back up the backup data storage destination.

2. Temporarily save backup data

Save the backup data to a different directory.

The reason for saving the backup data is so that the data in the data storage destination can be recovered even if it is corrupted before you perform the next step. If there is no disk at the save destination and you consider that there is no risk of corruption at the data storage destination, delete the backup data.

The following example saves backup data from the backup data storage destination directory (E:\backup\inst1) under F:\mnt\usb \backup.

## Example

```
> mkdir F:\mnt\usb\backup
> move E:\backup\inst1\* F:\mnt\usb\backup
```

3. Confirm that transaction log mirroring has stopped

Use the following SQL function to confirm that transaction log mirroring has stopped.

```
postgres=# SELECT pgx_is_wal_multiplexing_paused();
pgx_is_wal_multiplexing_paused
-----
t
(1 row)
```

If transaction log mirroring has not stopped, then stop it using the following SQL function.

4. Stop output of archive logs

Transaction logs may accumulate during replacement of backup storage disk, and if the data storage destination disk or the transaction log storage destination disk becomes full, there is a risk that operations may not be able to continue.

To prevent this, use the following methods to stop output of archive logs.

- Changing the archive\_command parameter

Specify a command that will surely complete normally, so that archive logs will be regarded as having been output.

If you specify echo, a message is output to the server log, so it may be used as a reference when you conduct investigations.

- Reloading the configuration file

Run the pg\_ctl reload command or the pg\_reload\_conf SQL function.

If you simply want to stop output of errors without the risk that operations will not be able to continue, specify an empty string (") in archive\_command and reload the configuration file.

5. Replace with a larger capacity disk

Replace the disk. Then, recover the volume configuration information.

6. Create a backup data storage destination

Create a backup data storage destination.

In [Properties] in Windows(R) Explorer, set appropriate permissions so that only the instance administrator can access the backup data storage destination directory. (Refer to [Help and Support] in Windows(R) for information on [Properties].)

Refer to "3.2.2 Using Server Commands" for details.

7. Resume output of archive logs

Return the archive\_command setting to its original value, and reload the configuration file.

8. Resume transaction log mirroring

Execute the pgx\_resume\_wal\_multiplexing SQL function.

### Example

```
SELECT pgx_resume_wal_multiplexing()
```

## 9. Run backup

Use the pgx\_dmpall command to back up the database cluster.

Specify the following value in the pgx\_dmpall command:

- Specify the data storage destination in the -D option. If the -D option is omitted, the value of the PGDATA environment variable is used by default.

### Example

```
> pgx_dmpall -D D:\database\inst1
```

## 10. Restore files

Restore the files backed up in step 1.

### 11. Delete temporarily saved backup data

If backup completes normally, the temporarily saved backup data becomes unnecessary and is deleted.

The following example deletes backup data that was temporarily saved in F:\mnt\usb.

## Example

```
> rmdir /S /Q F:\mnt\usb\backup
```

### If an instance has been stopped

### 1. Back up files

If the disk at the backup data storage destination contains any required files, back up the files. It is not necessary to back up the backup data storage destination.

# 2. Temporarily save backup data

Save the backup data to a different directory.

The reason for saving the backup data is so that the data in the data storage destination can be recovered even if it is corrupted before you perform the next step. If there is no disk at the save destination and you consider that there is no risk of corruption at the data storage destination, delete the backup data.

The following example saves backup data from the backup data storage destination directory (E:\backup\inst1) under F:\mnt\usb \backup.

## Example

```
> mkdir F:\mnt\usb\backup
> move E:\backup\inst1\* F:\mnt\usb\backup
```

Note: Place the temporary backup destination directory in a location where it will not impact on operating system resources or FUJITSU Enterprise Postgres resources.

# 3. Stop applications

Stop applications that are using the database.

## 4. Stop the instance

Stop the instance. Refer to "2.1.2 Using Commands" for information on how to stop an instance.

If the instance fails to stop, refer to "14.11 Actions in Response to Failure to Stop an Instance".

5. Replace with a larger capacity disk

Replace the disk. Then, recover the volume configuration information.

## 6. Create a backup data storage destination

Create a backup data storage destination.

In [Properties] in Windows(R) Explorer, set appropriate permissions so that only the instance administrator can access the backup data storage destination directory. (Refer to [Help and Support] in Windows(R) for information on [Properties].)

Refer to "3.2.2 Using Server Commands" for details.

7. Start the instance

Start the instance. Refer to "2.1.2 Using Commands" for information on how to start an instance.

8. Run backup

Use the pgx\_dmpall command to back up the database cluster.

Specify the following value in the pgx\_dmpall command:

- Specify the data storage destination in the -D option. If the -D option is omitted, the value of the PGDATA environment variable is used by default.

#### Example

```
> pgx_dmpall -D D:\database\inst1
```

9. Resume applications

Resume applications that are using the database.

10. Restore files

Restore the files backed up in step 1.

11. Delete temporarily saved backup data

If backup completes normally, the temporarily saved backup data becomes unnecessary and is deleted.

The following example deletes backup data that was temporarily saved in F:\mnt\usb.

## Example

> rmdir /S /Q F:\mnt\usb\backup



- $\ Refer to \ "pgx\_rcvall" \ and \ "pgx\_dmpall" \ in \ the \ Reference \ for \ information \ on \ the \ pgx\_rcvall \ command \ and \ pgx\_dmpall \ command.$
- Refer to "Write Ahead Log" under "Server Administration" in the PostgreSQL Documentation for information on archive\_command.
- Refer to "B.1 WAL Mirroring Control Functions" for information on the pgx\_is\_wal\_multiplexing\_paused and pgx\_resume\_wal\_multiplexing.

# 14.8 Actions in Response to Insufficient Space on the Transaction Log Storage Destination

If the transaction log storage destination runs out of space, check if the disk contains any unnecessary files and delete them so that operations can continue.

If deleting unnecessary files does not solve the problem, you must migrate data to a disk with larger capacity.

# 14.8.1 Replacing the Disk with a Larger Capacity Disk

Before replacing the disk with a larger capacity disk, migrate resources at the transaction log storage destination using the backup and recovery features.

There are two methods of performing backup and recovery:

- 14.8.1.1 Using WebAdmin
- 14.8.1.2 Using Server Commands

The following sections describe procedures that use each of these methods to replace the disk and migrate resources at the transaction log storage destination.



- Before replacing the disk, stop applications that are using the database.
- It is recommended that you back up the database cluster following recovery. Backup deletes obsolete archive logs (transaction logs copied to the backup data storage destination), freeing up disk space and reducing the recovery time.

# 14.8.1.1 Using WebAdmin

Follow the procedure below to replace the disk and migrate resources at the transaction log storage destination by using WebAdmin.

1. Back up files

If the disk at the transaction log storage destination contains any required files, back up the files. It is not necessary to back up the transaction log storage destination.

2. Back up the database cluster

Back up the latest data storage destination resources and transaction log storage destination resources (refer to "3.2.1 Using WebAdmin" for details).

3. Stop applications

Stop applications that are using the database.

4. Stop the instance

Stop the instance. Refer to "2.1.1 Using WebAdmin" for information on how to stop an instance. WebAdmin automatically stops instances if recovery of the database cluster is performed without stopping the instance.

5. Replace with a larger capacity disk

Replace the disk. Then, recover the volume configuration information.

6. Create a tablespace directory

If a tablespace was defined after backing up, create a directory for it.

7. Recover the keystore, and enable automatic opening of the keystore

Do the following if the data in the database has been encrypted:

- Restore the keystore to its state at the time of the database backup.
- Enable automatic opening of the keystore.
- 8. Recover the database cluster

Log in to WebAdmin, and perform recovery operations. Refer to steps 4 ("Create a tablespace directory") to 7 ("Run Recovery") under "If failure occurred in the data storage disk or the transaction log storage disk "in "14.1.1 Using WebAdmin" for information on the procedure. An instance is automatically started when recovery is successful.

9. Resume applications

Resume applications that are using the database.

10. Restore files

Restore the files backed up in step 1.

# 14.8.1.2 Using Server Commands

Follow the procedure below to replace the disk and migrate resources at the transaction log storage destination by using server commands.

## 1. Back up files

If the disk at the transaction log storage destination contains any required files, back up the files. It is not necessary to back up the transaction log storage destination.

## 2. Back up the database cluster

Use server commands to back up the latest data storage destination resources and transaction log storage destination resources. Refer to "3.2.2 Using Server Commands" for information on how to perform backup.

## 3. Stop applications

Stop applications that are using the database.

## 4. Stop the instance

After backup is complete, stop the instance. Refer to "2.1.2 Using Commands" for information on how to stop an instance.

If the instance fails to stop, refer to "14.11 Actions in Response to Failure to Stop an Instance".

### 5. Replace with a larger capacity disk

Replace the disk. Then, recover the volume configuration information.

### 6. Create a transaction log storage destination

Create a transaction log storage destination. If a tablespace was defined, also create a directory for it.

In [Properties] in Windows(R) Explorer, set appropriate permissions so that only the instance administrator can access the transaction log destination directory. (Refer to [Help and Support] in Windows(R) for information on [Properties].)

### 7. Recover the keystore, and enable automatic opening of the keystore

When the data in the database has been encrypted, restore the keystore to its state at the time of the database backup. Configure automatic opening of the keystore as necessary.

## 8. Recover the database cluster

Use the pgx\_rcvall command to recover the database cluster.

- Specify the data storage destination in the -D option. If the -D option is omitted, the value of the PGDATA environment variable is used by default.
- Specify the backup storage directory in the -B option.

# Example

> pgx\_rcvall -D D:\database\inst1 -B E:\backup\inst1



If recovery fails, remove the cause of the error in accordance with the displayed error message and then re-execute the pgx\_rcvall command.

If the message "pgx\_rcvall: an error occurred during recovery" is displayed, then the log recorded when recovery was executed is output after this message. The cause of the error is output in around the last fifteen lines of the log, so remove the cause of the error in accordance with the message and then re-execute the pgx\_rcvall command.

The following message displayed during recovery is output as part of normal operation of pgx\_rcvall command (therefore the user does not need not be concerned).

FATAL: The database system is starting



Refer to "pgx\_rcvall" in the Reference for information on the pgx\_rcvall command.

### 9. Start the instance

Start the instance.

Refer to "2.1.2 Using Commands" for information on how to start an instance.



The pgx\_revall command cannot accurately recover a hash index. If you are using a hash index, wait for the instance to start and then execute the REINDEX command for the appropriate index.

# 10. Resume applications

Resume applications that are using the database.

### 11. Restore files

Restore the files backed up in step 1.

# 14.9 Errors in More Than One Storage Disk

If an error occurs in the storage destination disks or resources are corrupted, determine the cause of the error from event logs and server logs and remove the cause.

If errors occur in either of the following combinations, you cannot recover the database.

Recreate the instance, and rebuild the runtime environment.

Data storage disk	Transaction log storage disk	Backup data storage disk
Error	-	Error
-	Error	Error



Refer to "Setup" in the Installation and Setup Guide for Server for information on how to create an instance and build the runtime environment.

# 14.10 Actions in Response to Instance Startup Failure

If an instance fails to start, refer to the event log and the server log, and determine the cause of the failure.

If using WebAdmin, remove the cause of the error. Then, click [Solution] and [Recheck the status] and confirm that the instance is in the normal state.

The following sections describe common causes of errors and the actions to take.

# 14.10.1 Errors in the Configuration File

If you have directly edited the configuration file using a text editor or changed the settings using WebAdmin, refer to the event log and the server log, confirm that no messages relating to the files below have been output.

- postgresql.conf
- pg\_hba.conf



Refer to the following for information on the parameters in the configuration file:

- "Configuring Parameters" in the Installation and Setup Guide for Server
- "Appendix A Parameters"
- "Server Configuration" and "Client Authentication" under "Server Administration" in the PostgreSQL Documentation

# 14.10.2 Errors Caused by Power Failure or Mounting Issues

If mounting is cancelled after restarting the server, for example, because the disk device for each storage destination disk was not turned on, or because automatic mounting has not been set, then starting an instance will fail.

Refer to "14.14.2 Errors Caused by Power Failure or Mounting Issues", and take actions accordingly.

# 14.10.3 Other Errors

This section describes the recovery procedure to be used if you cannot take any action or the instance cannot start even after you have referred to the event log and the server log.

There are two methods of recovery:

- 14.10.3.1 Using WebAdmin
- 14.10.3.2 Using Server Commands

Note that recovery will not be possible if there is an error at the backup data storage destination. If the problem cannot be resolved, contact Fujitsu technical support.

# 14.10.3.1 Using WebAdmin

Follow the procedure below to perform recovery.

- Delete the data storage destination directory and the transaction log storage destination directory
   Back up the data storage destination directory and the transaction log storage destination directory before deleting them.
- 2. Reconfirm the status

Log in to WebAdmin, and in the [Instances] tab, click [Solution] for the error message.

Click [Recheck the status] to reconfirm the storage destination resources.

3. Run recovery

Restore the database cluster after WebAdmin detects an error.

Refer to "14.2.1 Using WebAdmin" for details.

# 14.10.3.2 Using Server Commands

Follow the procedure below to recover the database.

- Delete the data storage destination directory and the transaction log storage destination directory
   Save the data storage destination directory and the transaction log storage destination directory, and then delete them.
- 2. Execute recovery

Use the pgx\_rcvall command to recover the database cluster.

Refer to "14.2.2 Using the pgx\_rcvall Command" for details.

# 14.11 Actions in Response to Failure to Stop an Instance

If an instance fails to stop, refer to the event log and the server log, and determine the cause of the failure.

If the instance cannot stop despite taking action, perform the following operation to stop the instance.

There are two methods of recovery:

- 14.11.1 Using WebAdmin
- 14.11.2 Using Server Commands

### 14.11.1 Using WebAdmin

In the [Instances] tab, click and select the Fast stop mode or the Immediate stop mode to stop the instance. Forcibly terminate the server process from WebAdmin if the instance cannot be stopped.

Refer to "2.1.1 Using WebAdmin" for information on the stop modes.

### 14.11.2 Using Server Commands

There are three methods:

- Stopping the Instance Using the Fast Mode

If backup is in progress, then terminate it, roll back all executing transactions, forcibly close client connections, and then stop the instance.

- Stopping the Instance Using the Immediate Mode

Forcibly terminate the instance immediately. A crash recovery is run when the instance is restarted.

- Forcibly Stopping the Server Process

Reliably stops the server process when the other methods are unsuccessful.

#### 14.11.2.1 Stopping the Instance Using the Fast Mode

Specify "-m fast" in the pg\_ctl command to stop the instance.

If the instance fails to stop when you use this method, stop the instance as described in "14.11.2.2 Stopping the Instance Using the Immediate Mode" or "14.11.2.3 Forcibly Stopping the Server Process".



> pg\_ctl stop -D D:\database\inst1 -m fast

#### 14.11.2.2 Stopping the Instance Using the Immediate Mode

Specify "-m immediate " in the pg\_ctl command to stop the instance.

If the instance fails to stop when you use this method, stop the instance as described in "14.11.2.3 Forcibly Stopping the Server Process".

.....



> pg\_ctl stop -D D:\database\inst1 -m immediate

### 14.11.2.3 Forcibly Stopping the Server Process

If both the Fast mode and the Immediate mode fail to stop the instance, use the kill parameter of the pg\_ctl command to forcibly stop the server process.

.....

The procedure is as follows:

1. Execute the wmic command to identify the process ID of the server process.

```
:

"C:\Program Files\Fujitsu\fsepv<x>server64\bin\postgres.exe" -D "D:\database\inst1"

postgres.exe 896
:
```

The postgres.exe process ID(896) that indicates the data storage destination directory of the applicable instance in the -D option becomes the server process.

2. Forcibly stop the server process

As instance manager, forcibly stop the server process using the pg\_ctl command.

```
c:\>pg_ctl kill QUIT 896
```

# 14.12 Actions in Response to Failure to Create a Streaming Replication Standby Instance

When creating a streaming replication standby instance using WebAdmin, if the instance creation fails, refer to the event log and the server log, and determine the cause of the failure.

When an error occurs in the creation of the standby instance using WebAdmin, it is unlikely that the partially created standby instance can be resumed to complete the operation.

In such a scenario, fix the cause of the error, delete the partially created standby instance, and then create a new standby instance. This recommendation is based on the following assumptions:

- As the instance is yet to be created completely, there are no applications connecting to the database.
- The standby instance is in error state and is not running.
- There are no backups for the standby instance and as a result, it cannot be recovered.



Refer to "Deleting Instances" in the Installation and Setup Guide for details on how to delete an instance.

# 14.13 Actions in Response to Error in a Distributed Transaction

If a system failure (such as server failure) occurs in an application that uses distributed transactions (such as .NET TransactionScope), then transactions may be changed to the in-doubt state. At that point, resources accessed by the transaction will be locked, and rendered unusable by other transactions.

The following describes how to check for in-doubt transactions, and how to resolve them.

#### How to check for in-doubt transactions

The following shows how to check for them:

If the server fails

1. An in-doubt transaction will have occurred if a message similar to the one below is output to the log when the server is restarted. Example

```
LOG: Restoring prepared transaction 2103.
```

2. Refer to system view pg\_prepared\_xacts to obtain information about the prepared transaction.

If the transaction identifier of the prepared transaction in the list (in the transaction column of pg\_prepared\_xacts) is the same as the identifier of the in-doubt transaction obtained from the log output when the server was restarted, then that row is the information about the in-doubt transaction.

#### Example

Information about the in-doubt transaction is output to the row with the transaction ID 2103 in the transaction column.

#### If the client fails

If there are no clients connected and there is a prepared transaction in pg\_prepared\_xacts, then you can determine that the transaction is in the in-doubt state.

If at least one client is connected and there is a prepared transaction in pg\_prepared\_xacts, you cannot determine whether there is a transaction in the in-doubt state. In this case, use the following query to determine the in-doubt transaction from the acquired database name, user name, the time PREPARE TRANSACTION was executed, and the information about the table name accessed.

```
\label{lem:select} select \ gid, x. database, owner, prepared, l.relation:: regclass \ as \ relation \ from \ pg\_prepared\_xacts \ x \\ left \ join \ pg\_locks \ l \ on \ l.virtual transaction = '-1/' | | x. transaction \ and \ l.locktype='relation'; \\ \\ \\
```

If it still cannot be determined from this information, wait a few moments and then check pg\_prepared\_xacts again.

If there is a transaction that has continued since the last time you checked, then it is likely that it is the one in the in-doubt state.



As you can see from the explanations in this section, there is no one way to definitively determine in-doubt transactions.

Consider collecting other supplementary information (for example, logging on the client) or performing other operations (for example, allocating database users per job).

......

#### How to resolve in-doubt transactions

From the system view pg\_prepared\_xacts mentioned above, obtain the global transaction identifier (in the gid column of pg\_prepared\_xacts) for the in-doubt transaction, and issue either a ROLLBACK PREPARED statement or COMMIT PREPARED statement to resolve the in-doubt transaction.

# Example

- Rolling back in-doubt transactions

```
postgres=# rollback prepared '374cc221-f6dc-4b73-9d62-d4fec9b430cd';
ROLLBACK PREPARED
```

- Committing in-doubt transactions

```
postgres=# commit prepared '374cc221-f6dc-4b73-9d62-d4fec9b430cd';
COMMIT PREPARED
```

### 14.14 I/O Errors Other than Disk Failure

Even if a disk is not defective, the same input-output error messages, as those generated when the disk is defective, may be output.

A few examples of such errors are given below. The appropriate action for each error is explained respectively.

- 14.14.1 Network Error with an External Disk
- 14.14.2 Errors Caused by Power Failure or Mounting Issues

### 14.14.1 Network Error with an External Disk

This is an error that occurs in the network path to/from an external disk.

Determine the cause of the error by checking the information in the event log and the server log, the disk access LED, network wiring, and network card status. Take appropriate action to remove the cause of the error, for example, replace problematic devices.

### 14.14.2 Errors Caused by Power Failure or Mounting Issues

These are errors that occur when the disk device is not turned on, automatic mounting of the disk was not set, or mounting was accidentally cancelled.

In this case, check the information in the event log and the server log, the disk access LED, and whether the disk is mounted correctly. If problems are detected, take appropriate action.

If mounting has been cancelled, it is possible that mounting was accidentally cancelled, or the existing setting (automatic mounting at the time of starting the operating system) has been changed so that mounting is not performed automatically. In this case, set the mounting to be performed automatically.

# 14.15 Anomaly Detection and Resolution

The following operations performed via the command line interface will result in an anomaly in WebAdmin:

- Changes to the port and backup\_destination parameters in postgresql.conf
- Changes to Mirroring Controller configuration of cluster replication added via WebAdmin

This section describes when WebAdmin checks for such anomalies, and what takes place when an anomaly is detected.

### 14.15.1 Port Number and Backup Storage Path Anomalies

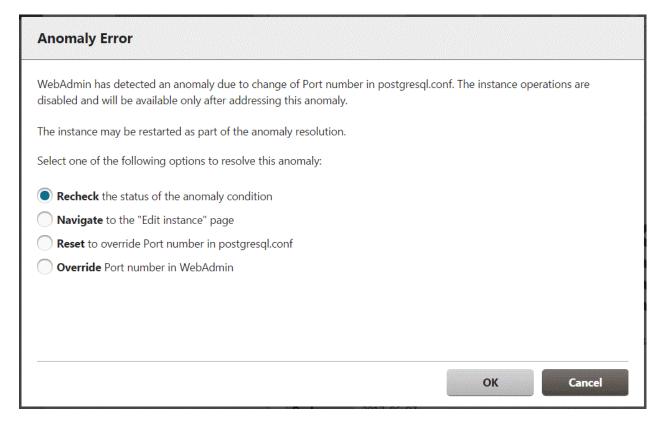
An anomaly occurs when the value of [Port number] and/or [Backup storage path] in WebAdmin is different from the value of its corresponding parameter in postgresql.conf - port and backup\_destination, respectively.

WebAdmin checks for anomalies when an instance is selected for viewing or any instance operation is performed. Anomalies will be identified for the selected instance only.

The following occurs when an anomaly is detected in port number and/or backup storage path:

- All instance operation buttons are disabled, except for "Edit instance", "Refresh instance", and "Delete Mirroring Controller"
- A red error status indicator is displayed on the instance icon
- For an anomaly specific to backup storage path, a red error status indicator is displayed on the [Backup storage] disk icon, and [Backup storage status] is set to "Error"
- The message, "WebAdmin has detected an anomaly with...", is displayed in the [Message] section along with an associated [Solution] button

Click [Solution]. The [Anomaly Error] dialog box is displayed.



Select the required option, click [OK], and then resolve the anomaly error.

Refer to "Editing instance information" in the Installation and Setup Guide for Server for information on the [Edit instance] page.



Critical errors encountered during anomaly resolution will be displayed, however, rollback of the instance to its previous state is not supported.

# 14.15.2 Mirroring Controller Anomalies

The following conditions will cause a Mirroring Controller anomaly:

- The Mirroring Controller management folder or configuration files have been deleted
- The permissions to the Mirroring Controller management folder or configuration files have been changed such that:
  - The instance administrator's access to Mirroring Controller configuration is denied
  - Users other than an instance administrator have access privileges to Mirroring Controller configuration files

WebAdmin checks for anomalies when Mirroring Controller status check is performed.

The following occurs when a Mirroring Controller anomaly is detected:

- All Mirroring Controller functionality is disabled for the replication cluster, except for "Delete Mirroring Controller"
- [Mirroring Controller status] is set to "Error"
- Either of the following messages is displayed in the [Message] section

"Failed to access the Mirroring Controller management folder or configuration files 'path'. Mirroring Controller functionality has been disabled. Consider deleting Mirroring Controller and adding it again."

"Failed to find the Mirroring Controller management folder or configuration files 'path'. Mirroring Controller functionality has been disabled. Consider deleting Mirroring Controller and adding it again."

# Appendix A Parameters

This appendix describes the parameters to be set in the postgresql.conf file of FUJITSU Enterprise Postgres.

The postgresql.conf file is located in the data storage destination.



Note the following when specifying the path:

- Specify \\ as the path delimiter.
- Enclose the path in double quotes (") if it contains spaces.
- core\_directory (string)

This parameter specifies the directory where the corefile is to be output. If this parameter is omitted, the data storage destination is used by default. This parameter can only be set when specified on starting an instance. It cannot be changed dynamically, while an instance is active.

- core\_contents (string)

This parameter specifies the contents to be included in the corefile.

- full: Outputs all contents of the server process memory to the corefile.
- none: Does not output a corefile.
- minimum: Outputs only non-shared memory server processes to the corefile. This reduces the size of the corefile. However, in some cases, this file may not contain sufficient information for examining the factor that caused the corefile to be output.

If this parameter is omitted, "minimum" is used by default. This parameter can only be set when specified on starting an instance. It cannot be changed dynamically, while an instance is active.

- keystore\_location (string)

This parameter specifies the directory that stores the keystore file. Specify a different location from other database clusters. This parameter can only be set when specified on starting an instance. It cannot be changed dynamically, while an instance is active.

- tablespace\_encryption\_algorithm (string)

This parameter specifies the encryption algorithm for tablespaces that will be created. Valid values are "AES128", "AES256", and "none". If you specify "none", encryption is not performed. The default value is "none". To perform encryption, it is recommended that you specify "AES256". Only superusers can change this setting.

- backup\_destination (string)

This parameter specifies the absolute path of the directory where pgx\_dmpall will store the backup data. Specify a different location from other database clusters. This parameter can only be set when specified on starting an instance. It cannot be changed dynamically, while an instance is active.

Place this directory on a different disk from the data directory to be backed up and the tablespace directory. Ensure that users do not store arbitrary files in this directory, because the contents of this directory are managed by the database system.

- search\_path (string)

When using the SUBSTR function compatible with Oracle databases, set "oracle" and "pg\_catalog" in the search\_path parameter. You must specify "oracle" before "pg\_catalog".



search\_path = '"\$user", public, oracle, pg\_catalog'



- The search\_path feature specifies the priority of the schema search path. The SUBSTR function in Oracle database is defined in the oracle schema.
- Refer to "Statement Behavior" under "Server Administration" in the PostgreSQL Documentation for information on search\_path.
- track\_waits (string)

This parameter enables collection of statistics for pgx\_stat\_lwlock and pgx\_stat\_latch.

- on: Enables collection of statistics.
- off: Disables collection of statistics.

If this parameter is omitted, "on" is assumed.

Only superusers can change this setting.

- track\_sql (string)

This parameter enables collection of statistics for pgx\_stat\_sql.

- on: Enables collection of statistics.
- off: Disables collection of statistics.

If this parameter is omitted, "on" is assumed.

Only superusers can change this setting.

#### Parameters for the in-memory feature

- reserve\_buffer\_ratio (numerical value)

This parameter specifies the proportion of shared memory to be used for a stable buffer table.

- Minimum value: 0
- Maximum value: 80

If this parameter is omitted, 0 will be used.

- vci.cost\_threshold (numerical value)

This parameter specifies the lowest cost that selects an execution plan that uses a VCI. If the cost of the best execution plan that does not use a VCI is lower than this value, that execution plan will be selected.

- Minimum value: 0
- Maximum value: Maximum value that can be expressed as a 4-byte signed integer

If this parameter is omitted or a value outside this range is specified, 18000 will be used.

- vci.control\_max\_workers (numerical value)

This parameter specifies the number of background workers that manage VCI. The number of workers for the entire instance is limited by max\_worker\_processes, so add the value specified here to max\_worker\_processes.

- Minimum value: 1
- Maximum value: 8388607

If this parameter is omitted or a value outside this range is specified, 8 will be used.

- vci.enable (string)

This parameter enables or disables VCI.

- on: Enables VCI.
- off: Disables VCI.

If this parameter is omitted, "on" will be used.

- vci.log\_query (string)

This parameter enables or disables log output when VCI is not used due to insufficient memory specified by vci.max\_local\_ros.

- on: Enables log output.
- off: Disables log output.

If this parameter is omitted, "off" will be used.

- vci.maintenance\_work\_mem (numerical value)

This parameter specifies the maximum memory size used for maintenance of VCI (when executing CREATE INDEX, for example).

- Minimum value: 1 MB
- Maximum value: Maximum value that can be expressed as a 4-byte signed integer / 1024

If this parameter is omitted or a value outside this range is specified, 256 MB will be used.

- vci.max\_local\_ros (numerical value)

This parameter specifies the maximum memory size used for VCI scan.

- Minimum value: 64 MB
- Maximum value: Maximum value that can be expressed as a 4-byte signed integer

If this parameter is omitted or a value outside this range is specified, 64 MB will be used.



The maximum value that can be expressed as a 4-byte signed integer changes according to the operating system. Follow the definition of the operating system in use.

vci.max\_parallel\_degree (numerical value)

This parameter specifies the maximum number of background workers used for parallel scan. The number of workers for the entire instance is limited by max\_worker\_processes, so add the value specified here to max\_worker\_processes.

A value from -8388607 to 8388607 can be specified.

- Integer (1 or greater): Parallel scan is performed using the specified degree of parallelism.
- 0: Stops the parallel scan process.
- Negative number: The specified value minus the maximum number of CPUs obtained from the environment is used as the degree of parallelism and parallel scan is performed.

If this parameter is omitted or a value outside this range is specified, 0 will be used.

- vci.shared\_work\_mem (numerical value)

This parameter specifies the maximum memory size used for VCI parallel scan.

- Minimum value: 32 MB
- Maximum value: Maximum value that can be expressed as a 4-byte signed integer

If this parameter is omitted or a value outside this range is specified, 1 GB will be used.

- vci.smc\_directory (string)

This parameter specifies a directory name in which a temporary file is created as the dynamic shared memory during a scan using a VCI.

If this parameter is omitted, a directory (dataStorageDir\base\\pgsql\_tmp) under the data storage directory will be used.

Parameters for the Global Meta Cache

- pgx\_global\_metacache (numerical value)

Specifies the memory size of the GMC area.

Specify a value calculated by the formula below.

A value lower than the calculated value will still work, but the meta cache may not be able to fit into the GMC area.

In this case, the system will discard the meta cache it thinks it is no longer needed, but if it is needed again, the meta cache will need to be expanded and will not perform well.

If the value is less than 10 MB and is set to a nonzero value that disables the feature, the database startup fails because the Global Meta Cache feature cannot operate.

A setting of 0 disables the Global Meta Cache feature. The default is 0.

Changing this setting requires restarting the database.

#### \*1) Safety Factor (1.5)

This value takes into account the case where both GMC before and after the change temporarily exist at the same time in shared memory when the table definition is changed or the row of the system catalog is changed.

- track\_gmc (string)

This parameter enables collection of statistics for pgx\_stat\_gmc.

- on: Enables collection of statistics.
- off: Disables collection of statistics.

If this parameter is omitted, "on" is used.

Only superusers can change this setting.



Refer to "Server Configuration" under "Server Administration" in the PostgreSQL Documentation for information on other postgresql.conf parameters.

# Appendix B System Administration Functions

This appendix describes the system administration functions of FUJITSU Enterprise Postgres.



Refer to "System Administration Functions" under "The SQL Language" in the PostgreSQL Documentation for information on other system administration functions.

# **B.1 WAL Mirroring Control Functions**

The following table lists the functions that can be used for backup and recovery based on WAL mirroring.

Table B.1 WAL mirroring control functions

Name	Return type	Description		
pgx_pause_wal_multiplexing()	void	Stops WAL multiplexing		
pgx_resume_wal_multiplexing()	void	Resumes WAL multiplexing		
pgx_is_wal_multiplexing_paused()	boolean	Returns true if WAL multiplexing has stopped		

If WAL multiplexing has not been configured, these functions return an error. Setting the backup\_destination parameter in postgresql.conf configures WAL multiplexing.

Only superusers can execute these functions.

# **B.2 Transparent Data Encryption Control Functions**

The following table lists the functions that can be used for transparent data encryption.

Table B.2 Transparent data encryption control functions

Name	Return type	Description
pgx_open_keystore(passphrase)	void	Opens the keystore
pgx_set_master_key(passphrase)	void	Sets the master encryption key
pgx_set_keystore_passphrase(oldPassphrase, newPassphrase)	void	Changes the keystore passphrase

The pgx\_open\_keystore function uses the specified passphrase to open the keystore. When the keystore is opened, the master encryption key is loaded into the database server memory. In this way, you can access the encrypted data and create encrypted tablespaces. If the keystore is already open, this function returns an error.

Only superusers can execute this function. Also, this function cannot be executed within a transaction block.

The pgx\_set\_master\_key function generates a master encryption key and stores it in the keystore. If the keystore does not exist, this function creates a keystore. If the keystore already exists, this function modifies the master encryption key. If the keystore has not been opened, this function opens it.

The passphrase is a string of 8 to 200 bytes.

Only superusers can execute this function. Also, this function cannot be executed within a transaction block. Processing is not affected by whether the keystore is open.

The pgx\_set\_keystore\_passphrase function changes the keystore passphrase. Specify the current passphrase in *oldPassphrase*, and a new passphrase in *newPassphrase*.

The passphrase is a string of 8 to 200 bytes.

Only superusers can execute this function. Also, this function cannot be executed within a transaction block. Processing is not affected by whether the keystore is open.

# **B.3 Data Masking Control Functions**

The table below lists the functions that can be used for data masking.

Table B.3 Data masking control functions

Name	Return type	Description		
pgx_alter_confidential_policy	boolean	Changes masking policies		
pgx_create_confidential_policy	boolean	Creates masking policies		
pgx_drop_confidential_policy	boolean	Deletes masking policies		
pgx_enable_confidential_policy	boolean	Enables or disables masking policies		
pgx_update_confidential_values	boolean	Changes replacement characters when full masking is specified for masking type		

### B.3.1 pgx\_alter\_confidential\_policy

#### Description

Changes masking policies

#### **Format**

The format varies depending on the content to be changed. The format is shown below.

- Common format

- Add a masking target to a masking policy

```
pgx_alter_confidential_policy(
commonArg,
[action
             := 'ADD_COLUMN', ]
column_name := 'colName'
[, function_type := 'FULL']
[, function_type := 'PARTIAL', partialOpt] |
[, function_type := 'REGEXP', regexpOpt]
partialOpt:
function_parameters := 'maskingFmt'
regexpOpt:
regexp_pattern
                   := 'regexpPattern',
regexp_replacement := 'regexpReplacementChar',
[, regexp_flags
                    := 'regexpFlags']
```

- Delete a masking target from a masking policy

- Change the masking condition

- Change the content of a masking policy set for a masking target

- Change the masking policy description

```
pgx_alter_confidential_policy(
commonArg,
action := 'SET_POLICY_DESCRIPTION',
policy_description := 'policyDesc'
)
```

- Change the masking target description

```
pgx_alter_confidential_policy(
commonArg,
action := 'SET_COLUMN_DESCRIPTION',
column_name := 'colName',
column_description := 'colDesc'
)
```

#### **Argument**

The argument varies depending on the content to be changed. Details are as follows.

- Common arguments

Masking type for which an argument can be specified	Argument	Data type	Description	Default value
All	schema_name	varchar(63)	Schema name of table for which a masking policy is applied	'public'
	table_name	varchar(63)	Name of table for which a masking policy is applied	Mandatory
	policy_name	varchar(63)	Masking policy name	Mandatory

#### - Add a masking target to a masking policy

Masking type for which an argument can be specified	Argument	Data type	Description	Default value
All	action	varchar(63)	'ADD_COLUMN'	'ADD_COLU MN'
	column_name	varchar(63)	Masking target name	Mandatory
	function_type	varchar(63)	Masking type	'FULL'
			- 'FULL': Full masking	
			- 'PARTIAL': Partial masking	
			- 'REGEXP': Regular expression masking	
Partial masking	function_parameters	varchar(1024)	Masking format for partial masking	Mandatory
Regular expression	regexp_pattern	varchar(1024)	Search pattern for regular expression masking	Mandatory
masking	regexp_replacement	varchar(1024)	Replacement character/string for regular expression masking	Mandatory
	regexp_flags	Regular expression flags	NULL	

#### - Delete a masking target from a masking policy

Masking type for which an argument can be specified	Argument	Data type	Description	Default value
All	action	varchar(63)	'DROP_COLUMN'	Mandatory
	column_name	varchar(63)	Masking target name	Mandatory

#### - Change the masking condition

Masking type for which an argument can be specified	Argument	Data type	Description	Default value
All	action	varchar(63)	'MODIFY_EXPRESSION'	Mandatory
	expression	varchar(1024)	Masking condition to be changed	Mandatory

- Change the content of a masking policy set for a masking target

Masking type for which an argument can be specified	Argument	Data type	Description	Default value
All	action	varchar(63)	'MODIFY_COLUMN'	Mandatory
	column_name	varchar(63)	Masking target name	Mandatory
	function_type	varchar(63)	Masking type	'FULL'
			- 'FULL': Full masking	
			- 'PARTIAL': Partial masking	
			- 'REGEXP': Regular expression masking	
Partial masking	function_parameters	varchar(1024)	Masking format for partial masking	Mandatory
Regular expression	regexp_pattern	varchar(1024)	Search pattern for regular expression masking	Mandatory
masking	regexp_replacement	varchar(1024)	Replacement character/string for regular expression masking	Mandatory
	regexp_flags	varchar(20)	Regular expression flags	NULL

### - Change the masking policy description

Masking type for which an argument can be specified	Argument	Data type	Description	Default value
All	action	varchar(63)	'SET_POLICY_DESCRIPTION'	Mandatory
	policy_description	varchar(1024)	Masking policy description	Mandatory

#### - Change the masking target description

Masking type for which an argument can be specified	Argument	Data type	Description	Default value
All	action	varchar(63)	'SET_COLUMN_DESCRIPTION'	Mandatory
	column_name	varchar(63)	Masking target name	Mandatory
	column_description	varchar(1024)	Masking target description	Mandatory

Details about whether arguments can be omitted are as follows.

Argument		Mandatory or optional									
	ADD_COLUMN			DROP_C	MODIFY	MOE	OIFY_C	OLUM	SET_POLICY	SET_COLUMN	
			OLUMN	_EXPRE		N		_DESCRIPTI	_DESCRIPTIO		
	Full	Partia	Regul		SSION	Full	Partia	Regul	ON	N	
	mas	1	ar			mas	- 1	ar			
	king	maski	expre			king	maski	expre			
		ng	ssion				ng	ssion			
			maski					maski			
			ng					ng			
schema_name	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	
table_name	N	N	N	N	N	N	N	N	N	N	

Argument					Mand	latory	or optic	nal		
	ADD_COLUMN		DROP_C OLUMN	MODIFY _EXPRE	MODIFY_COLUM N			SET_POLICY _DESCRIPTI	SET_COLUMN _DESCRIPTIO	
	Full mas king	- 1	Regul ar expre ssion maski ng		SSION	Full mas king	Partia I maski ng	Regul ar expre ssion maski ng	ON	N
policy_name	N	N	N	N	N	N	N	N	N	N
action	Y	Y	Y	N	N	N	N	N	N	N
column_name	N	N	N	N	-	N	N	N	-	N
function_type	Y	N	N	-	-	Y	N	N	-	-
expression	-	-	-	-	N	-	-	-	-	-
policy_description	-	-	-	-	-	-	-	-	N	-
column_description	-	-	-	-	-	-	-	-	-	N
function_parameters	-	N	-	-	-	-	N	-	-	-
regexp_pattern	-	-	N	-	-	-	-	N	-	-
regexp_replacement	-	-	N	-	-	-	-	N	-	-
regexp_flags	-	-	Y	-	-	-	-	Y	-	-

Y: Can be omitted; N: Cannot be omitted; -: Ignored when specified

#### Return value

Return value	Description
TRUE	Ended normally
FALSE	Ended abnormally

#### **Execution example 1**

Adding masking policy p1 to masking target c2

#### **Execution example 2**

Deleting masking target c1 from masking policy p1

```
postgres=# select pgx_alter_confidential_policy(table_name := 't1', policy_name := 'p1', action :=
'DROP_COLUMN', column_name := 'c1');
pgx_alter_confidential_policy
-----t
t
(1 row)
```

#### **Execution example 3**

Changing the masking condition for masking policy p1

```
postgres=# select pgx_alter_confidential_policy(table_name := 't1', policy_name := 'p1', action :=
'MODIFY_EXPRESSION', expression := 'false');
pgx_alter_confidential_policy
-----t
(1 row)
```

#### **Execution example 4**

Changing the content of masking policy p1 set for masking target c2

```
postgres=# select pgx_alter_confidential_policy(table_name := 't1', policy_name := 'p1', action :=
'MODIFY_COLUMN', column_name := 'c2', function_type := 'FULL');
pgx_alter_confidential_policy
-----t
t
(1 row)
```

#### **Execution example 5**

Changing the description of masking policy p1

#### **Execution example 6**

Changing the description of masking target c2

```
postgres=# select pgx_alter_confidential_policy(table_name := 't1', policy_name := 'p1', action :=
'SET_COLUMN_DESCRIPTION', column_name := 'c2', column_description := 'c2 column is FULL.');
pgx_alter_confidential_policy
------t
(1 row)
```

#### Description

- The arguments for the pgx\_alter\_confidential\_policy system management function can be specified in any order.
- The action parameters below can be specified. When action parameters are omitted, ADD\_COLUMN is applied.

Parameter	Description
ADD_COLUMN	Adds a masking target to a masking policy.
DROP_COLUMN	Deletes a masking target from a masking policy.
MODIFY_EXPRESSION	Changes expression.
MODIFY_COLUMN	Changes the content of a masking policy set for a masking target.
SET_POLICY_DESCRIPTION	Changes policy_description.
SET_COLUMN_DESCRIPTION	Changes column_description.

- The function\_parameters argument is enabled when the function\_type is PARTIAL. If the function\_type is other than PARTIAL, it will be ignored.

- The arguments below are enabled when the function\_type is REGEXP. If the function\_type is other than REGEXP, these arguments will be ignored.
  - regexp\_pattern
  - regexp\_replacement
  - regexp\_flags



#### See

- Refer to "String Constants" in the PostgreSQL Documentation for information on the strings to specify for arguments.
- Refer to "POSIX Regular Expressions" in the PostgreSQL Documentation and check pattern, replacement, and flags for information on the values that can be specified for regexp\_pattern, regexp\_replacement, and regexp\_flags.

......

### B.3.2 pgx\_create\_confidential\_policy

#### Description

Creates masking policies

#### **Format**

The format varies depending on the masking type. The format is shown below.

```
pgx_create_confidential_policy(
[schema_name := 'schemaName',]
               := 'tableName',
table_name
policy_name
                 := 'policyName',
              := 'expr'
expression
[, enable
                 := 'policyStatus']
[, policy_description
                      := 'policyDesc']
[, column_name := 'colName'
                     := 'FULL']
    [, function_type
    [, function_type := 'PARTIAL', partialOpt] |
    [, function_type := 'REGEXP', regexpOpt]
                          := 'colDesc']
    [, column_description
])
partialOpt:
function_parameters := 'maskingFmt'
regexpOpt:
              := 'regexpPattern',
regexp_pattern
                 := 'regexpFlags']
regexp_replacement
                       := 'regexpReplacementChar',
[, regexp_flags
```

#### **Argument**

Details are as follows.

Masking type for which an argument can be specified	Argument	Data type	Description	Default value
All	schema_name	varchar(63)	Schema name of table for which the masking policy is created	'public'
	table_name	varchar(63)	Name of table for which the masking policy is created	Mandatory

Masking type for which an argument can be specified	Argument	Data type	Description	Default value
	policy_name	varchar(63)	Masking policy name	Mandatory
	expression	varchar(1024)	Masking condition	Mandatory
	enable	boolean	Masking policy status	't'
			- 't': Enabled	
			- 'f': Disabled	
	policy_description	varchar(1024)	Masking policy description	NULL
	column_name	varchar(63)	Masking target name	NULL
	function_type	varchar(63)	Masking type	'FULL'
			- 'FULL': Full masking	
			- 'PARTIAL': Partial masking	
			- 'REGEXP': Regular expression masking	
	column_description	varchar(1024)	Masking target description	NULL
Partial masking	function_parameters	varchar(1024)	Masking format for partial masking	Mandatory
Regular expression masking	regexp_pattern	varchar(1024)	Search pattern for regular expression masking	Mandatory
	regexp_replacement	varchar(1024)	Replacement character/string for regular expression masking	Mandatory
	regexp_flags	varchar(20)	Regular expression flags	NULL

Details about whether arguments can be omitted are as follows.

Argument	Mandatory or optional		
	Full masking	Partial masking	Regular expression masking
schema_name	Y	Y	Y
table_name	N	N	N
policy_name	N	N	N
expression	N	N	N
enable	Y	Y	Y
policy_description	Y	Y	Y
column_name	Y	Y	Y
function_type	Y	Y	Y
column_description	Y	Y	Y
function_parameters	-	N	-
regexp_pattern	-	-	N
regexp_replacement	-	-	N
regexp_flags	-	-	Y

Y: Can be omitted; N: Cannot be omitted; -: Ignored when specified

#### Return value

Return value	Description
TRUE	Ended normally
FALSE	Ended abnormally

#### **Execution example 1**

Creating masking policy p1 that does not contain a masking target

#### **Execution example 2**

Creating masking policy p1 that contains masking target c1 of which the masking type is full masking

```
postgres=# select pgx_create_confidential_policy(schema_name := 'public', table_name := 't1',
policy_name := 'p1', expression := '1=1', enable := 't', policy_description := 'this policy is an
example.', column_name := 'c1', function_type := 'FULL', column_description := 'c1 column is FULL.');
pgx_create_confidential_policy
------t
(1 row)
```

#### **Execution example 3**

Creating masking policy p1 that contains masking target c2 of which the masking type is partial masking

#### **Execution example 4**

Creating masking policy p1 that contains masking target c3 of which the masking type is regular expression masking

```
postgres=# select pgx_create_confidential_policy( table_name := 't1', policy_name := 'p1',
expression := '1=1', column_name := 'c3', function_type := 'REGEXP', regexp_pattern := '(.*)(@.*)',
regexp_replacement := 'xxx\2', regexp_flags := 'g');
pgx_create_confidential_policy
------t
(1 row)
```

#### Description

- The arguments for the pgx\_create\_confidential\_policy system management function can be specified in any order.
- If column name is omitted, only masking policies that do not contain masking target will be created.
- One masking policy can be created for each table. Use the pgx\_alter\_confidential\_policy system management function to add a masking target to a masking policy.

- The function\_parameters argument is enabled when the function\_type is PARTIAL. If the function\_type is other than PARTIAL, it will be ignored.
- The arguments below are enabled when the function\_type is REGEXP. If the function\_type is other than REGEXP, these arguments will be ignored.

......

- regexp\_pattern
- regexp\_replacement
- regexp\_flags



If a table for which a masking policy is to be applied is deleted, delete the masking policy as well.



#### See

- Refer to "String Constants" in the PostgreSQL Documentation for information on the strings to specify for arguments.
- Refer to "POSIX Regular Expressions" in the PostgreSQL Documentation and check pattern, replacement, and flags for information on the values that can be specified for regexp\_pattern, regexp\_replacement, and regexp\_flags.

### B.3.3 pgx\_drop\_confidential\_policy

#### Description

Deletes masking policies

#### **Format**

```
pgx_drop_confidential_policy(
[schema_name := 'schemaName', ]
table_name := 'tableName',
policy_name := 'policyName'
)
```

#### **Argument**

Details are as follows.

Argument	Data type	Description	Default value
schema_name	varchar(63)	Schema name of table for which a masking policy is deleted	'public'
table_name	varchar(63)	Name of table for which a masking policy is deleted	Mandatory
policy_name	varchar(63)	Masking policy name	Mandatory

Details about whether arguments can be omitted are as follows.

Argument	Mandatory or optional
schema_name	Y
table_name	N
policy_name	N

Y: Can be omitted; N: Cannot be omitted

#### Return value

Return value	Description
TRUE	Ended normally
FALSE	Ended abnormally

#### **Execution example**

Deleting masking policy p1

```
postgres=# select pgx_drop_confidential_policy(table_name := 't1', policy_name := 'p1');
    pgx_drop_confidential_policy
    t
    (1 row)
```

#### Description

The arguments for the pgx\_drop\_confidential\_policy system management function can be specified in any order.



If a table for which a masking policy is to be applied is deleted, delete the masking policy as well.



Refer to "String Constants" in the PostgreSQL Documentation for information on the strings to specify for arguments.

### B.3.4 pgx\_enable\_confidential\_policy

#### Description

Enables or disables masking policies

#### **Format**

#### Argument

Details are as follows.

Argument	Data type	Description	Default value
schema_name	varchar(63)	Schema name of table for which a masking policy is enabled or disabled	'public'
table_name	varchar(63)	Name of table for which a masking policy is enabled or disabled	Mandatory
policy_name	varchar(63)	Masking policy name	Mandatory
enable	boolean	Masking policy status	Mandatory
		- 't': Enabled	

Argument	Data type	Description	Default value
		- 'f': Disabled	

Details about whether arguments can be omitted are as follows.

Argument	Mandatory or optional
schema_name	Y
table_name	N
policy_name	N
enable	N

Y: Can be omitted; N: Cannot be omitted

#### Return value

Return value	Description
TRUE	Ended normally
FALSE	Ended abnormally

#### **Execution example**

Enabling masking policy p1

```
postgres=# select pgx_enable_confidential_policy(table_name := 't1', policy_name := 'p1', enable :=
't');
pgx_enable_confidential_policy
-----t
(1 row)
```

......

#### Description

 $The \ arguments \ for \ the \ pgx\_enable\_confidential\_policy \ system \ management \ function \ can \ be \ specified \ in \ any \ order.$ 



#### See

Refer to "String Constants" in the PostgreSQL Documentation for information on the strings to specify for arguments.

# B.3.5 pgx\_update\_confidential\_values

#### **Description**

Changes replacement characters when full masking is specified for masking type

#### **Format**

```
pgx_update_confidential_values(
[number_value := 'numberValue']
[, char_value := 'charValue']
[, varchar_value := 'varcharValue']
[, date_value := 'dateValue']
[, ts_value := 'tsValue']
)
```

#### **Argument**

Details are as follows.

Argument	Data type	Description
number_value	integer	Replacement character in numeric type
char_value	varchar(1)	Replacement character in char type
varchar_value	varchar(1)	Replacement character in varchar type
date_value	date	Replacement character in date type
ts_value	timestamp	Replacement character in timestamp type

#### Return value

Return value	Description
TRUE	Ended normally
FALSE	Ended abnormally

#### **Execution example**

Using '\*' as a replacement character in char type and varchar type

```
postgres=# select pgx_update_confidential_values(char_value := '*', varchar_value := '*');
  pgx_update_confidential_values
-----t
  t
  (1 row)
```

#### Description

- The arguments for the pgx\_update\_confidential\_values system management function can be specified in any order.
- Specify one or more arguments for the pgx\_update\_confidential\_values system management function. A replacement character is not changed for an omitted argument.

.....



Refer to "String Constants" in the PostgreSQL Documentation for information on the strings to specify for arguments.

#### **B.4 VCI Data Load Control Function**

The table below lists the function that loads VCI data to buffer cache.

#### Table B.4 VCI data load control function

Name	Return type	Description
pgx_prewarm_vci(vci_index regclass)	int8	Loads the VCI data to buffer cache.

pgx\_prewarm\_vci loads the specified VCI data to buffer cache and returns the number of blocks of the loaded VCI data.

The aggregation process using VCI may take time immediately after an instance is started, because the VCI data has not been loaded to buffer cache. Therefore, the first aggregation process can be sped up by executing pgx\_prewarm\_vci after an instance is started.

The amount of memory required for preloading is the number of blocks returned by pgx\_prewarm\_vci multiplied by the size of one block.

This function can only be executed if the user has reference privilege to the VCI index and execution privilege to the pg\_prewarm function.

# **B.5 High-Speed Data Load Control Functions**

The table below lists the functions that can be used for high-speed data load.

Table B.5 High-speed data load control functions

Name	Return type	Description
pgx_loader	bigint	Creates dynamic shared memory, starts parallel workers and loads data
pgx_loader_recovery	smallint	Resolves in-doubt transactions

The pgx\_loader command executes the above functions internally.

# Appendix C System Views

This appendix describes how to use the system views in FUJITSU Enterprise Postgres.



Refer to "System Views" under "Internals" in the PostgreSQL Documentation for information on other system views.

### C.1 pgx\_tablespaces

The pgx\_tablespaces catalog provides information related to the encryption of tablespaces.

Column	Туре	References	Description
spctablespace	oid	pg_tablespace.oid	Tablespace OID
spcencalgo	text		Tablespace encryption algorithm

The spcencalgo string displays one of the following values:

- none: Tablespace is not encrypted
- AES128: AES with key length of 128 bits
- AES256: AES with key length of 256 bits

### C.2 pgx\_stat\_lwlock

The pgx\_stat\_lwlock view displays statistics related to lightweight locks, with each type of content displayed on a separate line.

Table C.1 pgx\_stat\_lwlock view

Column	Туре	Description
lwlock_name	name	Name of the lightweight lock
total_waits	bigint	Number of waits caused by the lightweight lock
total_wait_time	double precision	Number of milliseconds spent in waits caused by the lightweight lock
stats_reset	timestamp with timezone	Last time at which this statistics was reset

# C.3 pgx\_stat\_latch

The pgx\_stat\_latch view displays statistics related to latches, with each type of wait information within FUJITSU Enterprise Postgres displayed on a separate line.

Table C.2 pgx stat latch view

Table 0.2 pgx_stat_later	VICVV	
Column	Туре	Description
latch_name	name	Name of the latch
total_waits	bigint	Number of waits caused a wait
total_wait_time	double precision	Number of milliseconds spent in waits caused by the latch
stats_reset	timestamp with timezone	Last time at which this statistic was reset

# C.4 pgx\_stat\_walwriter

The pgx\_stat\_walwriter view displays statistics related to WAL writing, in a single line.

Table C.3 pgx\_stat\_walwriter view

Column	Туре	Description
dirty_writes	bigint	Number of times old WAL buffers were written to the disk because the WAL buffer was full when WAL records were added
writes	bigint	Number of WAL writes
write_blocks	bigint	Number of WAL write blocks
total_write_time	double precision	Number of milliseconds spent on WAL writing
stats_reset	timestamp with timezone	Last time at which this statistic was reset

# C.5 pgx\_stat\_sql

 $The \ pgx\_stat\_sql\ view\ displays\ statistics\ related\ to\ SQL\ statement\ executions, with\ each\ type\ of\ SQL\ statement\ displayed\ on\ a\ separate\ line.$ 

Table C.4 pgx stat sql view

Column	Туре	Description
selects	bigint	Number of SELECT statements executed
		In database multiplexing mode, this number includes the SELECT statements executed in Mirroring Controller. Mirroring Controller executes the SELECT statement using the interval specified for the heartbeat_interval of the server definition file (milliseconds).
inserts	bigint	Number of INSERT statements executed
deletes	bigint	Number of DELETE statements executed
updates	bigint	Number of UPDATE statements executed
selects_with_parallelism	bigint	Number of times parallel scan was used in SELECT statements
inserts_with_parallelism	bigint	Not used
deletes_with_parallelism	bigint	Not used
updates_with_parallelism	bigint	Not used
copies_with_parallelism	bigint	Not used
declares	bigint	Number of DECLARE statements executed (number of cursor OPENs)
fetches	bigint	Number of FETCH statements executed
checkpoints	bigint	Number of CHECKPOINT statements executed
clusters	bigint	Number of CLUSTER statements executed
copies	bigint	Number of COPY statements executed
reindexes	bigint	Number of REINDEX statements executed
truncates	bigint	Number of TRUNCATE statements executed
locks	bigint	Number of times a lock occurred
stats_reset	timestamp with timezone	Last time at which this statistic was reset

# C.6 pgx\_stat\_gmc

The pgx\_stat\_gmc view provides information about the GMC areas.

Table C.5 pgx\_stat\_gmc view

Column	Туре	Description
searches	bigint	Number of times the cache table is searched.
hits	bigint	Number of times the cache table is hit.
size	bigint	The current amount of memory (bytes) used in the GMC area.
stats_reset	timestamp with timezone	Last time these statistics were reset.

# Appendix D Tables Used by Data Masking

This appendix explains tables used by the data masking feature.



These tables are updated by the data masking control function, so do not use SQL statements to directly update these tables.

# D.1 pgx\_confidential\_columns

This table provides information on masking target for which masking policies are set.

Column	Туре	Description
schema_name	varchar(63)	Schema name of table for which a masking policy is applied
table_name	varchar(63)	Name of table for which a masking policy is applied
policy_name	varchar(63)	Masking policy name
column_name	varchar(63)	Masking target name
function_type	varchar(63)	Masking type
		- 'FULL': Full masking
		- 'PARTIAL': Partial masking
		- 'REGEXP': Regular expression masking
function_parameters	varchar(1024)	Masking format for partial masking
regexp_pattern	varchar(1024)	Search pattern for regular expression masking
regexp_replacement	varchar(1024)	Replacement character/string for regular expression masking
regexp_flags	varchar(20)	Regular expression flags
column_description	varchar(1024)	Masking target description

#### **Execution example**

# D.2 pgx\_confidential\_policies

This table provides information on masking policies.

Column	Туре	Description
schema_name	varchar(63)	Schema name of table for which a masking policy is applied

Column	Type	Description
table_name	varchar(63)	Name of table for which a masking policy is applied
policy_name	varchar(63)	Masking policy name
expression	varchar(1024)	Masking condition
enable	boolean	Masking policy status
		- 't': Enabled
		- 'f': Disabled
policy_description	varchar(1024)	Masking policy description

#### **Execution example**

# D.3 pgx\_confidential\_values

This table provides information on replacement characters when full masking is specified for masking type.

Column	Data type	Description	Default value
number_value	integer	Numeric	0
char_value	varchar(1)	char type	Spaces
varchar_value	varchar(1)	varchar type	Spaces
date_value	date	date type	'1970-01-01'
timestamp_value	timestamp	timestamp type	'1970-01-01 00:00:00'

#### **Execution example**

# Appendix E Tables Used by High-Speed Data Load

This appendix describes the tables used by high-speed data load.

# E.1 pgx\_loader\_state

 $The \ pgx\_loader\_state \ table \ provides \ information \ about \ transactions \ prepared \ by \ high-speed \ data \ load.$ 

Column	Туре	Description
id	serial	Unique identifier.
		This value is assigned from the pgx_loader_state_id_seq sequence.
gid	text	Global transaction identifier assigned to a transaction.
state	text	State of the transaction.
		The value can be one of the following:
		- commit: The prepared transaction has been committed.
		- rollback: The prepared transaction is in in-doubt state.
master_pid	integer	Process ID of the backend process (master process) that executed the pgx_loader control function.
role_oid	integer	Role identifier (OID). A prepared transaction can only be completed by the same user who executed the original transaction or by a superuser.
relation_oid	integer	Object identifier (OID).



The pgx\_loader\_state table and pgx\_loader\_state\_id\_seq sequence are updated by high-speed data load. Do not update these database objects directly using SQL.

# Appendix F Starting and Stopping the Web Server Feature of WebAdmin

To use WebAdmin for creating and managing a FUJITSU Enterprise Postgres instance on a server where FUJITSU Enterprise Postgres is installed, you must first start the Web server feature of WebAdmin.

- Using WebAdmin in a single-server configuration

You must start the Web server on the server on which FUJITSU Enterprise Postgres and WebAdmin are installed..

- Using WebAdmin in a multiserver configuration

You must start the Web server on all servers on which WebAdmin has been installed.

This appendix describes how to start and stop the Web server feature of WebAdmin.



Refer to "Installing WebAdmin in a Multiserver Configuration" in the Installation and Setup Guide for Server for information on multiserver installation.

### F.1 Starting the Web Server Feature of WebAdmin

Follow the procedure below to start the Web server feature of WebAdmin:

- 1. Display the [Services] window
  - Windows Server(R) 2012 or Windows Server(R) 2012 R2:

In the [Start] screen, select [Administrative Tools], and then click [Services].

- All other operating systems:

In the [Start] menu, select [Administrative Tools], and then click [Services].

2. Start a service

Select the displayed name "FUJITSU Enterprise Postgres WebAdmin version", and then click [Start Service].

You can also start a service by specifying the service name of the Web server feature of WebAdmin in the net start command or sc start command.

# F.2 Stopping the Web Server Feature of WebAdmin

This section describes how to stop the Web server feature of WebAdmin.

Follow the procedure below to stop the Web server feature of WebAdmin:

- 1. Display the [Services] window
  - Windows Server(R) 2012 or Windows Server(R) 2012 R2:

In the [Start] screen, select [Administrative Tools], and then click [Services].

- All other operating systems:

In the [Start] menu, select [Administrative Tools], and then click [Services].

2. Stop a service

Select the displayed name "FUJITSU Enterprise Postgres WebAdmin version", and then click [Stop Service].

You can also stop a service by specifying the service name of the Web server feature of WebAdmin in the net stop command or sc stocommand.	р

# Appendix G WebAdmin Wallet

This appendix describes how to use the Wallet feature of WebAdmin.

When a remote instance or a standby instance is created, it is necessary to provide user name and password for authentication with the remote machine or the database instance.

The Wallet feature in WebAdmin is a convenient way to create and store these credentials.

Once created, these credentials can be repeatedly used in one or more instances.

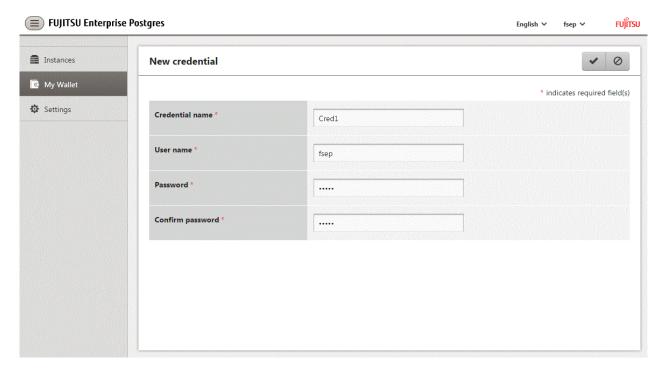


It is not mandatory to create a credential in the Wallet. It is possible to create a remote instance or a standby instance without creating any credential in the Wallet.

If no credential is created beforehand, a user name and password can be entered in the instance creation page. When creating a "Remote" instance, if operating system credentials are entered without using a credential stored in the Wallet, WebAdmin automatically creates a credential with the given user name and password, and stores it in the user's wallet for future use.

# G.1 Creating a Credential

- 1. In the [My Wallet] tab, click The [New credential] page will be displayed.
- 2. Enter the information for the credentials.



Enter the following items. Credential name, User name and Password should not contain hazardous characters. Refer to "Appendix H WebAdmin Disallow User Inputs Containing Hazardous Characters".

- [Credential name]: Name of the credential

The name must meet the conditions below:

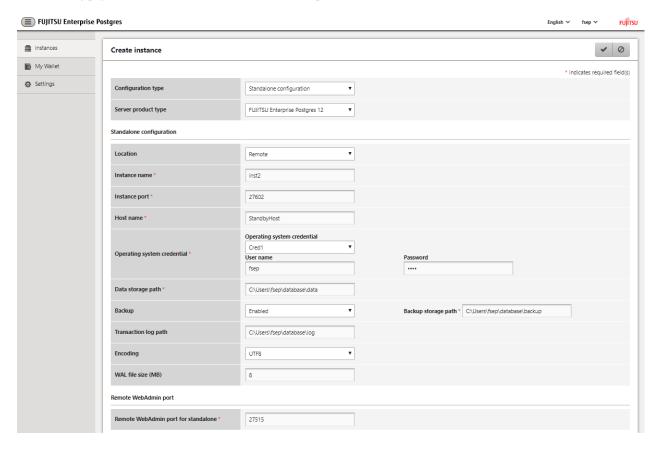
- Maximum of 16 characters
- The first character must be an ASCII alphabetic character
- The other characters must be ASCII alphanumeric characters

- [User name]: The operating system user name or database instance user name that will be used later
- [Password]: Password for the user
- [Confirm password]: Reenter the password.
- 3. Click **v** to store the credential.

# **G.2** Using a Credential

Once a credential is created in the Wallet, it can be used during remote instance creation or standby instance creation.

The following page uses the credential that was created in the previous section.



When "Cred1" is selected in [Operating system credential], the user name and password are automatically populated from the credential.

# Appendix H WebAdmin Disallow User Inputs Containing Hazardous Characters

WebAdmin considers the following as hazardous characters, which are not allowed in user inputs.

| (pipe sign)
& (ampersand sign)
; (semicolon sign)
\$ (dollar sign)
% (percent sign)
@ (at sign)
' (single apostrophe)
" (quotation mark)
\' (backslash-escaped apostrophe)
\" (backslash-escaped quotation mark)
<> (triangular parenthesis)
() (parenthesis)
+ (plus sign)
CR (Carriage return, ASCII 0x0d)

LF (Line feed, ASCII 0x0a)

, (comma sign)
\ (backslash)

# Appendix I Collecting Failure Investigation Data

If the cause of an error that occurs while building the environment or during operations is unclear, data must be collected for initial investigation.

This appendix describes how to collect data for initial investigation.

Use FJQSS (Information Collection Tool) to collect data for initial investigation.



#### See

Refer to the following manual for information on how to use FJQSS.

- In the [Apps] menu, select [FJQSS (Information Collection Tool)], and then click [FJQSS User's Guide].



- When using FJQSS to collect data for initial investigation, a window will be displayed for you to set the following environment variables:

.....

- PGDATA

Set the data storage destination.

- PGDATABASE

Set the database name from which you want to collect data for initial investigation.

PGPORT

Set the instance port number. This does not need to be set if the default port number (27500) has not been changed.

- PGUSER

Set the database superuser.

Set the database superuser so that client authentication is possible.

FJQSS establishes a TCP/IP connection with the template1 database and collects data from the database.

- Refer to "Collecting Failure Investigation Data" in the Cluster Operation Guide (Database Multiplexing) for information on how to collect failure investigation data when performing database multiplexing.

# <u>Index</u>

[A]	Encrypting Existing Data	36
Actions in Response to Instance Startup Failure116	Encryption mechanisms	
All user data within the specified tablespace27	Errors in More Than One Storage Disk	
Approximate backup time	· ·	
Approximate recovery time90	[F]	
Automatically opening the keystore38	Faster encryption and decryption based on hardware	
	File system level backup and restore	35
[B]	ri n	
Backing Up and Recovering the Keystore32	[H]	72
Backing Up and Restoring/Recovering the Database34	High-Speed Data Load	13
Backup/Recovery Using the Copy Command81	[1]	
Backup and recovery using the pgx_dmpall and pgx_rcvall	If failure occurred in the data storage disk or the transaction	on log
commands34	storage disk	
backup cycle	If failure occurred on the backup data storage disk	
Backup data	If failure occurred on the data storage disk or the transacti	
Backup operation20	storage directory	
Backup operation (file backup)21	Importing and Exporting the Database	
Backup status20,22	Installing and Operating the In-memory Feature	
Backup using the backup information file82		
Backup Using the Copy Command84	[K]	
backup_destination (string)	keystore_location (string)	123
Building and starting a standby server		
[C]	[L]	
Changing a Masking Policy46	Logging in to WebAdmin	
Changing the Keystore Passphrase	log in	4
Changing the Master Encryption Key	[M]	
Changing the master encryption key and the passphrase 38	Managing the Keystore	31
Checking an Encrypted Tablespace30	Masking Condition	
Checking backup status	Masking Format	
Checking the operating state of an instance	Masking Policy	
Checking the operating status of an instance	Masking Target	
Collecting Failure Investigation Data	Masking Type	
Configuration of the Copy Command	Monitoring Database Activity	
Configuration of the copy command for backup83		
Configuration of the copy command for recovery	[O]	
Confirming a Masking Policy46	Opening the Keystore	28
Continuous archiving and point-in-time recovery35	Operating FUJITSU Enterprise Postgres	1
Copy Command for Backup86	<b>(D)</b>	
Copy Command for Recovery88	[P]	==
Copy Command Interface86	Parallel Query	
core_contents (string)	Performing backup	
core_directory (string)123	Perform recovery	
Creating a Masking Policy45	Periodic Backup	
Cyclic usage of the backup area81	pgx_global_metacache (numerical value)	
	pgx_stat_gmc view	
[D]	pgx_stat_latch view	
Data Masking40	pgx_stat_lwlock view	
Data Types for Masking48	pgx_stat_sql view	
Deleting a Masking Policy48	pgx_stat_walwriter view	
Determining the backup area of the latest backup85	pgx_tablespaces	
r=1	Placement and automatic opening of the keystore file	
[E]	Placing the keystore file	
Enabling and Disabling a Masking Policy	Preparing for backup	84
Enabling Automatic Opening of the Keystore	[R]	
Encrypting a Tablespace	Recovery Using the Copy Command	85

reserve_buffer_ratio (numerical value)	124
[S]	
Scope of encryption	27
search_path (string)	123
Security-Related Notes	
Security Notes	49
Setting a restore point	
Setting the Master Encryption Key	28
Starting and Stopping the Web Server Feature of WebA	
Starting an instance	13,15
Starting pgAdmin	
Startup URL for WebAdmin	
Stopping an instance	
Streaming replication support	
Streaming Replication Using WebAdmin	
Strong encryption algorithms	
System Administration Functions	
System Views	
[T]	
tablespace_encryption_algorithm (string)	123
Tables Used by Data Masking	145
Tips for Installing Built Applications	39
track_gmc (string)	126
track_sql (string)	124
track_waits (string)	124
Transparent Data Encryption Control Functions	127
Two-layer encryption key and the keystore	27
[U]	
User environment	3
[V]	
vci.control_max_workers (numerical value)	124
vci.cost_threshold (numeric)	
vci.enable (string)	
vci.log_query (string)	
vci.maintenance_work_mem (numerical value)	
vci.max_local_ros (numerical value)	
vci.max_parallel_degree (numerical value)	
vci.shared_work_mem (numerical value)	
ver.shared_work_mem (numerical value)	123
[W]	
WAL and temporary files	27
WAL Mirroring Control Functions	127
WebAdmin Wallet	



# FUJITSU Enterprise Postgres 12



# Security Operation Guide

# **Preface**

#### Purpose of this document

This document describes security when building and operating a FUJITSU Software Enterprise Postgres (hereinafter referred to as "FUJITSU Enterprise Postgres") database system.

#### Intended readers

This document is intended for those who are:

- Considering installing FUJITSU Enterprise Postgres
- Designing, building, and operating the security operating environment in FUJITSU Enterprise Postgres
- Accessing FUJITSU Enterprise Postgres database systems Readers of this document are assumed to have general knowledge of:
  - Business operations
  - FUJITSU Enterprise Postgres
  - Linux

#### Structure of this document

This document is structured as follows:

#### Chapter 1 Overview of Security

Provides an overview of the security system, and explains the security features provided by FUJITSU Enterprise Postgres.

#### Chapter 2 Overview of Security Operation

Provides an overview of security operation.

#### Chapter 3 Tasks of the Manager

Explains the tasks for security measures to be implemented by the manager.

#### Chapter 4 Tasks of Administrators

Explains the tasks for security measures to be implemented by administrators.

#### Chapter 5 Tasks of Users

Explains the tasks for security measures to be implemented by users.

#### Chapter 6 Audit Log Feature

Explains the audit log feature provided by FUJITSU Enterprise Postgres.

#### References

This document contains abstracts from the following document:

 Database Security Guideline Version 2.0 (Database Security Consortium (DBSC))

#### **Export restrictions**

Exportation/release of this document may require necessary procedures in accordance with the regulations of your resident country and/or US export control laws.

### Issue date and version

Edition 1.0: July 2020

### Copyright

Copyright 2018-2020 FUJITSU LIMITED

# Contents

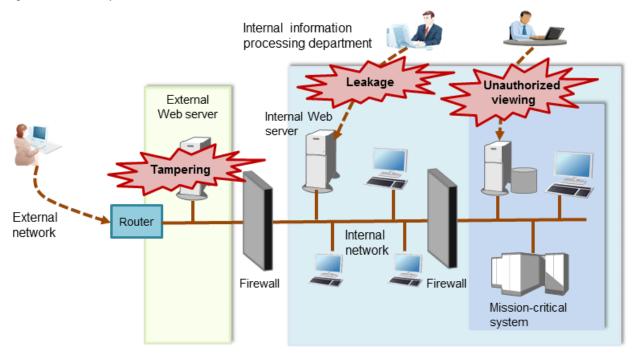
Chapter 1 Overview of Security	1
1.1 What is Security?	1
1.2 Security Requirements	
1.3 Security Threats	
1.4 Security Scope	
1.5 Security Provided by FUJITSU Enterprise Postgres	e
1.5.1 Roles Targeted For Security	(
1.5.2 Security Features.	
Chapter 2 Overview of Security Operation	
2.1 Security Operation Flow	9
Chapter 3 Tasks of the Manager	1.
3.1 Defining Important Information and Risk Analysis.	
3.2 Formulating Account Management Policies.	
3.3 Formulating Log Retrieval Policies.	
3.4 Formulating Rules	
3.5 Implementing Training	
3.6 Checking the Database Management Operations	
5.7 I chould Diagnosis of the Status of Security Measures	13
Chapter 4 Tasks of Administrators	14
4.1 Receiving Training	
4.2 Initial Setup.	
4.3 Authentication.	
4.3.1 Managing Accounts.	
4.3.2 Managing Passwords	
4.3.3 Configuring Connections and Authentication.	
4.4 Access Control	
4.5 Encryption.	
4.6 Controlling Use of External Media.	
4.7 Security Measures for Servers/Applications.	
4.8 Log Management.	
4.8.1 Retrieving Logs.	
4.8.2 Maintaining Logs.	
4.9 Detecting Unauthorized Access.	
4.10 Analyzing Logs.	
4.10 Analyzing Logs	,20
Chapter 5 Tasks of Users	21
5.1 Receiving Training	21
5.2 Managing Accounts/Passwords	21
Chapter 6 Audit Lea Feature	20
Chapter 6 Audit Log Feature.	
6.1 Audit Log Output Modes.	
6.2 Setup	
6.3 pgaudit Configuration File	
6.4 Session Audit Logging	
6.5 Object Audit Logging	
6.6 Database Multiplexing	
6.6.1 Setup	
6.6.2 Configuring Audit Log Retrieval	
6.7 Failover	
6.7.1 Configuring Audit Log Retrieval	
6.8 View Audit Logs Using SQL	
6.9 Removing Setup.	39

# Chapter 1 Overview of Security

# 1.1 What is Security?

Computer security is the protection of information systems and data from risks such as leakage or tampering of information, attacks, intrusions, eavesdropping from external sources, and interference with information services. Security measures are essential for the advance prevention of security threats in order for information systems to gain trust as social infrastructure.

Figure 1.1 Security threats



The security measures in information systems can be classified as follows:

- Network
- Web
- Application
- Database
- PC

This document focuses on database security measures when using FUJITSU Enterprise Postgres.

# 1.2 Security Requirements

Below are the necessary security requirements for information systems.

#### Maintenance of security policies

A security policy clarifies the approach the company should take in relation to information assets, and the actions employees should take.

It is necessary to undertake security of information systems while maintaining security policies.

#### Integrated security management

Security has the aspects below. It is necessary to manage information in an integrated manner based on these aspects.

#### Confidentiality

Access to the information is restricted to prevent leakage of information outside of the company

Example measures: Prevention of information leakage or setup of access privileges

#### Integrity

Integrity is guaranteed, ensuring information does not become corrupted or tampered with

Example measures: Prevention or detection of tampering

#### Availability

Failure is prevented and normal operation is maintained so that information can be used when needed

Example measures: Power supply measures, system mirroring

# 1.3 Security Threats

A security threat is defined as something that threatens the confidentiality, integrity, and availability indicated in "1.2 Security Requirements" in respect to information assets. This includes technical threats such as accessing a database, but does not include physical destruction.

Threats are considered to be a combination of type of user who is the source of the threat, information assets that need to be protected, techniques, and unauthorized actions. For example, a threat might be a general user exploiting a database vulnerability to obtain database management information, and then tampering with that information.

When considering security measures, it is firstly necessary to clarify what kind of threats there are. A list of possible threats is shown in the table below. Refer to "Types of user" and "Information assets" for details on the definition of each type of user and information assets that should be protected.

#### Possible threats

Type of user	Information asset	Technique	Unauthorized action
General user	Database	Eavesdropping of packets	Unauthorized acquisition
Internal user System manager	management information	Dictionary attack of passwords	(viewing) of information Unauthorized tampering or
System developer System		Unauthorized acquisition of IDs/passwords through social engineering	destruction (updating) of information
administrator  System operator		Unauthorized acquisition of information through misuse of settings	
		Unauthorized acquisition of information through exploiting a database vulnerability	
		Acquisition by an unauthorized route	
General user Internal user	General database information	Acquisition by a normal route	Misuse of information that can be acquired normally (taking data outside of the company)
		SQL issued with the aim of obstructing a job	Obstructing a job (resource depletion)

Type of user	Information asset	Technique	Unauthorized action
General user Internal user	General database information	Eavesdropping of packets  Dictionary attack of passwords	Unauthorized tampering or destruction (updating) of information
		Unauthorized acquisition of IDs/passwords through social engineering	
		Unauthorized acquisition of information through exploiting configuration errors	
		Unauthorized acquisition of information through exploiting a database vulnerability	
		Acquisition by an unauthorized route	
System manager	General	Eavesdropping of packets	Unauthorized acquisition
System developer System	database information	Dictionary attack of passwords	(viewing) of information Unauthorized tampering or
administrator System operator		Unauthorized acquisition of IDs/passwords through social engineering	destruction (updating) of information
		Unauthorized acquisition of information through exploiting configuration errors	
		Unauthorized acquisition of information through exploiting a database vulnerability	
		Acquisition by an unauthorized route	
System developer	Database management information	Creation of a backdoor	Unauthorized acquisition (viewing) of information
	General database information		Unauthorized tampering or destruction (updating) of information
System manager System	Database management	Unauthorized acquisition of information by creating an	Unauthorized acquisition (viewing) of information
administrator	information  General database information	unauthorized database administrator account	Unauthorized tampering or destruction (updating) of information
System manager System operator	Database management information	Unauthorized acquisition of information by tampering with database-related files	Unauthorized acquisition (viewing) of information

Type of user	Information asset	Technique	Unauthorized action
	General database information	(definition file, physical file, and so on)	Unauthorized tampering or destruction (updating) of information
Database administrator	Database management information	Misuse of information (taking information outside of the company) after obtaining it through the normal route	Misuse of information that can be acquired normally (taking information outside of the company)
		Unauthorized use of IDs/ passwords from the management information	Tampering with or destroying information that can be acquired
		Unauthorized acquisition of information by tampering with management information	
		SQL issued with the aim of obstructing a job	Obstructing a job (resource depletion)
	General	Eavesdropping of packets	Unauthorized acquisition
	database information	Misuse of information (taking information outside of the company) after obtaining it through an unauthorized route	(viewing) of information Unauthorized tampering or destruction (updating) of information
Database operator	Database	Eavesdropping of packets	Unauthorized acquisition
	management information	Dictionary attack of passwords	(viewing) of information Unauthorized tampering or
		Unauthorized acquisition of IDs/passwords through social engineering	destruction (updating) of information
		Unauthorized acquisition of information by exploiting configuration errors	
		Unauthorized acquisition of information through exploiting a database vulnerability	
		Acquisition by an unauthorized route	
	General database information	Acquisition by a normal route	Misuse of information that can be acquired normally (taking data outside of the company)
		SQL issued with the aim of obstructing a job	Obstructing a job (resource depletion)

### Types of user

In database security, the persons involved with databases and their roles are defined below.

Type of user	Role
System manager	Manages developers, administrators, and operators
System developer	Builds the network around the database server
	Builds the database server
System administrator	Operates devices of the surrounding database network
	Operates the database server
System operator	Operates the surrounding database network
Database administrator	Builds the database system
	Operates the database system
Database operator	Performs business operations
Internal user	End user inside the company
General user	End user outside the company

#### Information assets

In database security, it is necessary to protect the information assets to be stored on the database server.

Such assets are defined below.

Database management information

- Database configuration information (system catalog, user ID/password, and so on)
- Database logs (such as access logs)

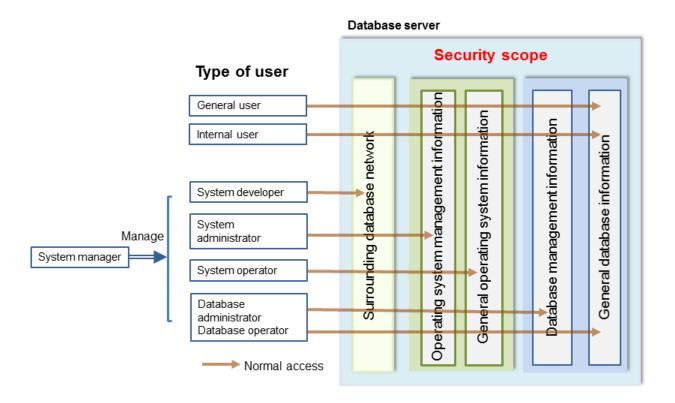
General database information

- Job data
- Applications

# 1.4 Security Scope

In database systems, both the database server and the surrounding database network are part of the security scope. It is necessary to clarify the extent of the security scope that each type of user is involved with, and consider security measures for the same.

The relationship of the security scope and the types of user is shown below.



# 1.5 Security Provided by FUJITSU Enterprise Postgres

FUJITSU Enterprise Postgres provides security features that satisfy the security requirements indicated in "1.2 Security Requirements".

This section describes security provided by FUJITSU Enterprise Postgres.

# 1.5.1 Roles Targeted For Security

In FUJITSU Enterprise Postgres database systems, the roles targeted in relation to security are "Manager", "Administrator", and "User". In order to build a robust security system, it is necessary to put security measures in place for each role.

The roles targeted for security and the mapping of Types of user indicated in "1.3 Security Threats" are shown in the table below.

Role targeted for security	Type of user
Manager	System manager
Administrator	System developer
	System administrator
	System operator
	Database administrator
	Database operator
User	General user
	Internal user

#### Manager

The manager establishes a security policy and decides on an operations policy for the organization as a whole.

Refer to "Chapter 3 Tasks of the Manager" for details.

#### Administrator

Administrators design, build and operate a system. While doing this, the administrators must implement the security measures in accordance with the security policy established by the manager.

Refer to "Chapter 4 Tasks of Administrators" for details.

#### User

A user is a person other than the manager or an administrator who accesses a database. There may be any number of users. It is necessary for users to be registered in the database system, and that access to the database is restricted according to the access privileges.

Refer to "Chapter 5 Tasks of Users" for details.

### 1.5.2 Security Features

FUJITSU Enterprise Postgres provides the following security features:

- Authentication
- Access control
- Encryption
- Audit log
- Data masking

This section describes each of these features.

#### **Authentication**

The databases that can be accessed can be restricted by authenticating the database users who access the database. Additionally, authentication of the server can be performed to prevent spoofing of the database server.

Refer to "Client Authentication" in "Server Administration" in the PostgreSQL Documentation for details on authentication.

Refer to "Secure TCP/IP Connections with SSL" in "Server Setup and Operation" in the PostgreSQL Documentation for details on server authentication.

#### **Access control**

Database objects can only be used by the object creator or database user who was specified as the owner when the object was created (both persons are hereinafter referred to as "owner"), or instance administrator, when objects are in their initial state. By having the object owner or instance administrator control access privileges for database users, it is possible to control what kind of tables the database users who connect to the database can access, and what kind of operations they can perform.

Refer to "Privileges" in "The SQL Language" in the PostgreSQL Documentation for details on object access control.

#### **Encryption**

FUJITSU Enterprise Postgres provides a transparent data encryption feature that satisfies the requirements below.

- Confidential information can be changed into an unidentifiable state.
- The encryption key and data are managed separately.
- The encryption key is replaced at regular intervals.

PostgreSQL provides an encryption feature called "pgcrypto" that can also be used in FUJITSU Enterprise Postgres, however, it is recommended to use the transparent data encryption features because it will otherwise be necessary to modify the applications that consider encryption. Refer to "Protecting Storage Data Using Transparent Data Encryption" in the Operation Guide for details.

Additionally, if communication data transferred between a client and a server contains confidential information, it is necessary to encrypt the communication data to protect it against threats, such as eavesdropping on the network.

Refer to "Configuring Secure Communication Using Secure Sockets Layer" in the Operation Guide for details on encryption of communication data.

#### **Audit log**

A feature that addresses threats such as misuse of administrator privileges, unauthorized access to a database by a user, and other such threats. Information for tracing the processing of administrators and users is retrieved and stored as an audit log.

By periodically viewing and monitoring audit logs, the administrators can detect events that are impacting on the system in some way, or are depleting system resources as a result of incorrect operations by users, and can take appropriate measures to prevent information leakages or system failures in advance.

Refer to "Chapter 6 Audit Log Feature" for details.

#### **Data masking**

A feature that changes part of the data to make it available for reference in response to queries issued by an application.

For example, for a query of employee data, digits except the last four digits of an eight-digit employee number can be changed to "\*" so that it can be used for reference without exposing the actual data.

Specifically, the data changed by the data masking feature can be transferred to a test database so that users who perform testing or development can reference the data. During testing, it is desirable to use the data that will be used on a production environment database. However, actual production data should not be used as is for testing because of the risk of leakage of confidential data. This feature enables data that is similar to actual production data to be safely used in test and development environments.

Refer to "Data Masking" in the Operation Guide for details on data masking.

# **Chapter 2 Overview of Security Operation**

# 2.1 Security Operation Flow

This section shows the flow of work when building a security environment and performing security operation in FUJITSU Enterprise Postgres.

When performing security operation, there are technical measures to be implemented to address security threats by equipping the system with security features, and manual work, such as the implementation of security guidelines, a training system, and the establishment of usage rules.

Figure 2.1 Security operation flow Administrator Manager User Define important information / perform risk analysis Formulate security policies Formulate account management policies Formulate log retrieval policies Formulate rules Implement training Receive training Receive training Check database management operations Perform initial setup Authenticate Control access Design and build security system Encrypt  $\Psi$ Control use of external media Restrict access Restrict resources Manage logs  $\Psi$ Detect unauthorized access Analyze logs

Operation

Periodically diagnose the status of security measures

Manage account/password

# Chapter 3 Tasks of the Manager

The manager formulates security policies, which become guidelines for security measures.

# 3.1 Defining Important Information and Risk Analysis

Before formulating security policies, define important information and perform risk analysis. Based on the importance of the information and the result of risk analysis, decide what kind of security measures to put in place.

In defining the important information, identify what should be protected and classify it by importance in order to effectively implement the security measures. Information that should be protected includes "database management information" and "general database information", as indicated in "Information assets". Examples of information classifications are "personal information" and "confidential information".

In the risk analysis, refer to "Possible threats" to identify threats that may arise, and analyze the risks in respect to such threats.

Additionally, by performing a risk analysis once annually as a guide, it is possible to identify threats that may adversely impact the business and related vulnerabilities.

# 3.2 Formulating Account Management Policies

In formulating an account management policy, implement the following and document the formulated policy.

#### Organize system users and roles

Identify the necessary roles of the relevant system based on "Types of user". Additionally, organize personnel for each role.

#### Organize accounts

Organize accounts with the appropriate privileges for each role, and decide on account policies.

- Database administrator account
  - Organize separate accounts for database administrators and database operators
  - Ensure that the database administrator account can only be used by specific persons
  - Perform tasks that do not require database administrator privileges using a separate account without database administrator privileges
- General account

Create an account for general users by application usage.

#### Review account management policy

Review the accounts in order to effectively implement security measures.

- Regularly check the accounts mentioned above and their privileges, and determine if they are still appropriate
- If there have been system or operational changes, review the accounts and privileges
- If unsuitable accounts and privileges are discovered, modify them as required

# 3.3 Formulating Log Retrieval Policies

In formulating a log retrieval policy, implement the following and document the formulated policy.

#### Organize the purpose of log retrieval

To clarify what logs will be retrieved for, define their reason for retrieval.

Examples of the purpose might include, "To use for investigation in the event of unauthorized access", and "To submit to investigating authorities as evidence if any issues arise".

#### Decide on the types of logs to be retrieved

In order to retrieve appropriate logs, organize the types of logs that can be retrieved in the target system, and decide on the logs to be retrieved.

Examples of log types are "operating system logs", "application run logs", and "database audit logs".

#### Organize log retrieval target access

In order to decide on access for log retrieval targets, organize what kind of access will take place.

For example, the following access is possible:

- Access related to important information
  - Access to personal information, confidential information, and database management information
  - Access outside of business hours
  - Login
  - Specific SQL
- Access suspected to be unauthorized
  - Large amount of search access
  - Access from different locations
  - Access outside of business hours

#### Decide on the log retrieval content

In order to effectively use retrieved logs, organize the required content as a log, and decide on the retrieval content.

For example, the following output content is possible:

- When (time)
- Who (database account, application user)
- What (object ID, table name)
- Where from (machine name, IP address)
- How (SQL type, SQL statement)
- Execution result (success/fail)

#### Formulate log maintenance policy

In order to use the logs as purposed, formulate the log maintenance policy.

For each log, define its location, storage medium, retention period, access control, and so on.

# 3.4 Formulating Rules

Formulate the rules that will become the standard for security measures of the target system. Additionally, prescribe penalties for security violations. For example, formulate rules and penalties as below:

- Rules
  - Applying security patches and update programs
  - Prohibiting unauthorized acquisition of information from the database
  - Prohibiting the saving of acquired information to media that is not permitted for use
- Penalties
  - Prescribe penalties in the company's employment policies and procedures

# 3.5 Implementing Training

In order to have administrators and users recognize the importance and necessity of information security, and to prevent unauthorized access due to operational omissions and mistakes, implement and promote security-related training for administrators and users.

For example, implement promotion of security policies, formulation of training schedules, and formulation of training materials.

# 3.6 Checking the Database Management Operations

In order to prevent operational errors and unauthorized actions by administrators, implement the measures below:

- Always collect the latest information on security incidents and vulnerabilities related to databases
- Implement management operations only after providing advance notice
- Retain records of management operations

# 3.7 Periodic Diagnosis of the Status of Security Measures

In order to check if the security measures are effective, periodically diagnose if the security measures have been put in place appropriately based on the security threats.

Additionally, evaluate if the current security measures and policies are effective for the threats and vulnerabilities, and if there are any issues, review the security policies and security measures.

# Chapter 4 Tasks of Administrators

Administrators perform the actions below as security measures when designing, building, and operating the system in accordance with the security policies formulated by the manager.

#### Preparation

- Implement training

#### Measures to protect against unauthorized behavior

- Perform initial setup
- Authenticate
- Control access
- Encrypt
- Control use of external media
- Restrict access
- Restrict resources

#### Measures to detect and trace unauthorized behavior

- Manage logs
- Detect unauthorized access
- Analyze logs

# 4.1 Receiving Training

Administrators receive security-related training in accordance with the training schedule formulated by the manager. Additionally, administrators instruct users to receive training.

# 4.2 Initial Setup

To minimize database vulnerabilities and the possibility of unauthorized access, implement the security measures below in the initial stage of system building. Additionally, configure the database server so that it primarily operates the database system only.

#### Making the server more robust

Configure the operating system and network to prevent intrusion into or destruction of a database server, so that the system operates on a secure server.

- Remove unnecessary features or services on the operating system
- Enable only the necessary protocols
- Implement the security features for services, protocols, and daemons considered to have a relatively low security level, such as file sharing and FTP

#### Installing the latest version

Always download and apply the latest patches in order to reflect the latest security measures.

#### Installing the minimum necessary features

Install only the necessary features in order to prevent unauthorized use of the system.

Additionally, delete or disable features and services that will not be used.

#### Changing the port

To prevent unauthorized use of the system, change the default port that is set during installation.



Specify the port during setup of FUJITSU Enterprise Postgres. Refer to the Installation and Setup Guide for Server for details.

#### Access restrictions for communication features

To prevent unauthorized use of the system using the communication features, implement access restrictions for communication features

#### Settings for prohibiting the access path to database configuration files

To prevent database destruction, implement the measures below:

- Restrict users who are permitted to access database configuration files, and periodically review the permissions
- Allow only administrators to access table or definition scripts

#### Restrictions on the access path to the database

To prevent unauthorized use or operating errors for the database, restrict the distribution range of applications used to access the database only to devices used by users who are permitted access.

#### Dealing with unauthorized programs

To prevent unauthorized intrusions into a system through a backdoor, such as by tampering with the program source code of an application, document the author of the program to be run and perform checking and testing so that the program will not be tampered with. Additionally, employ safe coding techniques so that issues with general coding vulnerabilities can be addressed.

#### System security settings

In cases where it is clear that the system security settings will impact security, set reliable security settings in the initial setup stage, such as setting appropriate security parameters.

### 4.3 Authentication

When accessing a database, authentication must always be performed in order to prevent tampering or information leakage from spoofing by a malicious user.

Password authentication is used when logging on to a database, and the account and password used for authentication are to be strictly managed by administrators.

Additionally, authentication must also be implemented reliably for connections to a database from clients, so that only permitted users can access the database.

# 4.3.1 Managing Accounts

For account management, perform the actions below.

#### Create the required accounts

To prevent unauthorized use of accounts, such as spoofing, implement the measures below when creating an account:

- Select the required account
- Specify the user privileges
- Create database administrator accounts and general user accounts separately according to the privileges



Accounts are created using the CREATE ROLE statement. Refer to "CREATE ROLE" in the PostgreSQL Documentation for details.

#### Delete unnecessary accounts

Remove accounts not used on a daily basis, such as unused accounts and accounts not needed for operations that are created by default during product installation.



Accounts are deleted using the DROP ROLE statement. Refer to "DROP ROLE" in the PostgreSQL Documentation for details.

•••••••••

#### Set up account lockout

The usage frequency of accounts is to be checked periodically, and if there are any accounts that have not been used for a long period, lock those accounts. Set a limit for failed login attempts, and if this limit is exceeded, lock the account. Additionally, set the period until a locked account is reenabled.



Account locking can be performed by using LDAP authentication. Refer to "LDAP Authentication" in the PostgreSQL Documentation for details.

#### Manage database administrator accounts

Manage database administrator accounts in accordance with the account management policy formulated by the manager.

#### Manage development environment and production environment accounts

To prevent unauthorized use of accounts used in a development environment, delete accounts used in the development environment before operation starts in the production environment. In cases where it is unavoidable to use an account used in the development environment in the production environment, use different passwords in each environment.

#### Set up a temporary use account

If a temporary user will use the system, either provide a shared account with a temporary password for each use, or create a temporary account.

# 4.3.2 Managing Passwords

Manage passwords as below.

#### Make strong passwords

The use of account passwords that can easily be guessed by others, such as a password that matches the ID, or the default password provided during installation, is prohibited. Set a complex and strong password.

#### Change passwords regularly

Change passwords regularly to prevent others from accessing the account in case the password is obtained by unauthorized means. Additionally, configure the settings to force a password change when prompted after the first use.

#### Set the password expiry period

To encourage regular changing of passwords, set a password expiry period.



Password setting and changing is specified using the CERATE ROLE statement or ALTER ROLE statement. Refer to "CREATE ROLE" and "ALTER ROLE" in the PostgreSQL Documentation for details.

Additionally, by using passwordcheck and LDAP authentication, the actions below can be performed:

- The default password set during installation can be changed
- The password expiry period can be set

- The number and types of characters used for the password can be checked

Refer to "passwordcheck" and "LDAP Authentication" in the PostgreSQL Documentation for details.

### 4.3.3 Configuring Connections and Authentication

Configure connections and authentication so that the database can only be accessed by permitted users.



### Point

Client authentication is configured in pg\_hba.conf. Refer to "Client Authentication" in the PostgreSQL Documentation for

### 4.4 Access Control

If appropriate access privileges are not set for administrators and users, security incidents may occur, such as information leakage resulting from access to information by an unauthorized person. To minimize such incidents, it is necessary to implement the security measures below for the access privileges and perform rule-based access control.



#### Notes when setting access privileges

- The creation of a special account that allows granting of privileges to all users is prohibited
- The creation of a general account that allows access to general information such as operations data is prohibited

#### Identifying the database access requirements

To set the appropriate access privileges for each usage purpose for the database, follow the procedure below to identify the access requirements:

- 1. Classify the usage purpose of the account, such as "For database management", "For object management", and "For data access".
- 2. Classify the required privileges for each usage purpose, such as "By feature" and "By object".
- 3. Categorize the accounts based on each privilege.
- 4. Identify the minimum necessary range of data and minimum necessary access content (view, update, create, delete) to be accessed for each categorized account, and decide on the database access requirements.

#### Setting the access privileges

Assign the minimum necessary privileges based on the database access requirements for each categorized account. Additionally, restrict accounts when assigning administrator privileges.

#### Reviewing access privileges

To reflect changes in access requirements in the system, periodically review the access privileges and check if there are any access privileges that are no longer needed. If any unnecessary access privileges have been set, promptly modify the access privileges.



Access privileges are set using the GRANT statement or REVOKE statement. Refer to "GRANT" and "REVOKE" in the PostgreSQL Documentation for details.

# 4.5 Encryption

To prevent unauthorized usage of data in the event information leakage occurs due to data theft, eavesdropping of communication, and other such activities, implement the encryption measures below.

#### **Encrypt communication**

To protect data from eavesdropping over the network between a database server and clients, use the encryption feature to encrypt communications.

Refer to "Configuring Secure Communication Using Secure Sockets Layer" in the Operation Guide for details.

#### Encrypt data

To protect data from theft, use the encryption feature to encrypt the data. The data below is targeted for encryption:

- Data to be stored on the database
- Backup data
- Data files

Refer to "Protecting Storage Data Using Transparent Data Encryption" in the Operation Guide for details.

#### Manage encryption keys

Restrict the persons who can access the encryption key to a minimum number of database administrators.

Additionally, to ensure the encrypted information will not be easily decrypted, create a mechanism for appropriately managing the encryption key for the entire life cycle (generation, distribution, saving, and disposal), and strictly manage the encryption key.

Refer to "Configuring Secure Communication Using Secure Sockets Layer" and "Protecting Storage Data Using Transparent Data Encryption" in the Operation Guide for details.

# 4.6 Controlling Use of External Media

Information leakage can be prevented by controlling use of external media (such as CD/DVD, USB drive, and external hard disk) and PCs that are connected to the database, and restricting the removal of data from the database.

#### Restricting connection of external media

Remove external media and printers that will not be used in operations, and restrict connection of external media to which information may be written.

#### Restricting use of external media

Restrict connections for external media and printers to control the writing of information to these devices.

#### Controlling use of connected PCs

Prevent leakage of information from PCs connected to the database:

- Limit connections of external media to PCs
- Implement security measures to make the PC robust
- Implement individual authentication for access from the PC
- Manage the installed software and monitor the software usage status
- Limit connections to printers

# 4.7 Security Measures for Servers/Applications

An even more robust security system can be achieved by strengthening security for servers and applications in addition to the security measures for databases. Implement the security measures below for servers and applications:

#### Restrict access

Implement the measures below and restrict access to the database server:

- Install the database server inside the firewall to prevent direct access to the database server from many unspecified PCs.
- In the local network, implement measures such as using the router to restrict IP addresses, and restrict PCs and segments that can directly access the database server.

#### Restrict resources

Restrict excessive use of CPU resources by general users to prevent the disruption of service and extraction of large amounts of data.

### 4.8 Log Management

Logs are a feature that addresses threats such as misuse of administrator privileges, and unauthorized access to a database by a user. Information for investigating/tracing processes and operations performed for the database is retrieved and managed as logs for identifying the cause in the event information leakage or unauthorized access occurs.

FUJITSU Enterprise Postgres provides the audit log feature for retrieving and managing logs. Refer to "Chapter 6 Audit Log Feature" for details.

This section describes the information that should be obtained as logs and how to maintain logs, as a measure for managing information leakage and unauthorized access.

### 4.8.1 Retrieving Logs

The audit logs below are retrieved in accordance with the log retrieval policy formulated by the manager.

#### Login information

Retrieves logs during login and logout.

#### Database access information (view/update)

Retrieves all access relating to the information below:

- General database information (such as personal information and confidential information used in the business)
- Database management information (system catalog, user ID/password, and so on)

#### Changed information of database objects

Retrieves logs related to creating, changing, and deleting database objects such as database accounts and tables.

#### Operation logs for audit logs

To prevent suppression of retrieved audit logs, operations such as initialization of audit logs, and stoppage of the audit log feature are retrieved as logs.

# 4.8.2 Maintaining Logs

Logs are maintained in accordance with the log retrieval policy formulated by the manager.

#### Storing logs

Perform the actions below and store logs securely so that the retrieved logs will not be updated by others:

- Save logs to external media, and store the external media in a secure location, such as lockable storage
- Restrict the viewing of logs to administrators only, and set access restrictions for logs, such as not assigning update rights
- Decide on the log retention period, with consideration to cases where investigation tracing back to the time of discovery of an issue is required

#### Preventing tampering of logs

Implement measures to prevent tampering of logs, such as retaining multiple copies of logs and using storage that cannot be rewritten

#### **Encrypting logs**

Encrypt logs so that logs are not easily viewed.

# 4.9 Detecting Unauthorized Access

To address unauthorized access, it is necessary to establish a mechanism for detecting unauthorized access to databases and monitor access.

#### Communicating unauthorized access

Create a mechanism that notifies of detected unauthorized access, such as notifying the manager and the administrator, if an account lock occurs due to the limit for failed login attempts being exceeded.

#### Checking access times

Create a mechanism that can check for suspicious access to the information below outside of normal access hours, together with implementing measures to address such access.

#### Detecting access to database management information

- Monitor logs and detect access during timeframes that have not been applied for
- In the event a request for access permission outside of normal access hours is made, the log is checked for any discrepancies in the requested content and work result

#### Detecting access to general database information

- Decide on the timeframes during which access to the database is permitted for each general account
- Detect access outside of normal access hours from session information logs

#### Checking the connection source where access is not permitted

To detect access from connection sources that are not permitted, define the sources from where access is permitted, and detect access from connection sources that are not permitted.

Define the access patterns (connection source, operating system user and account) of database administrator accounts and general accounts, and check for access outside of these patterns.

# 4.10 Analyzing Logs

Create a mechanism that analyzes logs to detect unauthorized behavior in cases where information leakage, unauthorized access, or other such activity, is suspected. Analyses should include those shown below.

#### Periodic analysis of session information

Analyze session information of logs from the perspectives below to detect unauthorized logins:

- Trend of sessions with a large number of failed login attempts
- Trend of sessions with accounts that are logged in for long periods of time
- Trend of sessions in which a large amount of resources are used

#### Periodic analysis of database access information

Analyze SQL statements from the perspectives below to detect unauthorized access to the database:

- Trend of SQL being executed over a long period of time
- Trend of SQL using a large amount of resources

# Chapter 5 Tasks of Users

The user performs the actions below as security measures when using the system.

# 5.1 Receiving Training

The user must receive security-related training as instructed by the manager or the administrator to learn about security. By having users with a common awareness relating to security, an even more stable security system can be established.

# 5.2 Managing Accounts/Passwords

Users can use the database system by using the account and password provided by the administrator. At such times, the user is to implement the measures below so that the account and password are not misused by others:

- Be responsible for managing the ID and password in a manner that ensures the account does not become locked during login
- Change the password regularly
- Comply with the expiry period set for the password by promptly changing the password when it is about to expire

# Chapter 6 Audit Log Feature

In PostgreSQL, logs output as server logs can be used as audit logs by using the log output feature. There are, however, logs that cannot be analyzed properly, such as SQL runtime logs, which do not output the schema name. Additionally, because the output conditions cannot be specified in detail, log volumes can be large, which may lead to deterioration in performance.

The audit log feature of FUJITSU Enterprise Postgres enables retrieval of details relating to database access as an audit log by extending the feature to pgaudit. Additionally, audit logs can be output to a dedicated log file or server log. This enables efficient and accurate log monitoring.



The audit log feature cannot be used if PostgreSQL is running in single-user mode.

# 6.1 Audit Log Output Modes

In pgaudit, the two types of audit log below can be output.

#### Session Audit Logging

Session Audit Logging outputs information related to SQL executed in backend processes (processes generated when connection requests are received from clients), information related to starting and connecting databases, and information related to errors, as a log. In Session Audit Logging, by specifying the log output conditions and filtering the logs to be output, performance degradation due to outputting large volumes of logs can be prevented.

Refer to "6.4 Session Audit Logging" for details.

#### Object Audit Logging

When SELECT, INSERT, UPDATE, and DELETE are executed for specific objects (tables, columns), Object Audit Logging outputs these as a log. TRUNCATE is not supported. Object Audit Logging outputs object operations for which privileges have been assigned to specified roles, as a log. Object Audit Logging can control log output at an even finer level of granularity than Session Audit Logging.

Refer to "6.5 Object Audit Logging" for details.



Depending on the application or command, FUJITSU Enterprise Postgres may execute SQL internally and the audit logs may be retrieved

Also, the audit logs of multiple SQLs with the same statement ID may be retrieved. This is because before the user executes the SQL, another SQL is executed internally by FUJITSU Enterprise Postgres.

# 6.2 Setup

This section describes the setup method of pgaudit.

1. Copy the pgaudit files

As superuser, run the following command. Note that " $\langle x \rangle$ " in paths indicates the product version.

```
$ su -
Password:*****
# cp -r /opt/fsepv<x>server64/OSS/pgaudit/* /opt/fsepv<x>server64
```

Open a command prompt as administrator privileges, run the following command. Note that "<*x*>" in paths indicates the product version.

> xcopy /E "c:\Program Files\Fujitsu\fsepv<x>server64\OSS\pgaudit\\*" "c:\Program Files\Fujitsu\fsepv<x>server64"

#### 2. Create the pgaudit configuration file

Create the pgaudit configuration file, which describes the information required for pgaudit actions. Create the file using the same encoding as used for the database.

In addition, set write permissions for the database administrator only in the pgaudit configuration file so that policies related to the audit log are not viewed by unintended users.

Refer to "6.3 pgaudit Configuration File" for details.



Do not define the rule section in the pgaudit configuration file at this point.

#### Example of a pgaudit configuration file

[output] logger = 'auditlog'

#### 3. Configure postgresql.conf

Configure the parameters below in postgresql.conf to use audit logs:

shared\_preload\_libraries

Specify "pgaudit".

pgaudit.config\_file

Specify the deployment destination path of the pgaudit configuration file.

If a relative path is specified, the path will be relative to the data storage directory.

log\_replication\_commands

Specify "on".

log\_min\_messages

Check if "ERROR" or higher has been specified.

If outputting an audit log to a server log ("serverlog" is specified in the logger parameter of the pgaudit configuration file), check the parameters below relating to server logs.

logging\_collector

Check if "on" has been specified.

log\_destination

Check if "stderr" has been specified.

#### log\_file\_mode

Check if the server log permissions are appropriate, so that only the permitted persons can access it.



The default for the log\_file\_mode parameter is 0600, which only allows the database administrator to have access.

For example, to permit other members of the group to which the database administrator belongs to view the audit logs, specify 0640 for log\_file\_mode.

- 23 -

#### Example

```
log_file_mode = 0640
```

The database administrator can also be prevented from viewing audit logs by specifying 0000. However, write privileges are assigned for outputting logs.

If outputting an audit log to a dedicated log file ("auditlog" is specified in the logger parameter of the pgaudit configuration file), check the parameter below.

max\_worker\_processes

If the max\_worker\_processes parameter has been set, add 1 to the specified value.



#### See

Refer to "Error Reporting and Logging" in the PostgreSQL Documentation for details on server logs.

If using database multiplexing, refer to "6.6 Database Multiplexing" for details.

#### Example of postgresql.conf

In the example below, only the parameters that need to be configured when using the audit log feature are described.

```
shared_preload_libraries = 'pgaudit'
pgaudit.config_file = 'pgaudit.conf'
log_replication_commands = on
log_min_messages = WARNING
```

#### 4. Start the instance

Start the instance and check if the message below is output.

```
LOG: pgaudit extension initialized
```

5. Create the pgaudit extension

Execute CREATE EXTENSION to create the pgaudit extension.

6. Configure the parameters in the pgaudit configuration file

Add or change the parameters in the pgaudit configuration file as required.

Refer to "6.3 pgaudit Configuration File" for details.

7. Restart the instance

Restart the instance to apply the changes to the pgaudit configuration file. After restarting, check if the changes have been reflected correctly.

Linux

```
LOG: log_catalog = 1
LOG: log_level_string =
LOG: log_level = 15
     log_parameter = 0
LOG: log_statement_once = 0
LOG: role =
LOG: logger = auditlog
LOG: log_directory = pgaudit_log
LOG: log_filename = pgaudit-%Y-%m-%d_%H%M%S.log
LOG: log_file_mode = 0600
LOG: log_rotation_age = 1440
LOG: log_rotation_size = 10240
LOG: log_truncate_on_rotation = 0
LOG: fifo_directory = /tmp
LOG: Rule 0
LOG: pgaudit extension initialized
```

W

#### Windows

```
LOG: log_catalog = 1
LOG: log_level_string =
LOG: log_level = 15
LOG: log_parameter = 0
LOG: log_statement_once = 0
LOG: role =
LOG: logger = auditlog
LOG: log_directory = pgaudit_log
LOG: log_filename = pgaudit_*Y-%m-%d_%H%M%S.log
LOG: log_file_mode = 0600
LOG: log_rotation_age = 1440
LOG: log_rotation_size = 10240
LOG: log_truncate_on_rotation = 0
LOG: Rule 0
LOG: pgaudit extension initialized
```

# 6.3 pgaudit Configuration File

In the pgaudit configuration file, specify the information required for pgaudit actions. The pgaudit configuration file comprises three sections: "output section", "option section", and "rule section".

#### output section

The output section is specified using the format below:

- paramName = 'value'

The valid parameters in the output section are shown in the table below.

L

#### Linux

Parameter name	Description	Remarks
logger	Dedicated log file (auditlog)/serverLog (serverlog) that will be the output destination of the audit log  The default is "auditlog" (dedicated log file).	The dedicated log file is output using the same encoding as used for the database.
log_directory	Directory where the audit log is to be created  Specify the full path or the relative path from the data storage directory.	Enabled only if "auditlog" is

Parameter name	Description	Remarks
	The default is "pgaudit_log".	specified for the logger parameter
log_filename	File name of the audit log  Specify a file name that varies according to the time, in the same manner as for log_filename in the postgresql.conf file.  The default is "pgaudit-%Y-%m-%d_%H%M%S.log".	Enabled only if "auditlog" is specified for the logger parameter
log_file_mode	Specify the permissions of the audit log so that only permitted persons can access it.  The parameter value is the numeric mode specified in the format permitted in chmod and umask system calls. The default is "0600".  Refer to "log_file_mode" in "6.2 Setup" for information on audit log file permissions.	Enabled only if "auditlog" is specified for the logger parameter
log_rotation_age	Maximum age of the audit log file  A new audit log file is generated when the time (minute units) specified here elapses. To disable generation of new log files based on time, specify "0".  The valid units are "min" (minutes), "h" (hours), and "d" (days). If the unit is omitted, "min" will be used.  The default is "1d" (1 day).	Enabled only if "auditlog" is specified for the logger parameter
log_rotation_size	Maximum size of the audit log file  A new log file will be generated after logs of the size specified here are output to a log file. To disable generation of new log files based on size, specify "0".  The valid units are "kB" (kilobytes), "MB" (megabytes), and "GB" (gigabytes). If the unit is omitted, "kB" will be used.  The default is "10MB".	Enabled only if "auditlog" is specified for the logger parameter
log_truncate_on_rotation	If rotating audit log files based on time, this parameter is used to specify whether to overwrite (on)/not overwrite (off) existing audit log files of the same name. For example, if "on" is specified, and "pgaudit-%H.log" is specified for log_filename, 24 separate log files will be generated based on time, and those files will be cyclically overwritten.  The default is "off". If "off" is specified, the logs will be written to the existing audit log files.	Enabled only if "auditlog" is specified for the logger parameter
fifo_directory	FIFO (named pipe) directory to be used between the daemon process that outputs audit log files and the backend process  FIFO named p.PGAUDIT.nnnn (nnnn is the postmaster PID) are created in the fifo_directories directory. The files cannot be deleted manually.  The default is "/tmp".	Enabled only if "auditlog" is specified for the logger parameter

W

Parameter name	Description	Remarks
logger	Dedicated log file (auditlog)/serverLog (serverlog) that will be the output destination of the audit log  The default is "auditlog" (dedicated log file).	The dedicated log file is output using the same encoding as used for the database.
log_directory	Directory where the audit log is to be created  Specify the full path or the relative path from the data storage directory.  The default is "pgaudit_log".	Enabled only if "auditlog" is specified for the logger parameter
log_filename	File name of the audit log  Specify a file name that varies according to the time, in the same manner as for log_filename in the postgresql.conf file.  The default is "pgaudit-%Y-%m-%d_%H%M%S.log".	Enabled only if "auditlog" is specified for the logger parameter
log_file_mode	This parameter is ignored in Windows.	
log_rotation_age	Maximum age of the audit log file  A new audit log file is generated when the time (minute units) specified here elapses. To disable generation of new log files based on time, specify "0".  The valid units are "min" (minutes), "h" (hours), and "d" (days). If the unit is omitted, "min" will be used.  The default is "1d" (1 day).	Enabled only if "auditlog" is specified for the logger parameter
log_rotation_size	Maximum size of the audit log file  A new log file will be generated after logs of the size specified here are output to a log file. To disable generation of new log files based on size, specify "0".  The valid units are "kB" (kilobytes), "MB" (megabytes), and "GB" (gigabytes). If the unit is omitted, "kB" will be used.  The default is "10MB".	Enabled only if "auditlog" is specified for the logger parameter
log_truncate_on_rotation	If rotating audit log files based on time, this parameter is used to specify whether to overwrite (on)/not overwrite (off) existing audit log files of the same name. For example, if "on" is specified, and "pgaudit-%H.log" is specified for log_filename, 24 separate log files will be generated based on time, and those files will be cyclically overwritten.  The default is "off". If "off" is specified, the logs will be written to the existing audit log files.	Enabled only if "auditlog" is specified for the logger parameter



If the logger parameter is set to "serverlog", audit logs will be output to the server log as log messages, therefore the status information and message severity level according to the log\_line\_prefix parameter in postgresql.conf will be output to the beginning of the audit log.

If the logger parameter is omitted or set to "auditlog", audit logs will be output to a dedicated log file as dedicated logs, therefore the status information and message severity level according to the log\_line\_prefix parameter in the postgresql.conf file will not be output.

Refer to "Output format" in "6.4 Session Audit Logging" or "Output format" in "6.5 Object Audit Logging" for information on the output format of audit logs.

L



The pgaudit log\_file\_mode configuration parameter setting is separate from, and unaffected by, the log\_file\_mode GUC parameter setting and the -g/-allow-group-access initdb option.

When using a dedicated pgaudit log\_file, since the pgaudit log\_directory location defaults to inside the data storage directory, it is possible for the pgaudit log\_file\_mode permissions to conflict with the intended file permissions specified by the -g/-allow-group-access initdb option. In this case, the pgaudit log\_directory should be specified to be a directory located outside of the data storage directory.

W



If the file permissions for the server log do not satisfy the security requirements of the system, specify a dedicated log file for the output destination of the audit log, and set the access permissions for the output destination directory so that the security requirements are satisfied. At that time, assign read and write permissions to the database administrator.

#### option section

The option section is specified using the format below:

- paramName = 'value'

The valid parameters in the option section are shown in the table below.

Parameter name	Description	Remarks
role	Name of roles used in Object Audit Logging  If specifying a name containing uppercase characters, key words, multibyte characters and commas, enclose the name in double quotation marks.	Parameter used in Object Audit Logging only
log_catalog	Whether to enable (on)/disable (off) log output for pg_catalog  If pgAdmin, and so on, will not be retrieving audit logs that access pg_catalog, specify "off".  The default is "on" (enabled).	
log_parameter	Whether to enable (on)/disable (off) output of values passed by parameters in SQL execution  The default is "off" (disabled).	
log_statement_once	Whether to control (on)/not control (off) output for the second and subsequent SQL statements if the same SQL statement is the log output target  The default is "off" (do not control).	

Parameter name	Description	Remarks
log_level	Log level of audit logs  The valid values are "DEBUG5", "DEBUG4", "DEBUG3", "DEBUG2", "DEBUG1", "INFO", "NOTICE", "WARNING", and "LOG".  The default is "LOG".	Enabled only if "serverlog" is specified for the logger parameter

#### rule section

The rule section is used in Session Audit Logging. Refer to "6.4 Session Audit Logging" for details.



Do not specify the rule section if the role parameter has been specified in the option section. If you specify the rule section, the audit logs of Object Audit Logging and Session Audit Logging will be output intermingled and you will be unable to view the logs in CSV format.

# 6.4 Session Audit Logging

In Session Audit Logging, specify the rules for filtering logs to be output in the rule section in the pgaudit configuration file.

Rules are specified using the formats below. Multiple values can be specified, using a comma as the delimiter.

- paramName = 'value'
- paramName != 'value'

If [rule] is described on its own in the rule section with no parameters specified, all audit logs will be output.

#### Example

```
[output]
logger = 'auditlog'
[rule]
```

If the rule section is not described ([rule] is not described), the audit logs will not be output.

#### Example

```
[output]
logger = 'auditlog'
```

The valid parameters in the rule section are shown in the table below.

Parameter name	Description	Example
timestamp	Timestamp range  Refer to "timestamp" for details on how to specify timestamps.	timestamp = '09:00:00 - 10:00:00, 18:00:00 - 18:30:00'
database	Database name  To specify a blank value, use a pair of double quotation marks (""). When specifying a name containing uppercase characters, key words, multibyte	database = 'prodcut_db'

Parameter name	Description	Example
	characters and commas, enclose the name in double quotation marks.	
audit_role	Role name	audit_role = 'appuser1'
	To specify a blank value, specify use a pair of double quotation marks (""). When specifying a name containing uppercase characters, key words, multibyte characters and commas, enclose the name in double quotation marks.	
class	Operation class	class = 'READ, WRITE'
	Select from the values below. Multiple values can be specified. Refer to "class" for details on the meaning of each class.	
	- BACKUP	
	- CONNECT	
	- DDL	
	- ERROR	
	- FUNCTION	
	- MISC	
	- READ	
	- ROLE	
	- WRITE	
	- SYSTEM	
object_type	Object type	object_type = 'TABLE, INDEX'
	This parameter is enabled when the class parameter is "READ" and "WRITE".	
	Select from the values below. Multiple values can be specified.	
	- TABLE	
	- INDEX	
	- SEQUENCE	
	- TOAST_VALUE	
	- VIEW	
	- MATERIALIZED_VIEW	
	- COMPOSITE_TYPE	
	- FOREIGN_TABLE	
	- FUNCTION	
object_name	Object name	object_name = 'myschema.tbl1'
	This parameter is enabled when the class parameter is "READ" and "WRITE".	object_name = 'myschema.tbl1, " <i>mySchema</i> .TABLE"'

Parameter name	Description	Example
	For objects that can be modified by a schema, such as a table, modify such objects by schema name.	
When specifying a name containing uppercase characters, key words, multibyte characters and commas, enclose the name in double quotation marks. When specifying the object name "schemaName.tableName", enclose the entire object name modified by schema name in double quotation marks.		
	To specify a blank value, use a pair of double quotation marks ("").	
application_name	Application name	application_name = 'myapp'
	To specify a blank value, use a pair of double quotation marks ("").	
remote_host	Connection destination host name or IP address	remote_host = 'ap_server'
	If "on" is specified for the log_hostname parameter in the postgresql.conf file, specify a host name. Otherwise, specify the IP address. If using a local host, specify "[local]".	
	To specify a blank value, use a pair of double quotation marks ("").	

#### timestamp

Specify a timestamp range from "*startTime*" to "*endTime*" for the log output target. The timestamp format is 'hh:mm:dd-hh:mm:dd' (hh is expressed in 24-hour notation, and hh, mm, and dd are expressed in two-digit notation).

The start time must be earlier than the end time. If specifying multiple ranges, specify each start and end timestamp using a comma as the delimiter.

End timestamps consider milliseconds. For example, if '11:00:00 - 11:59:59' is specified for the timestamp, "11:00:00:000" to "11:59:59:999" will be the target range.

The timestamps used by evaluation in the rule section of pgaudit are different to the timestamps issued in the log entries. That is because log entries are output after evaluation by pgaudit, with the timestamp being generated at that time.

#### class

The meaning of each class specified in the class parameter is below:

- READ: SELECT, COPY FROM
- WRITE: INSERT, UPDATE, DELETE, TRUNCATE, COPY TO
- FUNCTION: Function call, DO
- ROLE: GRANT, REVOKE, CREATE ROLE, ALTER ROLE, DROP ROLE
- DDL: All DDLs (such as CREATE and ALTER) other than the DDLs of the ROLE class
- CONNECT: Events relating to connecting (request, authenticate, and disconnect)
- SYSTEM: Instance start, promotion to primary server
- BACKUP: pg\_basebackup

- ERROR: Event completed by an error (PostgreSQL error codes other than 00). This class can be used if ERROR or lower level is specified for the log\_min\_messages parameter in postgresql.conf.
- MISC: Other commands (such as DISCARD, FETCH, CHECKPOINT, and VACUUM)

#### Evaluation of the rule section

- When a log event occurs, all expressions in the rule section are evaluated at once. Log entries are only output if all parameters in the rule section are evaluated as being true.

For example, if the rule below has been set, of the operations performed by 'apserver' to 'myschema.tbl1', the operations applicable to classes other than 'WRITE" in the period from 10 a.m. to 11 a.m. will be output as audit logs.

```
[rule]
timestamp = '10:00:00-11:00:00'
remote_host = 'apserver'
object_name = 'myschema.tbll'
class != 'WRITE'
```

- Multiple rule sections can be defined in the pgaudit configuration file. Log events are evaluated using each rule section, and an audit log is output for each matching rule section.

For example, if the rules below are set, duplicated audit logs will be output.

```
[rule]
object_name = 'myschema.tbll'
[rule]
object_name = 'myschema.tbll'
```

- If the same parameter is specified multiple times in one rule section, the last specified parameter is effective.

For example, if the rule below has been set, "object\_name = 'myschema.tbl3'" will take effect.

```
[rule]
object_name = 'myschema.tbl1'
object_name = 'myschema.tbl2'
object_name = 'myschema.tbl3'
```

#### **Output format**

In Session Audit Logging, audit logs are output in the format below:

```
AUDIT: SESSION,READ,2017-03-12 19:00:58 PDT,

(1) (2) (3)

[local],19944,psql,appuser,postgres,2/8, 2, 1,SELECT,,TABLE,myschema.account,,

(4) (5) (6) (7) (8) (9)(10)(11)(12)(13)(14) (15) (16)

SELECT * FROM myschema.account;,<not logged>

(17) (18)
```

No	Content
(1)	Log header
	Fixed as "AUDIT: SESSION".
(2)	Class
(3)	SQL start time
(4)	Remote host name
	If using a local host, [local] is output.
(5)	Backend process ID

No	Content
(6)	Application name
	If an application name is not specified, [unknown] is output.
(7)	User name
(8)	Database name
(9)	Virtual transaction ID
(10)	Statement ID
(11)	Substatement ID
(12)	Command tag
(13)	SQLSTATE
(14)	Object type
(15)	Object name
(16)	Error message
(17)	SQL
	If the SQL contains a password, such as for CREATE ROLE, and so on, it will be replaced with " <redacted>".</redacted>
	Additionally, if "on" is specified for the log_statement_once parameter of the option section in the pgaudit configuration file, " <pre>previously logged&gt;"</pre> is output for the second and subsequent statements.
(18)	Depending on the log_parameter parameter value of the option section in the pgaudit configuration file, the output content will be as below.
	- log_parameter=on is specified
	If parameters are specified in the SQL, the parameters are concatenated and output, using a space as the delimiter.  If parameters are not specified in the SQL, " <none>" is output.</none>
	- log_parameter=off (default) is specified
	" <not logged="">" is output.</not>
	Additionally, if "on" is specified for the log_statement_once parameter of the option section in the pgaudit configuration file, " <pre>reviously logged&gt;"</pre> is output for the second and subsequent statements.



If accessing resources that use the features below, the command tag (12) may be output as "UNKNOWN".

- INSTEAD OF trigger
- RULE
- VIEW
- Security policy per row
- Table inheritance

# **Example**

Below is an example of retrieving audit logs in Session Audit Logging.

# 1. Settings

In the pgaudit configuration file, specify the rule section below.

```
[rule]
class = 'READ, WRITE'
object_name = 'myschema.account'
```

## 2. Retrieving logs

Execute the SQL below from the client.

```
CREATE TABLE myschema.account
(
    id int,
    name text,
    password text,
    description text
);
INSERT INTO myschema.account (id, name, password, description) VALUES (1, 'userl', 'HASHl', 'blah, blah');
SELECT * FROM myschema.account;
```

The audit log below can be retrieved.

'DDL' is not defined in the class parameter, so CREATE TABLE is not output as an audit log.

```
AUDIT: SESSION, WRITE, 2017-03-12 19:00:49 PDT, [local], 19944, psql, appuser, postgres, 2/7,1,1, INSERT, TABLE, myschema.account, "INSERT INTO myschema.account (id, name, password, description) VALUES (1, 'userl', 'HASHl', 'blah, blah'); ", <not logged>
AUDIT: SESSION, READ, 2017-03-12 19:00:58 PDT, [local], 19944, psql, appuser, postgres, 2/8, 2, 1, SELECT, TABLE, myschema.account, SELECT * FROM myschema.account; <not logged>
```

# 6.5 Object Audit Logging

In Object Audit Logging, retrieval of audit logs is achieved by using roles.

Roles are specified in the role parameter of the option section to retrieve audit logs. If there are privileges for commands executed by a role, or if privileges have been inherited from another role, those command operations are output as audit logs.

For example, after "auditor" is set for the role parameter of the option section, the SELECT and DELETE privileges for the account table are assigned to "auditor". In this case, when SELECT or DELETE is executed for the account table, audit logs are output.

# **Output format**

In Object Audit Logging, audit logs are output in the format below:

```
AUDIT: OBJECT,1,1,READ,SELECT,TABLE,public.account,SELECT password FROM account;,<not logged>
(1) (2)(3)(4) (5) (6) (7) (8) (9)
```

No	Content
(1)	Log header
	Fixed as "AUDIT: OBJECT".
(2)	Statement ID
(3)	Substatement ID
(4)	Class name

No	Content	
(5)	Command tag	
(6)	Object type	
(7)	Object name	
(8)	SQL	
	If "on" is specified for the log_statement_once parameter of the option section in the pgaudit configuration file, " <pre>reviously logged&gt;" is output for the second and subsequent statements.</pre>	
(9)	Depending on the log_parameter parameter value of the option section in the pgaudit configuration file, the output content will be as below.	
	- When log_parameter=on	
	If parameters are specified in the SQL, the parameters are concatenated and output, using a comma as the delimiter.  If parameters are not specified in the SQL, " <none>" is output.</none>	
	- When log_parameter=off (default)	
	" <not logged="">" is output.</not>	
	Additionally, if "on" is specified for the log_statement_once parameter of the option section in the pgaudit configuration file, " <pre>reviously logged&gt;"</pre> is output for the second and subsequent statements.	



If accessing resources that use the features below, the command tag (5) may be output as "UNKNOWN".

- INSTEAD OF trigger
- RULE
- VIEW
- Security policy per row
- Table inheritance

# Example

Below is an example of retrieving logs in Object Audit Logging.

By setting the target for assigning privileges to roles in detail, log output can be controlled.

In the example below, log retrieval of the account table is controlled by the privileges assigned to the columns, however, log retrieval of the account\_role\_map table is controlled by the privileges assigned to the table.

# 1. Settings

The role parameter below is specified for the option section in the pgaudit configuration file.

```
[option]
role = 'auditor'
```

# 2. Defining a role

A role is defined for Object Audit Logging.

```
CREATE USER auditor NOSUPERUSER LOGIN;
```

#### 3. Retrieving logs

Execute the SQL below from the client.

```
CREATE TABLE account
   id int,
   name text,
   password text,
   description text
GRANT SELECT (password) ON public.account TO auditor;
SELECT id, name FROM account;
SELECT password FROM account;
GRANT UPDATE (name, password) ON public.account TO auditor;
UPDATE account SET description = 'yada, yada';
UPDATE account SET password = 'HASH2';
CREATE TABLE account_role_map
   account id int,
   role_id int
GRANT SELECT ON public.account_role_map TO auditor;
SELECT account.password, account_role_map.role_id
  INNER JOIN account_role_map ON account.id = account_role_map.account_id;
```

The audit log below can be retrieved.

In the account table, only the operations for columns that privileges have been assigned to are output as logs.

In the account\_role\_map table, privileges are assigned to the table, so operations performed for the table are output as logs.

# 6.6 Database Multiplexing

This section describes audit log retrieval while in database multiplexing mode.

# 6.6.1 Setup

If setting up the audit log feature in a database multiplexing environment that has already been built, follow the procedure below.

#### 1. Copy the pgaudit files

Copy the pgaudit files on the primary server and standby server.

Refer to step 1 in "6.2 Setup" for details on copying the pgaudit files.

#### 2. Create the pgaudit configuration file

Create the pgaudit configuration file on the primary server. Copy the pgaudit configuration file you created to the standby server.

Refer to step 2 in "6.2 Setup" for details on creating the pgaudit configuration file.

#### 3. Configure postgresql.conf

In the postgresql.conf file on the primary server and standby server, configure the parameters for using audit logs. Set the same values for the parameters.

Refer to step 3 in "6.2 Setup" and "6.6.2 Configuring Audit Log Retrieval" for details on the parameters to configure.

#### 4. Configure the serverIdentifier.conf file of Mirroring Controller

In the *serverIdentifier*.conf file on the primary server and standby server, configure the parameters for using audit logs.

Refer to "6.6.2 Configuring Audit Log Retrieval" for details on the parameters to be set.

#### 5. Start the instance

Start the instance of the primary server and standby server.

#### 6. Create the pgaudit extension

Execute CREATE EXTENSION on the primary server to create a pgaudit extension.

Refer to step 5 in "6.2 Setup" for details on creating pgaudit extensions.

#### 7. Configure the parameters in the pgaudit configuration file

Add/change the parameters of the pgaudit configuration file on the primary server. Copy the pgaudit configuration file with the added/changed parameters to the standby server.

Refer to "6.3 pgaudit Configuration File" and "6.6.2 Configuring Audit Log Retrieval" for details on the parameters to set.

## 8. Restart the instance

Restart the instance of the primary server and standby server.

# 6.6.2 Configuring Audit Log Retrieval

In database multiplexing mode, Mirroring Controller periodically accesses the database to check the multiplexing status and detect failure. Due to this, audit logs are also periodically retrieved, so log files become used up. Therefore, set the parameters below so that audit logs are not retrieved by Mirroring Controller.

# postgresql.conf

log\_connections

Omit, or specify "off".

log\_disconnections

Omit, or specify "off".

serverIdentifier.conf file of Mirroring Controller

target\_db

Specify "template1".



If creating a new database, create it after stopping Mirroring Controller, or specify a name other than "template1" for the template database.

the template database.

pgaudit configuration file

rule section database

Specify database != 'template1'.

# 6.7 Failover

This section explains audit log retrieval performed by the failover feature integrated with the cluster software.

# 6.7.1 Configuring Audit Log Retrieval

The failover feature accesses the database periodically to check the instance status. Due to this, audit logs are also periodically retrieved, so log files become depleted. Therefore, set the parameters below so that audit logs are not retrieved.

postgresql.conf

log\_connections

Omit, or specify "off".

log\_disconnections

Omit, or specify "off".

pgaudit configuration file

rule section database

Specify database != 'template1'.

# 6.8 View Audit Logs Using SQL

By using file\_fdw of an additional module, audit logs can be accessed using SQL. This section describes how to view audit logs, using Session Audit Logging output to a dedicated log file as an example.

1. Install file\_fdw

Execute CREATE EXTENSION to install file\_fdw as an extension.

```
$ psql
=# CREATE EXTENSION file_fdw;
=# \dx
                 List of installed extensions
         | Version | Schema |
Name
                                    Description
file_fdw | 1.0
                  public
                               | foreign-data wrapper for flat file access
        1.0
                  public
pgaudit
                              provides auditing functionality
        1.0
                  | pg_catalog | PL/pgSQL procedural language
plpgsql
(3 rows)
```

2. Create an external server

Use CREATE SERVER to create an external server managed by file\_fdw.

```
$ psql
=# CREATE SERVER auditlog FOREIGN DATA WRAPPER file_fdw;
```

3. Create an audit log table.

Use CREATE FOREIGN TABLE to define the table columns of audit logs, CSV file name and format.

```
=# CREATE FOREIGN TABLE auditlog (
header text,
class text,
sql_start_time timestamp with time zone,
remote_host_name text,
backend_process_id text,
application_name text,
session_user_name text,
database_name text,
virtual_transaction_id text,
statement_id text,
substatement_id text,
command_tag text,
sqlstate text,
object_type text,
object_name text,
error_message text,
sql text,
parameter text
) SERVER auditlog
OPTIONS (filename '/database/inst1/pgaudit_log/pgaudit-2017-03-12.log', format
'csv' );
```



If an audit log file is rotated and multiple audit log files exist, it is necessary to create a table for each audit log file.

4. View audit logs

Use SELECT and view the audit logs.

# 6.9 Removing Setup

This section describes how to remove the setup of pgaudit.

- 1. Start the instance
- 2. Remove the pgaudit extension

Execute DROP EXTENSION to remove the pgaudit extension from the database.

```
$ psql -d <database name>
=# DROP EXTENSION pgaudit;
=# \q
```

3. Change postgresql.conf

Remove the parameters below relating to pgaudit.

- shared\_preload\_libraries
- pgaudit.config\_file
- 4. Restart the instance
- 5. Remove the pgaudit files

As superuser, run the following command. Note that "<x>" in paths indicates the product version.

\$ su Password:\*\*\*\*\*
# rm -rf /opt/fsepv<x>server64/filesCopiedDuringSetup



The files copied during setup can be checked below.

# find /opt/fsepv<x>server64/OSS/pgaudit

Open a command prompt as administrator privileges, run the following command. Note that "<x>" in paths indicates the product version.

> del "c:\Program Files\Fujitsu\fsepv<x>server64\filesCopiedDuringSetup"

**1** Information

The files copied during setup can be checked below.

> dir /b /s "c:\Program Files\Fujitsu\fsepv<x>server64\OSS\pgaudit"

4

L

W



# FUJITSU Enterprise Postgres 12



# Cluster Operation Guide

Database Multiplexing (Linux) > Database Multiplexing (Windows) > PRIMECLUSTER >



# FUJITSU Enterprise Postgres 12



# Cluster Operation Guide Database Multiplexing (Linux)

# **Preface**

#### Purpose of this document

This document describes the tasks required for using the database multiplexing feature of FUJITSU Enterprise Postgres.

#### Intended readers

This document is intended for those who set up and use the database multiplexing feature.

Readers of this document are also assumed to have general knowledge of:

- PostgreSQL
- SQL
- Linux

#### Structure of this document

This document is structured as follows:

#### Chapter 1 Overview of Database Multiplexing Mode

Provides an overview of database multiplexing mode.

#### Chapter 2 Setting Up Database Multiplexing Mode

Describes how to set up database multiplexing mode.

# Chapter 3 Operations in Database Multiplexing Mode

Explains periodic database multiplexing mode.

#### Chapter 4 Action Required when an Error Occurs in Database Multiplexing Mode

Explains the action required when an error occurs during a database multiplexing mode.

#### Chapter 5 Managing Mirroring Controller Using WebAdmin

Explains how to set up and manage Mirroring Controller in a streaming replication cluster using WebAdmin.

#### Appendix A Parameters

Explains the configuration files and parameters required for database multiplexing mode.

# Appendix B Supplementary Information on Building the Primary Server and Standby Server on the Same Server

Explains supplementary information on building the primary server and standby server on the same server.

## Appendix C User Commands

Explains the user commands.

# Appendix D Notes on Performing Automatic Degradation Immediately after a Heartbeat Abnormality

Provides notes when performing automatic degradation unconditionally after a heartbeat abnormality is detected during heartbeat monitoring of an operating system or server.

#### Appendix E Supplementary Procedure on Configuring for Operation in Database Multiplexing Mode

Explains supplementary procedure on the configuration required for operation in database multiplexing mode.

## Appendix F WebAdmin Disallow User Inputs Containing Hazardous Characters

Explains characters not allowed in WebAdmin.

#### Appendix G Collecting Failure Investigation Data

Explains how to collect data for initial investigation.

# **Export restrictions**

Exportation/release of this document may require necessary procedures in accordance with the regulations of your resident country and/or US export control laws.

# Issue date and version

Edition 1.0: July 2020

# Copyright

Copyright 2016-2020 FUJITSU LIMITED

# Contents

Chapter 1 Overview of Database Multiplexing Mode	′
1.1 What is Database Multiplexing Mode	1
1.1.1 Monitoring Using Database Multiplexing Mode	2
1.1.2 Referencing on the Standby Server	<i>6</i>
1.1.2.1 If Prioritizing the Main Job on the Primary Server	<i>6</i>
1.1.2.2 If Performing the Referencing Job on the Synchronous Standby Server	<del>(</del>
1.2 System Configuration for Database Multiplexing Mode	e
1.2.1 Mirroring Controller Resources.	8
1.2.1.1 Database Server Resources	8
1.2.1.2 Arbitration Server Resources	9
1.2.2 Mirroring Controller Processes	9
1.2.2.1 Database Server Processes	9
1.2.2.2 Arbitration Server Process.	9
1.2.3 Redundancy of the Admin and Log Transfer Networks	
1.2.4 Notes on CPU Architecture and Products	10
1.3 Deciding on Operation when a Heartbeat Abnormality is Detected	10
1.4 Security in Database Multiplexing	
1.4.1 Authentication of the Standby Server	13
1.4.2 Encryption of Transaction Logs Transferred to the Standby Server	
Chapter 2 Setting Up Database Multiplexing Mode	
2.1 Installation	
2.2 Preparing for Setup	
2.2.1 Preparing the Database Server	
2.2.1.1 Preparing the Backup Disk	
2.2.2 Preparing the Arbitration Server	
2.2.2.1 Preparing to Output Error Logs to the Event Log (Windows)	
2.2.2.2 Security Policy Settings (Windows)	
2.3 Setting Up the Arbitration Server	
2.3.1 Configuring the Arbitration Server.	
2.3.2 Creating a User Exit for the Arbitration Server.	
2.3.3 Starting the Mirroring Controller Arbitration Process	
2.4 Setting Up the Primary Server.	
2.4.1 Setting Up Database Multiplexing Mode on the Primary Server	
2.4.2 Creating, Setting, and Registering the Primary Server Instance.	
2.4.3 Starting Mirroring Controller on the Primary Server	
2.5 Setting Up the Standby Server	
2.5.1 Setting Up Database Multiplexing Mode on the Standby Server	
2.5.2 Creating, Setting, and Registering the Standby Server Instance	
2.5.3 Starting Mirroring Controller on the Standby Server	
2.6 Creating a User Exit for a Database Server.	
2.7 Confirming the Streaming Replication Status.	
2.8 Checking the Connection Status.	
2.8.1 Checking the Connection Status on a Database Server.	
2.8.2 Checking the Connection Status on the Arbitration Server	
2.9 Creating Applications.	
2.9.1 Application Connection Server Settings.	
2.10 Checking the Behavior	
2.11 Tuning.	
2.11.1 Tuning to Stabilize the Database Multiplexing Mode	
2.11.2 Tuning to Stabilize Queries on the Standby Server.	
2.11.3 Tuning to Stabilize Queries on the Standby Server (when Performing Frequent Updates on the Primary Server)	
2.11.4 Tuning for Optimization of Degradation Using Abnormality Monitoring.	
2.11.4.1 Tuning for Abnormality Monitoring of the Operating System or Server.	

2.11.4.1.2 Tuning Abnormality Monitoring for Operations that Perform Automatic Degeneration by Calling a Use	
Determines Degeneration	
2.11.4.1.3 Tuning Abnormality Monitoring for Operations that Notify Messages	
2.11.4.1.4 Tuning Abnormality Monitoring for Operations that Perform Automatic Degenerate Unconditionally due	
Abnormality	
2.11.4.2 Tuning for Abnormality Monitoring of Darabase Processes	
2.11.4.3 Tuning for Abnormality Monitoring of Streaming Replication	
2.11.4.4 Tuning for Disk Abnormality Monitoring	
2.12 Setting Automatic Start and Stop of Mirroring Controller and Multiplexed Instances	
2.13 Setting Automatic Start and Stop of the Mirroring Controller Arbitration Process	
2.14 Backup Operation	
2.14.1 Backing up Database Multiplexing Mode Information	
2.14.2 Database Backup Operation	59
Chapter 3 Operations in Database Multiplexing Mode	60
3.1 Starting and Stopping the Mirroring Controller Arbitration Process	
3.1.1 Starting the Mirroring Controller Arbitration Process	
3.1.2 Stopping the Mirroring Controller Arbitration Process.	
3.2 Starting and Stopping Mirroring Controller.	
3.3 Checking the Database Multiplexing Mode Status	
3.3.1 Checking the Status of the Database Server	
3.3.2 Checking the Status of the Arbitration Server.	
3.4 Manually Switching the Primary Server	
3.5 Manually Disconnecting the Standby Server	
3.6 Action Required when a Heartbeat Abnormality is Detected	
3.7 Monitoring Mirroring Controller Messages.	66
3.8 Server Maintenance	68
3.8.1 Rolling Updates	69
3.8.2 Stopping for Maintenance	74
3.8.3 Arbitration Server Maintenance	75
3.9 Changes in Operation	76
3.9.1 Changes Required when the Standby Server is Stopped	
3.9.2 Changing from Single Server Mode to Database Multiplexing Mode	
3.9.3 Changing from Database Multiplexing Mode to Single Server Mode	
3.9.4 Changing to Database Multiplexing Mode when the Arbitration Server is Used for Automatic Degradation	
3.9.5 Changing Parameters	
3.9.6 Uninstalling in Database Multiplexing Mode	
Chapter 4 Action Required when an Error Occurs in Database Multiplexing Mode	
4.1 Action Required when Server Degradation Occurs	
4.1.1 Operations when the Server has Started Degrading after a Switch has Occurred	
4.1.1.1 Identify Cause of Error and Restore the Standby Server	
4.1.1.1 Stop Mirroring Controller.	
4.1.1.1.2 Recovery of the Mirroring Controller management directory	
4.1.1.1.3 Identify cause of error and perform recovery	
4.1.1.2 Rebuild the Standby Server.	
4.1.1.3 Failback of the Primary Server	
4.1.2 Operations when the Server has Started Degrading after a Disconnection has Occurred	
4.1.2.1 Identify Cause of Error and Restore the Standby Server	
4.1.2.1.1 Stop Mirroring Controller	
4.1.2.1.2 Recovery of the Mirroring Controller management directory	
4.1.2.1.3 Identify cause of error and perform recovery	
4.1.2.2 Rebuild the Standby Server	
4.1.3 Addressing Errors During Degrading Operation	
4.2 Action Required when Automatic Switch Fails	93
4.3 Action Required when Automatic Disconnection Fails	
4.4 Action Required when All Database Servers or Instances Stopped	
4.5 Recovering from an Incorrect User Operation	98

Chapter 5 Managing Mirroring Controller Using WebAdmin	100
5.1 Mirroring Controller Setup	101
5.2 Edit Mirroring Controller Setup	103
5.3 Mirroring Controller Configuration	104
5.4 Stopping Mirroring Controller	105
5.5 Starting Mirroring Controller	106
5.6 Disabling Failover Mode	106
5.7 Enabling Failover Mode	106
5.8 Deleting Mirroring Controller Setup	106
5.9 Status Update after Failover	107
5.10 Action Required when an Error Occurs in the Combined Admin Network and Log Transfer Network	107
5.11 Performing Automatic Degradation Using the Arbitration Server	108
Appendix A Parameters	110
A.1 Parameters Set on the Primary Server	110
A.2 Parameters Set on the Standby Server	112
A.3 Network Configuration File.	114
A.4 Server Configuration File	117
A.4.1 Server Configuration File for the Database Servers	117
A.4.2 Arbitration Configuration File	125
Appendix B Supplementary Information on Building the Primary Server and Standby Server on the Same Server	130
B.1 Backup Data Storage Destination Directory	130
B.2 How to Execute the mc_ctl Command	130
Appendix C User Commands	132
C.1 Fencing Command	132
C.2 Arbitration Command.	133
C.3 State Transition Commands	134
C.3.1 Post-switch Command.	134
C.3.2 Pre-detach Command	135
C.3.3 Post-attach Command	135
Appendix D Notes on Performing Automatic Degradation Immediately after a Heartbeat Abnormality	137
Appendix E Supplementary Procedure on Configuring for Operation in Database Multiplexing Mode	140
E.1 Security Policy Settings	140
E.2 Windows Firewall Settings	140
Appendix F WebAdmin Disallow User Inputs Containing Hazardous Characters	142
Appendix G Collecting Failure Investigation Data.	143
Index	144

# Chapter 1 Overview of Database Multiplexing Mode

This chapter provides an overview of database multiplexing mode.



In this and subsequent chapters, the word "Mirroring Controller" may be used in the process or management directory name or explanation.

# 1.1 What is Database Multiplexing Mode

Database multiplexing mode is an operation mode (log shipping mode) based on PostgreSQL streaming replication. Other software such as cluster software is not required.

This mode replicates the database on all servers that comprise the cluster system. It achieves this by transferring the updated transaction logs of the database from the server that receives the updates (primary server) to another server (standby server), and then reflecting them on the standby server. The client driver automatically distinguishes between the primary and standby servers, so applications can be connected transparently regardless of the physical server.

It consists of a feature that detects faults in the elements that are essential for the continuity of the database operation (such as the database process, disk, and network), as well as simplified switchover and standby server disconnection features. Furthermore, referencing can be performed on the standby server. The database will be copied in synchronous mode.



If using WebAdmin or Mirroring Controller, FUJITSU Enterprise Postgres supports cluster systems comprising one primary server and one standby server.

- Although it is possible to connect an asynchronous standby server to the cluster system as an additional server, the standby server is not targeted for monitoring by Mirroring Controller.
- A synchronous standby server cannot be connected to the cluster system as an additional server.



The streaming replication feature is not described in this manual.

Refer to "High Availability, Load Balancing, and Replication" in the PostgreSQL Documentation for information on the streaming replication feature.

Figure 1.1 Failover from the primary server to the standby server

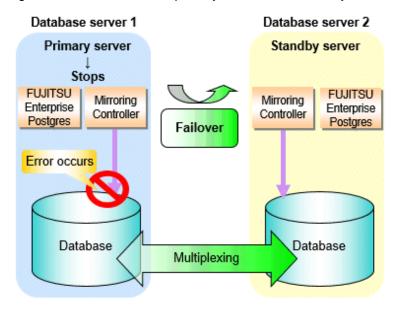
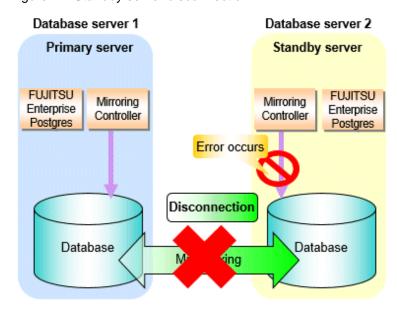


Figure 1.2 Standby server disconnection



# Database degradation using the arbitration server

FUJITSU Enterprise Postgres provides the Server Assistant that objectively determines the status of database servers as a third party, and if necessary, isolates affected databases if the database servers are unable to accurately ascertain their mutual statuses in database multiplexing mode, such as due to a network error between database servers, or server instability. Database degradation can be performed by using the server (arbitration server) on which the Server Assistant is installed.

For database degradation using the arbitration server, if the database servers are unable to check their mutual statuses (due to a network error between database servers or server instability), then the database server queries the arbitration server for the status of the other database server. If it is determined based on the heartbeat result that the status is unstable, the applicable database server will be isolated from the cluster system (fencing). The arbitration server periodically heartbeats the database server so that it can respond immediately to queries from the database server. The fencing process can be customized according to the environment where Mirroring Controller is used.

Additionally, the database servers are always performing their heartbeats for the arbitration server so that it can perform check requests any time.

Database server 1 Database server 2 Primary server Standby server Abnormality Network issue between servers, detection or one server is unresponsive **FUJITSU FUJITSU** Mirroring Enterprise Enterprise Controller Postgres Postgres Failover Database Database Check status Arbitration server Fencing FUJITSU Enterprise Postgres Server Assistant

Figure 1.3 Database degradation using the arbitration server

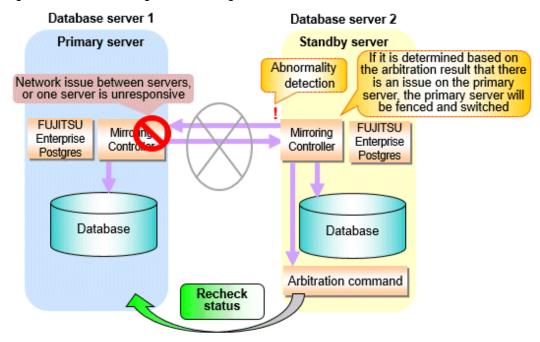


Install the arbitration server on a different physical server to that of the database server. Refer to "1.2 System Configuration for Database Multiplexing Mode" for information on the system configuration when using the arbitration server.

# Database degradation using the arbitration command

The arbitration command is a user command that performs arbitration processing in lieu of the arbitration server. If an arbitration server cannot be deployed, arbitration of the database server can be performed using the arbitration command.

Figure 1.4 Database degradation using the arbitration command





Refer to "2.6 Creating a User Exit for a Database Server" or "Appendix C User Commands" for information on user commands.

# 1.1.1 Monitoring Using Database Multiplexing Mode

In database multiplexing mode, perform the monitoring below.

- Operating system or server failures, and no-response state

By generating a heartbeat between Mirroring Controller on each server, operating system or server errors are detected and acknowledged between the relevant servers.

The optimal operating method for environments where database multiplexing mode is performed can be selected from the following:

- Use the arbitration server to perform automatic degradation (switch/disconnect)

This is the default method.

The arbitration server objectively determines the status of database servers, then isolates and degrades from the cluster system the ones with an unstable status.

Refer to "Database degradation using the arbitration server" for details.

- Call the user exit (user command) that will perform the degradation decision, and perform automatic degradation

If the arbitration server cannot be installed, select if arbitration processing can be performed by the user instead.

Mirroring Controller queries the user exit on whether to degrade. The user exit determines the status of the database server, and notifies Mirroring Controller whether to perform degradation.

Refer to "Database degradation using the arbitration command" for details.

- Notification messages

Use this method if using a two-database server configuration.

Mirroring Controller outputs messages to the system log when an abnormality is detected. This ensures that a split brain will not occur due to a heartbeat abnormality - however, automatic switching will not be performed if the primary server operating system or server fails or becomes unresponsive.

- Perform automatic degradation unconditionally after a heartbeat abnormality

This method is handled as in FUJITSU Enterprise Postgres 9.6 or earlier versions.

This method is not recommended, because Mirroring Controller unconditionally will perform automatic degradation after heartbeat abnormalities.

- Database process failures, and no-response state

Mirroring Controller periodically accesses the database processes and checks the status. A process error is detected by monitoring whether an access timeout occurs.

- Disk failure

Mirroring Controller periodically creates files on the data storage destination disk below. A disk error is detected when an I/O error occurs.

- Data storage destination disk
- Transaction log storage destination disk
- Tablespace storage destination disk

Failures that can be detected are those that physically affect the entire system, such as disk header or device power failures.

- Streaming replication issue

Mirroring Controller detects streaming replication issues (log transfer network and WAL send/receive processes) by periodically accessing the PostgreSQL system views.

- Mirroring Controller process failure and no response

In order to continue the monitoring process on Mirroring Controller, Mirroring Controller process failures and no responses are also monitored.

The Mirroring Controller monitoring process detects Mirroring Controller process failures and no responses by periodically querying the Mirroring Controller process. If an issue is detected, Mirroring Controller is automatically restarted by the Mirroring Controller monitoring process.



- If output of messages is selected as the operation to be performed when a heartbeat abnormality is detected during heartbeat monitoring of the operating system or server, automatic degradation will not be performed.
  - However, if an issue in the WAL send process is detected on the primary server, then the standby server will be disconnected, and as a result an automatic disconnection may be performed even if the standby server operating system or server fails or becomes unresponsive.
- You can select in the parameters if the primary server will be switched if a database process is unresponsive or if tablespace storage destination disk failure is detected. Refer to "Appendix A Parameters" for details.
- If the standby server was disconnected, Mirroring Controller will automatically comment out the synchronous\_standby\_names parameter in the postgresql.conf file of the primary server. Accordingly, you can prevent the application processing for the primary server being stopped.



If the role of primary server was switched to another server and then starts degrading, the original primary server will not become the standby server automatically. Remove the cause of the error, and then change the role of the original primary server to the server currently acting as standby server. Refer to "4.1 Action Required when Server Degradation Occurs" for details.

# 1.1.2 Referencing on the Standby Server

# 1.1.2.1 If Prioritizing the Main Job on the Primary Server

If a reference job is performed on the standby server and the primary server is switched, this may impact the main job from the point of view of load and conflict. This is because, on the new primary server (that is, the original standby server), both the main job that was being executed on the original primary server and the reference job that was being continued on the original standby server will be processed.

Therefore, to degrade the reference job (so that the impact on the main job is reduced), you can select the user exit below to disconnect the reference job that was performed on the original standby server.

- Post-switch command



If continuing with the referencing job after switching the primary server, give careful consideration to the server resource estimates, and the likely impact on performance.

# 1.1.2.2 If Performing the Referencing Job on the Synchronous Standby Server

If an issue such as a log transfer network failure obstructs the continuation of a job on the primary server, the standby server may be automatically disconnected from the cluster system.

Therefore, if operating the reference job on the assumption that the connection destination is the synchronous standby server, you can select to temporarily stop the job by using the user exit or the feature below, so that unexpected referencing of past data does not occur as a result of the disconnection.

- Pre-detach command
- Forced stoppage of the standby server instance on disconnection (specify in the parameter of the server configuration file)

Additionally, if the standby server is incorporated into the cluster system, reference jobs can be started or resumed by using the user exit below.

- Post-attach command



#### See

- Refer to "2.6 Creating a User Exit for a Database Server" or "Appendix C User Commands" for information on each user exit.
- Refer to "A.4.1 Server Configuration File for the Database Servers" for information on the server configuration file of the database server.



Mirroring Controller will continue processing regardless of the processing result of the above user exits and features.

# 1.2 System Configuration for Database Multiplexing Mode

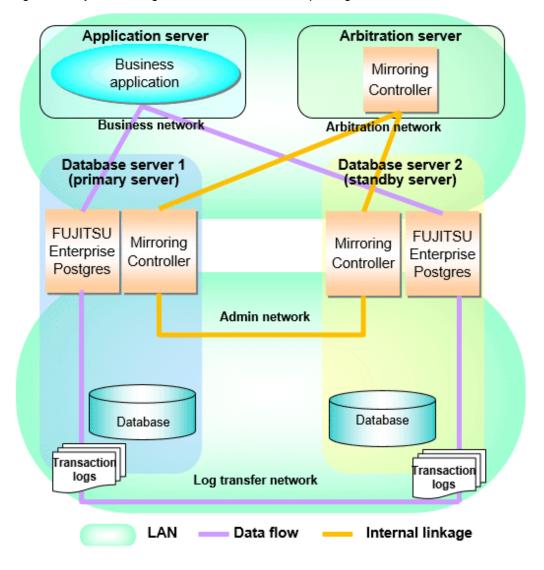
This section explains the products, features, and networks that are part of a database multiplexing system.

The following table shows the network types uses by database multiplexing systems.

Network type	Description
Job network	Network between the application that accesses the database, and the database
	server.

Network type	Description
Arbitration network	Network used by the arbitration server to check the status of the primary server and standby server, and communicate with Mirroring Controller of the database servers. Additionally, if the job network is disconnected from outside, it can also be used as the arbitration network. Refer to "1.4 Security in Database Multiplexing" for details on network security.
Admin network	Network used by the primary server and the standby server to monitor each other using Mirroring Controller, and to control Mirroring Controller of other servers.
Log transfer network	Network used to transfer the transaction logs of the database, which is part of database multiplexing.

Figure 1.5 System configuration for database multiplexing mode



The arbitration server is installed to check the database server status as a third party, and to perform fencing. Therefore, to obtain the intended benefits, consider the following.

- Install the arbitration server on a different server to that of the database server.
- For the arbitration network, use a network that will not be impacted by line faults or the load on the admin network or log transfer network. This is necessary to correctly determine issues on the admin network or log transfer network.



- The arbitration server can also be used as an application server. However, consider the server load.
- It is recommended to link the arbitration server with other cluster systems, in order to provide redundancy.
- Use the arbitration server in combination with the same version of FUJITSU Enterprise Postgres as that of the primary server and standby server.
- The arbitration server can be built on a different platform to that of the database server.



Because the ping command of the operating system is used for heartbeat monitoring of the database server, configure the network so that ICMP can be used on the admin network and the arbitration network.

# 1.2.1 Mirroring Controller Resources

This section describes the database server and arbitration server resources of Mirroring Controller.

#### 1.2.1.1 Database Server Resources

The only Mirroring Controller resource is the Mirroring Controller management directory, which stores the files that define the Mirroring Controller behavior, and the temporary files that are created when Mirroring Controller is active.

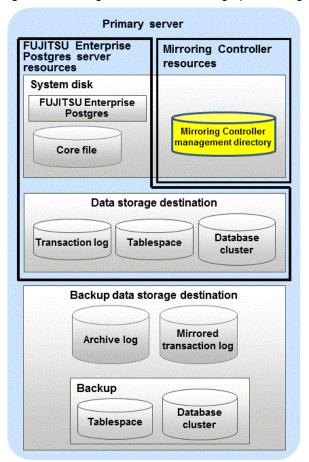


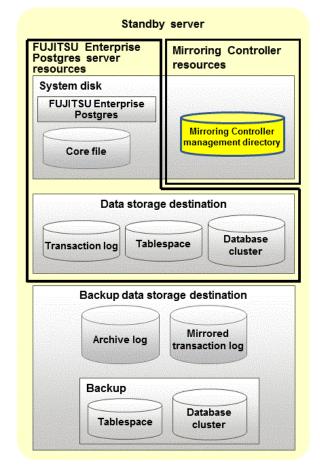
- Do not create the Mirroring Controller management directory in a directory managed by FUJITSU Enterprise Postgres, otherwise it may be deleted by mistake or may cause unexpected problems when FUJITSU Enterprise Postgres recovery is performed (such as old version of files being restored).

- Refer to "Preparing Directories for Resource Deployment" in the Installation and Setup Guide for Server for information on the directories managed by FUJITSU Enterprise Postgres.
- The backup methods described in "Backing Up the Database" in the Operation Guide cannot be used to back up the Mirroring Controller resources. Therefore, users must obtain their own backup of Mirroring Controller resources, in addition to FUJITSU Enterprise Postgres server resources. Retrieve backups after stopping Mirroring Controller.
- If the automatic switch/disconnection is enabled, do not edit synchronous\_standby\_names for the Mirroring Controller monitoring target instance. Otherwise, if Mirroring Controller is switched after editing, data may be lost or SQL access may stop.
- If you are building on a virtual machine or cloud, make sure the virtual machines are on different physical servers. Refer to your virtual machine software and cloud vendor documentation for instructions on how to deploy virtual machines.

The content on the primary server will be backed up. You cannot tell which server is the primary server to be backed up, because switching and failback may be performed between the servers. It is also impossible to tell which server is to be restored using the backed up data. Accordingly, ensure that you create a backup of each server when it is working as the primary server.

Figure 1.6 Configuration when backing up Mirroring Controller resources





# 1.2.1.2 Arbitration Server Resources

The only arbitration server resource is the Mirroring Controller arbitration process management directory. This directory stores the files that define the Mirroring Controller arbitration process behavior and the temporary files created when Mirroring Controller is active.

# 1.2.2 Mirroring Controller Processes

This section describes the database server and arbitration server processes of Mirroring Controller.

# 1.2.2.1 Database Server Processes

The database server processes comprise the Mirroring Controller process and Mirroring Controller monitoring process.

Process type	Description
Mirroring Controller process	Performs operating system/server and process heartbeat monitoring and disk abnormality monitoring between database servers. Additionally, it issues arbitration requests to the arbitration server.
Mirroring Controller monitoring process	Performs heartbeat monitoring of the Mirroring Controller process. If the Mirroring Controller process returns no response or is down, the monitoring process is restarted automatically.

# 1.2.2.2 Arbitration Server Process

The only arbitration process is the Mirroring Controller arbitration process.

Process type	Description
Mirroring Controller arbitration process	Performs rechecks for issues detected on the primary server or the standby server.  Additionally, this process performs fencing if it determines that there is an issue on the primary server or the standby server.

# 1.2.3 Redundancy of the Admin and Log Transfer Networks

The admin network is an important one, because it is used by Mirroring Controller to check the status of each server.

Additionally, the log transfer network is an important one, because it is necessary to ensure data freshness.

Accordingly, configure a failure-resistant network by implementing network redundancy via channel bonding provided by the operating system or network driver vendor.

# 1.2.4 Notes on CPU Architecture and Products

Use the same CPU architecture (endian) for the primary server, standby server, and the arbitration server.

A server using only PostgreSQL streaming replication cannot be specified as the database multiplexing system log transfer destination.

# 1.3 Deciding on Operation when a Heartbeat Abnormality is Detected

The operation to be performed when a heartbeat abnormality is detected using operating system/server heartbeat monitoring is decided on according to the environment where database multiplexing mode is performed or the operating method.

It is possible to select from the four operations below, and specify this in the parameters of Mirroring Controller:

- Use the arbitration server to perform automatic degradation (switch/disconnect)
- Call the user exit (user command) that will perform the degradation decision, and perform automatic degradation
- Notification messages
- Perform automatic degradation unconditionally (switch/disconnect)

The table below shows if jobs can be continued on the primary server when an issue is detected during heartbeat monitoring of the operating system/server.

# Continuation of jobs on the primary server when an issue is detected during heartbeat monitoring of the operating system/server

	Abnormal event					
Operation	Server/operating system failures or no responses		Admin network	Log transfer	Issue on a network for both admin and	
	Primary server	Standby server	issue	Hetwork issue	log transfer	
Automatic degradation using the arbitration server	Y (switch)	Y (disconnect)	Y	Y (disconnect)	Y (disconnect)	
Call a user exit and perform automatic degradation	Y (switch)	Y (disconnect)	Y	Y (disconnect)	Y (disconnect)	
Notification messages	N (message notification only)	N (message notification only)	Y	Y (disconnect)	Y (disconnect)	
Unconditional automatic degradation	Y (switch)	Y (disconnect)	N (split brain occurs)	Y (disconnect)	N (split brain occurs)	

# 1.4 Security in Database Multiplexing

The database server replicates the database on all servers that comprise the cluster system. It achieves this by transferring and reflecting the updated transaction logs of the database from the primary server to the standby server.

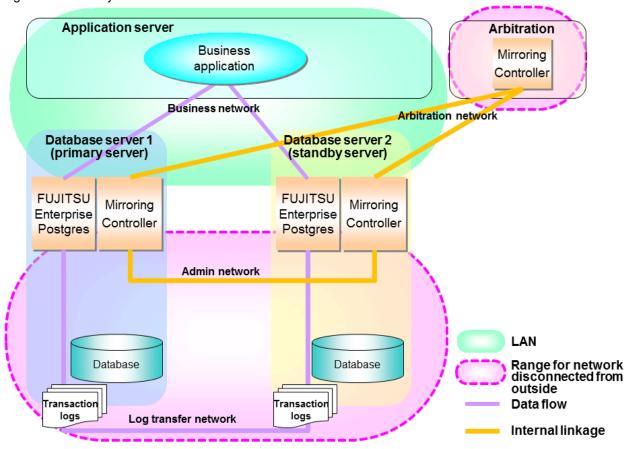
To safeguard the database against unauthorized access and preserve data confidentiality in transaction log transfers, carefully consider security and take note of the following when performing database multiplexing:

- Do not use trust authentication when using replication connection.
- Configure the admin network and the log transfer network so that they cannot be connected from the outside, as shown in Figure 1.7 Security.

Additionally, for the line on which Mirroring Controller connects from the database server to the arbitration server, take note of the following points and consider security carefully.

- Build a network with the arbitration server disconnected from outside, as shown in Figure 1.7 Security.

Figure 1.7 Security



However, it may not always be possible to adopt the configuration mentioned above. For example, you may want to place the servers in a nearby/neighboring office to minimize network delays.

In this case, combine the following features to enhance security:

- Authentication of the Standby Server
- Encryption of Transaction Logs Transferred to the Standby Server

When these features are combined, security will be achieved as shown below.

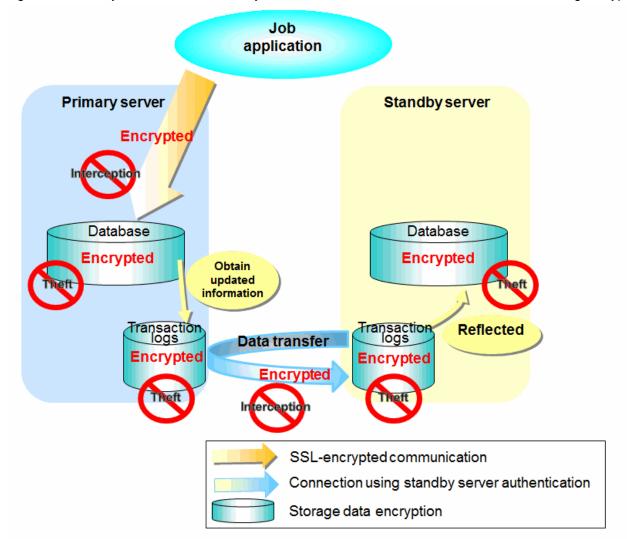


If the job network is disconnected from outside, it can be used as the arbitration network. However, if a network is to be used as both a job network and arbitration network, consider the load on the network.



If a port is blocked (access permission has not been granted) by a firewall, etc., enable use of the target port by granting access. Refer to the vendor document for information on how to open (grant access permission to) a port. Consider the security risks carefully when opening ports.

Figure 1.8 Security achieved when standby server authentication is combined with transaction log encryption





Refer to "Performing Database Multiplexing" under "Configuring Secure Communication Using Secure Sockets Layer" in the Operation Guide for information on encrypting SSL communications.

# 1.4.1 Authentication of the Standby Server

You can prevent spoofing connections from an external server purporting to be the standby server by using authentication with a user name and password.

Configure the setting in the primary server pg\_hba.conf file so that authentication is performed for connections from the standby server in the same way as for connections from the client.



Refer to "Client Authentication" in the PostgreSQL Documentation for information on content that can be configured in pg\_hba.conf.

# 1.4.2 Encryption of Transaction Logs Transferred to the Standby Server

In case the authentication of the standby server is breached so that a malicious user purporting to be the standby server can spoof data, the transaction log data can be encrypted to prevent it from being deciphered. The transparent data encryption feature is used to encrypt the data.

.....



Refer to "Protecting Storage Data Using Transparent Data Encryption" in the Operation Guide for details.

# Chapter 2 Setting Up Database Multiplexing Mode

This chapter describes how to set up database multiplexing mode, and how to check it.

# Users who perform setup and operations on the database server

Setup and operations of the database server must be performed by the instance administrator user.

# Users who perform setup and operations on the arbitration server

The following users may perform setup and operations on the arbitration server when it is used for automatic degradation.

Linux

Any operating system user.

Windows

Any user with administrator privileges. This user must be assigned the "Log on as a service" user right.



- Mirroring Controller selects a database superuser as the user who will connect to the database instance. This enables instance administrator users and database superusers who operate the Mirroring Controller commands to run database multiplexing mode in different environments.
- The application name for connecting to the database instance is "mc\_agent".

# Matching the system times

Before starting the setup, ensure that the times in the primary server, standby server and arbitration server match, by using the operating system time synchronization feature, for example.

The tolerated difference is approximately one second.

If the system times are not synchronized (because the tolerated difference is exceeded, for example), problem investigation may be affected.

#### Configuring ICMP

Because the ping command of the operating system is used for heartbeat monitoring of the database server, configure the network so that ICMP can be used on the admin network and the arbitration network. Refer to the relevant operating system procedure for details.

# Setup

The setup procedure is shown in the table below. However, the procedure on the arbitration server should be performed only when the arbitration server is used for automatic degradation. A distinction is made between the procedures on the primary server and standby server according to whether the arbitration server is used.

Ston		Refer to		
Step	Primary server	Standby server	Arbitration server	Kelei to
1	Installation			2.1 Installation
2	Preparing the database server		Preparing the arbitration server	2.2 Preparing for Setup
3			Configuring the arbitration server	2.3.1 Configuring the Arbitration Server
4			Creating a user exit	2.3.2 Creating a User Exit for the Arbitration Server

Cton		Defer to		
Step	Primary server	Standby server	Arbitration server	Refer to
5			Starting the arbitration process	2.3.3 Starting the Mirroring Controller Arbitration Process
6	Setting up database multiplexing mode			2.4.1 Setting Up Database Multiplexing Mode on the Primary Server
7	Creating, setting, and registering the instance			2.4.2 Creating, Setting, and Registering the Primary Server Instance
8	Creating a user exit			2.6 Creating a User Exit for a Database Server
9	Starting Mirroring Controller			2.4.3 Starting Mirroring Controller on the Primary Server
10		Setting up database multiplexing mode		2.5.1 Setting Up Database Multiplexing Mode on the Standby Server
11		Creating, setting, and registering the instance		2.5.2 Creating, Setting, and Registering the Standby Server Instance
12		Creating a user exit		2.6 Creating a User Exit for a Database Server
13		Starting Mirroring Controller		2.5.3 Starting Mirroring Controller on the Standby Server
14	Confirming the streaming replication status			2.7 Confirming the Streaming Replication Status
15	Checking the connection status			2.8.1 Checking the Connection Status on a Database Server
16		Checking the connection status		2.8.1 Checking the Connection Status on a Database Server
17			Checking the connection status	2.8.2 Checking the Connection Status on the Arbitration Server
18	Creating applications			2.9 Creating Applications
19	Checking the behavior			2.10 Checking the Behavior

Explanations for each step are provided below.



- The setup procedure is also the same when changing the mode on a single server to database multiplexing mode. In this case, omit the installation of FUJITSU Enterprise Postgres and the creation of the instance.

Refer to "3.9.2 Changing from Single Server Mode to Database Multiplexing Mode" for details.

- The primary and standby server can be pseudo-configured on the same server for system testing, for example. In this case, the setup can be performed using the same procedure, however there will be some supplementary steps.

Before performing the setup, refer to "Appendix B Supplementary Information on Building the Primary Server and Standby Server on the Same Server".

# 2.1 Installation

Refer to the manuals below, and then install the product.



- Refer to the Installation and Setup Guide for Server for details on how to install FUJITSU Enterprise Postgres.
- Refer to the Installation and Setup Guide for Server Assistant for information on installing the Server Assistant on the arbitration server.



Do not use the arbitration server also as a database server. The arbitration server is installed to check the database server status as a third party, and to perform fencing. Using the arbitration server also as a database server nullifies the effectiveness of the arbitration server.

# 2.2 Preparing for Setup

This section describes the preparation required before setting up Mirroring Controller.

# 2.2.1 Preparing the Database Server

# 2.2.1.1 Preparing the Backup Disk

In Mirroring Controller, by performing a backup, recovery is possible even if all server disks are corrupted.

The content on the primary server should be backed up. However, through switching and failback, the standby server may also become the primary server. Accordingly, prepare each of the backup disk devices for the primary and standby servers. Perform backup on the primary server used at the time of the backup.

# 2.2.2 Preparing the Arbitration Server

# 2.2.2.1 Preparing to Output Error Logs to the Event Log (Windows)

This section explains the preparatory tasks for outputting error logs to the event log.



If no event source name is registered, messages output to the event log may be incomplete.

# Configuring each server

Event logs for the Mirroring Controller commands on the arbitration server may be output with the default event source name "MirroringControllerArbiter". Therefore, register this default event source name beforehand.

Example)

The following is an example in which the DLL of a 64-bit product is registered under the default event source name. Note that "<x>" indicates the product version.

regsvr32 "c:\Program Files\Fujitsu\fsepv<x>assistant64\lib\mcarbevent.dll"

# Setting each instance

You can output messages to any event source named by the user, so that messages output to the event log can be identified by each instance.

Example)

The following is an example in which the DLL of a 64-bit product is registered under the event source name "Mirroring Controller arbiterl". Note that "<x>" indicates the product version.

> regsvr32 /n /i:"Mirroring Controller arbiter1" "c:\Program Files\Fujitsu\fsepv<x>assistant64\lib
\mcarbevent.dl1"

The parameter must be edited for each instance. Refer to "A.4.2 Arbitration Configuration File" and set the event\_source parameter.

# If installing multiple versions

If FUJITSU Enterprise Postgres is already installed on the same machine, search for the key below in Registry Editor, and make a note of the path of the registered DLL. Afterwards, register a new DLL using the default event source name.

Use the DLL path that you made a note of in the above step when re-registering the default event source name during uninstallation.

MirroringControllerArbiter

# 2.2.2.2 Security Policy Settings (Windows)

On the arbitration server, operating system user accounts that operate the Mirroring Controller arbitration process must be assigned the "Log on as a service" user right in order to use Windows Services to start and stop the Mirroring Controller arbitration process.

If the security settings to enable this have not been configured, refer to "E.1 Security Policy Settings" and configure the settings.

# 2.3 Setting Up the Arbitration Server

This section explains how to set up the arbitration server.

# 2.3.1 Configuring the Arbitration Server

This section explains how to set up database multiplexing mode on the arbitration server.

In database multiplexing mode, the files that are required for operations are managed in the Mirroring Controller arbitration process management directory.

There is one Mirroring Controller arbitration process management directory for each arbitration process.



The arbitration process for each database multiplexing system can be started on a single arbitration server.



- Refer to the Reference for information on the mc\_arb command.
- Refer to "Appendix A Parameters" for information on the parameters to be edited for the setup.

Perform the following procedure:

Linux

- 1. On the arbitration server, log in as any operating system user who starts and stops the arbitration process.
- 2. Configure the environment variables.

Set the following environment variables:

- PATH

Add the installation directory "/bin".

- MANPATH

Add the installation directory "/share/man".

#### - LD\_LIBRARY\_PATH

Add the installation directory "/lib".

#### Example

The following example configures environment variables when the installation directory is "/opt/fsepv<*x*>assistant".

Note that "<x>" indicates the product version.

#### sh, bash

```
$ PATH=/opt/fsepv<x>assistant/bin:$PATH ; export PATH
$ MANPATH=/opt/fsepv<x>assistant/share/man:$MANPATH ; export MANPATH
$ LD_LIBRARY_PATH=/opt/fsepv<x>assistant/lib:$LD_LIBRARY_PATH ; export LD_LIBRARY_PATH
```

#### csh, tcsh

```
$ setenv PATH /opt/fsepv<x>assistant/bin:$PATH
$ setenv MANPATH /opt/fsepv<x>assistant/share/man:$MANPATH
$ setenv LD_LIBRARY_PATH /opt/fsepv<x>assistant/lib:$LD_LIBRARY_PATH
```

- 3. Create the Mirroring Controller arbitration process management directory that will store the files required by the arbitration server.

  Use ASCII characters in the Mirroring Controller arbitration process management directory.
- 4. In the network configuration file (network.conf), define the Mirroring Controller network configuration that will be managed by the Mirroring Controller arbitration process.

Create network.conf in the Mirroring Controller arbitration process management directory, based on the sample file. For network.conf, set read and write permissions only for the operating system user who starts and stops the arbitration process in step 1

If users other than this are granted access permissions, the mc\_arb command will not work. Accordingly, users other than the operating system user who starts and stops the arbitration process in step 1 are prevented from operating the Mirroring Controller arbitration process.

Sample file

```
/installDir/share/mcarb_network.conf.sample
```

In network.conf, specify the IP address or host name and port number of the primary server and standby server, and define the Mirroring Controller network configuration that will be managed by the Mirroring Controller arbitration process.

Refer to "A.3 Network Configuration File" for details.

A definition example is shown below.

```
Example)
The IDs of the servers are set to "server1" and "server2", and their port numbers are set to "27541".
```

```
server1 192.0.3.100 27541
server2 192.0.3.110 27541
```

5. In the arbitration configuration file (arbitration.conf), define the information related to control of the Mirroring Controller arbitration process.

Create arbitration.conf in the Mirroring Controller arbitration process management directory, based on the sample file. For arbitration.conf, set read and write permissions only for the operating system user who starts and stops the arbitration process in step 1. If users other than this are granted access permissions, the mc\_arb command will not work.

Sample file

```
/installDir/share/mcarb_arbitration.conf.sample
```

Set the parameters shown in the table below in arbitration.conf.

Table 2.1 Parameters

Parameter	Content specified	Remarks
port	Port number of the Mirroring Controller arbitration process	The port number must be 0 to 65535. Ensure that the port number does not conflict with other software. Do not specify an ephemeral port that may temporarily be assigned by another program.
my_address	'ipAddrOrHostNameThatAcceptsConnectionFro mMirroringControllerProcessOnDbServer' [Setting example] my_address = '192.0.3.120'	IPv4 and IPv6 addresses can be specified.  Specify the IP address, enclosed in single quotation marks (').
syslog_ident	'programName'	Specify using single quotation marks (') to enclose the program name used to identify the Mirroring Controller arbitration process message in the system log. Use ASCII characters excluding spaces to specify this parameter.
		The default is 'MirroringControllerArbiter'.
fencing_command	'fencingCmdFilePath'  [Setting example]  fencing_command = '/arbiter/fencing_dir/ execute_fencing.sh'	Specify the full path of the fencing command that fences a database server where it is determined that an error has occurred.  Enclose the path in single quotation marks (').  Specify the path using less than 1024 bytes.
fencing_command_ti meout	Timeout for fencing command (seconds)	If the command does not respond within the specified number of seconds, it is determined that fencing has failed and a signal (SIGTERM) is sent to the fencing command execution process.  Specify a value between 1 and 2147483647.  The default is 20 seconds.



Refer to "A.4.2 Arbitration Configuration File" for information on the parameters and for other parameters.

#### Windows

- 1. On the arbitration server, log in as any operating system user who starts and stops the arbitration process.
- 2. Configure the environment variables.

Set the following environment variable:

- PATH

Add the installation folders "\bin" and "\lib".

# Example

The following example configures environment variables when the installation folder is "c:\Program Files\Fujitsu \fsepv<*x*>assistant64".

Note that "<*x*>" indicates the product version.

```
> SET PATH=c:\Program Files\Fujitsu\fsepv<x>assistant64\bin;c:\Program Files\Fujitsu\fsepv<x>assistant64\lib;%PATH%
```

3. Create the Mirroring Controller arbitration process management directory that will store the files required by the arbitration server.

Use ASCII characters in the Mirroring Controller arbitration process management directory.

4. In the network configuration file (network.conf), define the Mirroring Controller network configuration that will be managed by the Mirroring Controller arbitration process.

Create network.conf in the Mirroring Controller arbitration process management directory, based on the sample file.

Sample file

```
installDir\share\mcarb_network.conf.sample
```

In network.conf, specify the IP address or host name and port number of the primary server and standby server, and define the Mirroring Controller network configuration that will be managed by the Mirroring Controller arbitration process.

Refer to "A.3 Network Configuration File" for details.

A definition example is shown below.

#### Example)

The IDs of the servers are set to "server1" and "server2", and their port numbers are set to "27541".

```
server1 192.0.3.100 27541
server2 192.0.3.110 27541
```

5. Change the access permissions for network.conf.

For network.conf, set read and write permissions only for the operating system user who starts and stops the arbitration process in step 1. If users other than this are granted access permissions, the mc\_arb command will not work. Accordingly, users other than the operating system user who starts and stops the arbitration process in step 1 are prevented from operating the Mirroring Controller arbitration process.

#### Example)

The following is an execution example, in which the operating system user who starts and stops the arbitration process in step 1 is granted full access permissions as the owner when the user is "fsepuser". The following procedure applies when the user is logged in to the Windows server as "fsepuser".

```
> takeown /f network.conf
> icacls network.conf /reset
> icacls network.conf /inheritance:r
> icacls network.conf /grant fsepuser:F
```

6. In the arbitration configuration file (arbitration.conf), define the information related to control of the Mirroring Controller arbitration process.

Create arbitration.conf in the Mirroring Controller arbitration process management directory, based on the sample file.

Sample file

```
installDir\share\mcarb_arbitration.conf.sample
```

Set the parameters shown in the table below in arbitration.conf.

#### Table 2.2 Parameters

Parameter	Content specified	Remarks
port	Port number of the Mirroring Controller arbitration process	The port number must be 0 to 65535. Ensure that the port number does not conflict with other software. Do not specify an ephemeral port that

Parameter	Content specified	Remarks
		may temporarily be assigned by another program.
my_address	'ipAddrOrHostNameThatAcceptsConnectionF romMirroringControllerProcessOnDbServer' [Setting example] my_address = '192.0.3.120'	IPv4 and IPv6 addresses can be specified.  Enclose the parameter value in single quotation marks (').
service_name	'registeredServiceNameOfArbitrationProcessI nWindowsServices'	Use ASCII characters excluding forward slash (/) and backslash (\) to specify this parameter.  Enclose the parameter value in single quotation marks (').  The maximum length of the service name is 124 bytes.
event_source	'eventSourceNameUsedToIdentifyArbitration ProcessMsgInEventLog'	Use ASCII characters excluding spaces to specify this parameter.  Enclose the parameter value in single quotation marks (').  The maximum length of the event source name is 255 bytes.
fencing_command	'fencingCmdFilePath'  [Setting example]  fencing_command = 'c:\\arbiter\\fencing_dir\ \execute_fencing.bat'	Specify the full path of the fencing command that fences a database server where it is determined that an error has occurred.  Specify "\\" as the delimiter.  Enclose the path in single quotation marks (').  Specify the path using less than 260 bytes.  Any multibyte characters must use the same encoding as the operating system.
fencing_command_time out	Timeout for fencing command (seconds)	If the command does not respond within the specified number of seconds, it is determined that fencing has failed and a signal (SIGTERM) is sent to the fencing command execution process.  Specify a value between 1 and 2147483647.  The default is 20 seconds.



Refer to "A.4.2 Arbitration Configuration File" for information on the parameters and for other parameters.

#### 7. Change the access permissions for arbitration.conf.

For arbitration.conf, set read and write permissions only for the operating system user who starts and stops the arbitration process in step 1. If users other than this are granted access permissions, the mc\_arb command will not work.

#### Example

The following is an execution example, in which the operating system user who starts and stops the arbitration process in step 1 is granted full access permissions as the owner when the user is "fsepuser". The following procedure applies when the user is logged in to the Windows server as "fsepuser".

```
> takeown /f arbitration.conf
> icacls arbitration.conf /reset
> icacls arbitration.conf /inheritance:r
> icacls arbitration.conf /grant fsepuser:F
```

8. Configure Windows Firewall.

If Windows Firewall is used, enable the port number of Mirroring Controller specified in the network configuration file in step 3. Refer to "E.2 Windows Firewall Settings" for details.

9. Register the Mirroring Controller arbitration process as a Windows service.

Execute the mc\_arb command in register mode.

For the -P option of the mc\_arb command, specify the password of the operating system user who executes the command.

Example)

```
> mc_arb register -M D:\mcarb_dir\arbiter1 -P *******
```



When specifying the password in the -P option of the mc\_arb command, for security reasons, you should be careful not to allow other users to access it.



Use the mc\_arb command with the -S option to specify automatic start and stop of the Mirroring Controller arbitration process. Refer to "2.13 Setting Automatic Start and Stop of the Mirroring Controller Arbitration Process" for details.

The Mirroring Controller arbitration process is registered as a Windows service using the service name specified in the service\_name parameter of arbitration.conf in step 6.

You can execute the sc qc command to check the registration status.

## 2.3.2 Creating a User Exit for the Arbitration Server

The only user exit for the arbitration server is the fencing command.

The fencing command is a user command that is called by the Mirroring Controller arbitration process if Mirroring Controller performs arbitration processing and determines that a database server is unstable.

In the fencing command, the user implements a process that isolates a database server from a cluster system by, for example, stopping the target operating system or server. The fencing command that was created is to be specified for the parameter in the arbitration configuration file. Refer to "A.4.2 Arbitration Configuration File" for information on the parameters.

- Fencing the primary server during the switch
  - Prevent the Mirroring Controller management process on the primary server from communicating with the Mirroring Controller management process on the other server.
  - Prevent applications from connecting to the primary server instance.
- Fencing the standby server during disconnection
  - Prevent the Mirroring Controller management process on the standby server from communicating with the Mirroring Controller management process on the other server.
  - Prevent applications from connecting to the standby server instance.
  - Prevent the standby server from continuing streaming replication.



Refer to "Appendix C User Commands" for information on user exits.

## 2.3.3 Starting the Mirroring Controller Arbitration Process

This section explains how to start the Mirroring Controller arbitration process.

An operating system user who has logged in to the arbitration server can start the Mirroring Controller arbitration process by executing the mc\_arb command in start mode.

#### Linux

Example)

\$ mc\_arb start -M /mcarb\_dir/arbiter1

#### Windows

Example)

> mc\_arb start -M D:\mcarb\_dir\arbiter1

## 2.4 Setting Up the Primary Server

This section explains how to set up the primary server.

## 2.4.1 Setting Up Database Multiplexing Mode on the Primary Server

This section explains how to set up database multiplexing mode on the primary server.

In database multiplexing, the files that are required for operations are managed in the Mirroring Controller management directory.

There is one Mirroring Controller management directory for each instance.



- Do not place the Mirroring Controller management directory in a directory managed by FUJITSU Enterprise Postgres, otherwise it may be deleted by mistake with the directories managed by FUJITSU Enterprise Postgres, and an old version of files may be restored.

......



- Refer to "Preparing Directories for Resource Deployment" in the Installation and Setup Guide for Server for details on the directories that are managed by FUJITSU Enterprise Postgres.
- Refer to "mc\_ctl" in Reference for information on the command.
- Refer to "Appendix A Parameters" for details on each parameter to be edited for the setup.

Perform the following procedure:

- 1. Log in to the primary server.
- 2. Create the Mirroring Controller management directory that will store the files required by database multiplexing.

Use ASCII characters in the Mirroring Controller management directory.

Additionally, grant "Write" permission to the instance administrator user for the Mirroring Controller management directory.

3. In the network configuration file (network.conf), define the network configuration that will link between the Mirroring Controller processes.

Create the network.conf file in the Mirroring Controller management directory, based on the sample file. For network.conf, set read and write permissions for the instance administrator user only.

If users other than the instance administrator user are granted access, the mc\_ctl command will not work. In this way, users other than the instance administrator user are prevented from operating Mirroring Controller.

Sample file

```
/installDir/share/mc_network.conf.sample
```

In network.conf, specify the IP address or host name and port number of the primary server and standby server, and define the network configuration that will link between the Mirroring Controller processes, and between Mirroring Controller processes and the Mirroring Controller arbitration process.

Refer to "A.3 Network Configuration File" for details.

A definition example is shown below.

The content to be defined depends on the operation settings at the time a heartbeat abnormality is detected.

When automatic degradation by the arbitration server is selected

```
Example)
```

The IDs of the primary server and standby server are set to "server1" and "server2", and their port numbers are set to "27540" and "27541". The ID of the server of the Mirroring Controller arbitration process is set to "arbiter", and its port number is set to "27541".

```
server1 192.0.2.100,192.0.3.100 27540,27541 server
server2 192.0.2.110,192.0.3.110 27540,27541 server
arbiter 192.0.3.120 27541 arbiter
```

Ensure that the port numbers set for the primary server, standby server, and arbitration server do not conflict with other software. In addition, when the arbitration server is used for automatic degradation, use a network in which the arbitration network is not affected by a line failure in the admin network.

When the server type is "server", two IP addresses or host names, and two port numbers need to be specified in the following order:

- IP address or host name of the database server used as the admin network
- IP address or host name of the database server used as the arbitration network
- Port number of the database server used as the admin network
- Port number of the database server used as the arbitration network

If the server type is "arbiter", specify the IP address or host name set for the my\_address parameter and the port number set for the port parameter in arbitration.conf of the arbitration server.

When operation other than automatic degradation by the arbitration server is selected

```
Example) The IDs of the servers are set to "server1" and "server2", and their port numbers are set to "27540".
```

```
server1 192.0.2.100 27540
server2 192.0.2.110 27540
```

Ensure that the port numbers for the primary and standby server do not conflict with other software.

Register in /etc/services the port number of the primary server, because programs such as WebAdmin use it to search for available port numbers.

Register any name as the service name.

4. Define the information related to Mirroring Controller monitoring and control in the serverIdentifier.conf file.

Create the serverIdentifier.conf file in the Mirroring Controller management directory, based on the sample file.

For *serverIdentifier*.conf, set read and write permissions for the instance administrator user only. If users other than the instance administrator user are granted access, the mc\_ctl command will not work.

As the file name for the *serverIdentifier*.conf file, use the server identifier name that was specified in the network.conf file in step 3. Sample file

/InstallDir/share/mc\_server.conf.sample

Set the parameters shown in the table below in the serverIdentifier.conf file.

Table 2.3 Parameters

Parameter	Content specified	Remarks
db_instance	'dataStorageDestinationDir	Use ASCII characters, enclosed in single quotation marks (').
db_instance_password	'passwordOfInstanceAdminUset'	If password authentication is performed, you must specify this parameter in the settings used when Mirroring Controller connects to a database instance. Use ASCII characters, enclosed in single quotation marks (').
		If the specified value of this parameter includes ' or  write \' or \ respectively.
enable_hash_in_pass word	on or off	Specify on to treat the # in the db_instance_password specification as a password character, or off to treat it as a comment.
		The default is "off".
syslog_ident	'programName'	Specify the program name to be used to identify the Mirroring Controller messages in the system log.
		Use ASCII characters excluding spaces, enclosed in single quotation marks (').
		Use the same program name as the parameter in the postgresql.conf file ensures that the Mirroring Controller output content can be referenced transparently, so log reference is easy.
remote_call_timeout	Admin communication timeout	Specify the timeout value (milliseconds) of the Mirroring Controller agent process for communication between servers.
		Specify a value that is less than the operation system TCP connection timeout.
heartbeat_error_action	Operation when a heartbeat abnormality is detected using operating system or server heartbeat	arbitration: Perform automatic degradation using the arbitration server.
	monitoring	command: Call a user exit to determine degradation, and perform automatic degradation if required.
		message: Notify messages.
		fallback: Perform automatic degradation unconditionally.
		Set the same value on the primary server and standby server.

Parameter	Content specified	Remarks
heartbeat_interval	Interval time for abnormality monitoring during heartbeat monitoring of the operating system or server (milliseconds)	Abnormality monitoring of the operating system or server is performed at the interval (milliseconds) specified in
heartbeat_timeout	Timeout for abnormality monitoring during heartbeat monitoring of the operating system or server (seconds)	heartbeat_interval.  This parameter setting is used as the default for database process heartbeat monitoring,
heartbeat_retry	Number of retries for abnormality monitoring during heartbeat monitoring of the operating system or server (number of times)	streaming replication abnormality monitoring, and disk abnormality monitoring. When setting the monitoring time, there are some considerations to take into account to optimize degradation using abnormality monitoring. Refer to "2.11.4.1 Tuning for Abnormality Monitoring of the Operating System or Server" for details.
fencing_command	'fencingCmdFilePath' [Setting example]	Specify the full path of the fencing command that fences a database server where an error is determined to have occurred.
	fencing_command = '/mc/fencing_dir/ execute_fencing.sh'	Enclose the path in single quotation marks (').
		This parameter must be specified when "command" is set for heartbeat_error_action.
		Specify the path using less than 1024 bytes.
fencing_command_ti meout	Fencing command timeout (seconds)	If the command does not respond within the specified number of seconds, fencing is determined to have failed and a signal (SIGTERM) is sent to the fencing command execution process.
		Specify a value between 1 and 2147483647.
		The default is 20 seconds.
arbitration_timeout	Timeout for arbitration processing in the Mirroring Controller arbitration process (seconds)	The specified value must be at least equal to the heartbeat monitoring time of the operating system or server + fencing_command_timeout in the arbitration configuration file.
		If there is no response for at least the number of seconds specified, the primary server will not be switched and the standby server will not be disconnected. Therefore, perform degradation manually.
		Specify a value between 1 and 2147483647.
		This parameter does not need to be set for operation that does not use the arbitration server.
arbitration_command	'arbitrationCmdFilePath'  [Setting example]  arbitration_command = '/mc/arbitration_dir/ execute_arbitration_command.sh'	Specify the full path of the arbitration command to be executed when an abnormality is detected during heartbeat monitoring of the operating system or server.
	=	Enclose the path in single quotation marks (').

Parameter	Content specified	Remarks
		Specify the path using less than 1024 bytes.
arbitration_command _timeout	Timeout for arbitration command (seconds)	If the arbitration command does not respond within the specified number of seconds, it is determined that execution of the arbitration command has failed and a signal (SIGTERM) is sent to the arbitration command execution process.  Specify a value between 1 and 2147483647.  This parameter can be specified only when "command" is set for heartbeat_error_action.



Refer to "A.4.1 Server Configuration File for the Database Servers" for information on the parameters and for other parameters.

## 2.4.2 Creating, Setting, and Registering the Primary Server Instance

This section explains how to create, set, and register the primary server instance.



- Refer to "Client Authentication" in the PostgreSQL Documentation for information on the pg\_hba.conf file.
- Refer to "A.1 Parameters Set on the Primary Server" for information on the postgresql.conf file.
- Refer to "mc\_ctl" in Reference for information on the command.

#### Perform the following procedure:

1. Refer to "Setup" in the Installation and Setup Guide for Server, and then perform the FUJITSU Enterprise Postgres setup and create the FUJITSU Enterprise Postgres instance.

Use ASCII characters in the data storage destination directory.



- If degradation starts occurring due to an error during operations in database multiplexing mode, recovery is required for the standby server. There are some conditions to execute the pg\_rewind command to recover the standby server. One of the conditions can be satisfied by enabling checksums when executing the initdb command. This is not mandatory. Refer to "4.1.1.1.3 Identify cause of error and perform recovery" for details.
- 2. When using transparent data encryption, configure the encryption settings for the storage data.

Create the keystore file.

Refer to "Database Multiplexing Mode" in the Operation Guide for details, and then configure the settings.

3. Add the following entry to the pg\_hba.conf file to authenticate connections from the standby server.

Copy the file to the standby server later.

# TYPE	DATABASE	USER	ADDRESS	METHOD
host	replication	fsep	standbyServerAddress	authenticationMethod
host	replication	fsep	primary Server Address	authenticationMethod

For the primary and standby server addresses, specify the IP address that will connect to the log transfer network.

Additionally, all servers can be used as the primary server or the standby server, so add entries for the addresses of all servers that comprise the database multiplexing system.



#### Setting an authentication method other than trust authentication

If the primary server becomes the standby server, to perform automatic authentication of connections to the primary server, create the .pgpass file in the home directory of the instance administrator user, and then specify a password for the replication database. Accordingly, the instance administrator operating system user and the user registered in the database will be the same, so you can verify that the connection was not made by an unspecified user. Additionally, the password that was set beforehand will be used in the authentication, so that the connection will be automatic.



If trust authentication is set, all OS users who can log in to the primary server will be able to connect, and if one of these is a malicious user, then that user can corrupt the standby server data, or cause the job system to fail, by sending an erroneous transaction log. Therefore, decide on the authentication method according to the security requirements of the system using database multiplexing mode.

Refer to "Authentication Methods" in the PostgreSQL Documentation for details on the authentication methods that can be set.

4. Configure this setting to enable the instance administrator user of the primary server to connect as a database application.

This setting enables the connection to the instance using the user name of the instance administrator user, so that Mirroring Controller can monitor instance errors. Configure this setting to enable the connection to the postgres database.

- If password authentication is used

In the db\_instance\_password parameter of the *serverIdentifier*.conf file, specify the password for the instance administrator user. This password is used to connect to the database instance. If a password is not specified in the db\_instance\_password parameter, the connection to the database instance from Mirroring Controller will fail, and it will not be possible to perform the process monitoring of the instance.

- If password authentication is not used

There is no need to specify the password in the db\_instance\_password parameter.

Even if the password for the instance administrator user is specified in the db\_instance\_password parameter, it will be ignored.

- If certificate authentication using SSL is used

Specify connection parameters for SSL in the db\_instance\_ext\_pq\_conninfo parameter and db\_instance\_ext\_jdbc\_conninfo parameter in the *serverIdentifier*.conf file. If the pasameters are not specified, the connection to the database instance from Mirroring Controller will fail, and it will not be possible to perform the process monitoring of the instance. If certificate authentication using SSL is not performed, the parameters specification is not required.

An example of setting the authentication method is shown below.

# TYPE	DATABASE	USER	ADDRESS	METHOD
host	postgres	fsep	127.0.0.1/32	authenticationMethod



Mirroring Controller uses the PostgreSQL JDBC 4.2 driver to connect to the database instance. Therefore, for the authentication method, specify a method supported by the JDBC driver. If an authentication method not supported by the JDBC driver is specified, Mirroring Controller will fail to start. Refer to the PostgreSQL JDBC Driver Documentation for information on authentication methods supported by the JDBC driver.

5. To use database multiplexing mode, specify the parameters shown in the table below in the postgresql.conf file.

The postgresql.conf file is copied when the standby server instance is created. Accordingly, set the required parameters in the standby server.

To use database multiplexing mode, specify the parameters shown in the table below in the postgresql.conf file. After editing the postgresql.conf file, restart the instance.

Table 2.4 Parameters

Parameter	Content specified	Remarks
wal_level	replica or logical	Specify "logical" when logical decoding is also to be used.
max_wal_senders	2 or more	Specify "2" when building a Mirroring Controller cluster system.
		When additionally connecting asynchronous standby servers to the cluster system, add the number of simultaneous connections from these standby servers.
synchronous_standby_names	'standbyServerName'	Specify the name that will identify the standby server.
		Enclose the name in single quotation marks (').
		Do not change this parameter while Mirroring Controller is running.
		Do not specify multiple names to this parameter as the Mirroring Controller can manage only one standby server.
hot_standby	on	
wal_keep_segments	Number of file segments	If a delay exceeding the value set in this parameter occurs, the WAL segment required later by the primary server may be deleted.
		Additionally, if you stop a standby server (for maintenance, for example), consider the stop time and set a value that will not cause the WAL segment to be deleted.
		Refer to "Estimating Transaction Log Space Requirements" in the Installation and Setup Guide for Server for information on estimating the WAL segment.
wal_log_hints	on	When using the pg_rewind command to recover a standby server, specify this parameter or enable checksums when executing the initdb command.
wal_sender_timeout	Timeout (milliseconds)	Specify the time period after which it is determined that an error has occurred in the transaction log transfer on the primary server.
		By aligning this value with the value for the database process heartbeat monitoring time, you can unify the time after which it is determined that an error has occurred.
archive_mode	on	Specify the archive log mode.
archive_command	'installDir/bin/ pgx_walcopy.cmd "%p" "backupDataStorageDestin ationDirectory/ archived_wal/%f"	Specify the command and storage destination to save the transaction log.
backup_destination	Backup data storage destination directory	Specify the name of directory where to store the backup data.
		Set the permissions so that only the instance administrator user can access the specified directory.

Parameter	Content specified	Remarks
		Specify the same full path on all servers, so that the backup data of other servers can be used to perform recovery.
max_connections	Number of simultaneous client connections to the instance + superuser_reserved_conne ctions	The value specified is also used to restrict the number of connections from client applications and the number of connections for the management of instances.  Refer to "When an Instance was Created with the initdb Command" in the Installation and Setup Guide for Server,
		and "Connections and Authentication" in the PostgreSQL Documentation, for details.
superuser_reserved_connections	Add the number of simultaneous executions of	Specify the number of connections reserved for connections from database superusers.
	mc_ctl status (*1) + 2	Add the number of connections from Mirroring Controller processes. Also reflect the added value in the max_connections parameter.
wal_receiver_timeout	Timeout (milliseconds)	Specify the time period after which it is determined that an error has occurred when the transaction log was received on the standby server.
		By aligning this value with the value for the heartbeat monitoring time of the database process, you can unify the time after which it is determined that an error has occurred.
restart_after_crash	off	If "on" is specified, or the default value is used for this parameter, behavior equivalent to restarting FUJITSU Enterprise Postgres, including crash recovery, will be performed when some server processes end abnormally.
		However, when database multiplexing monitoring is used, a failover will occur after an error is detected when some server processes end abnormally, and the restart of those server processes is forcibly stopped. Specify "off" to prevent behavior such as this from occurring for no apparent reason.
synchronous_commit	on or remote_apply	Specify up to what position WAL send is to be performed before transaction commit processing returns a normal termination response to a client.
		The recommended value is "on" or "remote_apply" to prevent data loss caused by operating system or server down immediately after a switch or switch.
recovery_target_timeline	latest	Specify "latest" so that the new standby server (original primary server) will follow the new primary server when a switch occurs.
		This parameter is required when the original primary server is incorporated as a new standby server after the primary server is switched.

<sup>\*1:</sup> Number of simultaneous executions of the mc\_ctl command in the status mode.

## 2.4.3 Starting Mirroring Controller on the Primary Server

This section explains how to start Mirroring Controller on the primary server.

When the arbitration server is used for automatic degradation, start the Mirroring Controller arbitration process on the arbitration server in advance.

1. Start the Mirroring Controller process.

#### Enabling automatic switch/disconnection

As the instance administrator user, execute the mc\_ctl command in start mode.

Example)

\$ mc\_ctl start -M /mcdir/inst1

#### Disabling automatic switch/disconnection

As the instance administrator user, execute the mc\_ctl command in start mode with the -F option specified.

Example)

\$ mc\_ctl start -M /mcdir/inst1 -F



- When the arbitration server is used for automatic degradation, the database server must connect to the arbitration server, and as a result, Mirroring Controller startup may take longer than when the arbitration server is not used.
- If the parameter for heartbeat monitoring of operating systems or servers set by the abitration server is greater than parameter for heartbeat monitoring of operating systems and servers of the Mirroring Controller, the Mirroring Controller may fail to start. In this case, check the contents of the message notification and review the parameters for heartbeat monitoring of operating systems or servers for the arbitration server or Mirroring Controller.
- If the heartbeat\_error\_action parameter in serverIdentifier.conf is set to "message", even if automatic switch/disconnection is enabled and Mirroring Controller is started, only message output is performed when a heartbeat abnormality is detected during heartbeat monitoring of operating systems and servers - switch/disconnection is not performed.
- Mirroring Controller startup usually fails if the standby server is mistakenly started as the primary server or if the old primary server is not recovered after the switch and is then mistakenly started as the primary server. However, if the admin network is disconnected, then startup does not fail, and both servers may become primary servers. Therefore ensure that the admin network is connected before starting Mirroring Controller.



- The mc\_ctl command fails if the Mirroring Controller arbitration process has not been started on the arbitration server when the arbitration server is used for automatic degradation. However, if the Mirroring Controller arbitration process cannot be started in advance, it can be started by specifying the --async-connect-arbiter option in the mc\_ctl command.
- After Mirroring Controller is started, automatic switch/disconnection can be enabled or disabled using the enable-failover or disable-failover mode of the mc\_ctl command.
- 2. Obtain the backup.

Use the pgx\_dmpall command to collect the backup.

## 2.5 Setting Up the Standby Server

This section explains how to set up the standby server.

## 2.5.1 Setting Up Database Multiplexing Mode on the Standby Server

This section explains how to set up database multiplexing mode on the standby server.

In database multiplexing, the files that are required for operations are managed in the Mirroring Controller management directory.

There is one Mirroring Controller management directory for each instance.



- Do not place the Mirroring Controller management directory in a directory managed by FUJITSU Enterprise Postgres, otherwise it may be deleted by mistake with the directories managed by FUJITSU Enterprise Postgres, and an old version of files may be restored.
- When creating a standby server for a large database, stop job system operations, specify a large value for the wal\_keep\_segments parameter, or use replication slots.

This is because WALs generated after the standby server is built using the pg\_basebackup command, but before it is started, need to be retained. However, the number of WAL segments that can be retained is constrained by the wal\_keep\_segments parameter. Additionally, setting the wal\_keep\_segments parameter requires consideration regarding stabilization of the database multiplexing mode (refer to "2.11.1 Tuning to Stabilize the Database Multiplexing Mode" for details).



- Refer to "Preparing Directories for Resource Deployment" in the Installation and Setup Guide for Server for details on the directories that are managed by FUJITSU Enterprise Postgres.

......

- Refer to "pg\_basebackup" in "Reference" in the PostgreSQL Documentation for information on the pg\_basebackup command.
- Refer to "mc\_ctl" in Reference for information on the command.
- Refer to "Appendix A Parameters" for details on each parameter to be edited for the setup.
- Refer to "Replication Slots" in the PostgreSQL Documentation for information on replication slots.

#### Perform the following procedure:

- 1. Log in to the standby server.
- 2. Create the Mirroring Controller management directory that will store the files required by database multiplexing.

Use ASCII characters in the Mirroring Controller management directory.

Additionally, grant "Write" permission to the instance administrator user for the Mirroring Controller management directory.

 $3. \; \text{Copy}, \text{ and then deploy, the network.conf file of the primary server.}$ 

Copy the network.conf file that was defined in the primary server setup, and deploy it to the Mirroring Controller management directory of the standby server.

Set read and write permissions for the instance administrator user only. If users other than the instance administrator user are granted access, the mc\_ctl command will not work. Accordingly, users other than the instance administrator user are prevented from operating Mirroring Controller.

Register in /etc/services the port number of the standby server that was specified in the network.conf file, because programs such as WebAdmin use it to search for available port numbers.

Register any name as the service name.

4. Copy, and then deploy, the *serverIdentifier*.conf file of the primary server.

Copy the *serverIdentifier*.conf file that was defined in the primary server setup, and deploy it to the Mirroring Controller management directory of the standby server.

Set read and write permissions for the instance administrator user only. If users other than the instance administrator user are granted access permissions, the mc\_ctl command will not work.

## 2.5.2 Creating, Setting, and Registering the Standby Server Instance

This section explains how to create, set, and register the standby server instance.



- Refer to "Appendix A Parameters" for details on each parameter.
- Refer to "mc\_ctl" in Reference for information on the command.

#### Perform the following procedure:

1. Set the kernel parameters.

Refer to "Configuring Kernel Parameters" in the Installation and Setup Guide for Server for details.

2. When using transparent data encryption, configure the encryption settings for the storage data.

Deploy a copy of the keystore file of the primary server on the standby server.

Refer to "Database Multiplexing Mode" in the Operation Guide for details.

3. Execute the pg\_basebackup command to create a copy of the primary server instance on the standby server.

Example)

\$ pg\_basebackup -D /database/inst1 -X fetch --waldir=/transaction/inst1 --progress --verbose -R
--dbname='application\_name=standbyServerName' -h primaryServerIpAddress -p
primaryServerPortNumber



- Use the pg\_basebackup command with the -R option to create a standby.signal file. If you do not create the standby.signal file, the Mirroring Controller cannot be started as a standby server.

- If using a method that requires password authentication for connections to the primary server, you will need to ensure that authentication is performed automatically. If the -R option is specified for the pg\_basebackup command and the password parameter is specified for the --dbname option, the pg\_basebackup command will set the password in the primary\_conninfo parameter in postgresql.auto.conf file, enabling connections to be performed automatically.

If a password is not set in the primary\_conninfo parameter in postgresql.auto.conf file, it will be necessary to create a .pgpass file in the home directory of the instance administrator user, and specify a password for the replication database.

- The primary\_conninfo parameter should not be set in the postgresql.conf file, but only in the postgresql.auto.conf file using the pg\_basebackup command.
- When executing the  $pg\_basebackup$  command, consider the following for collection of transaction logs.
  - When "fetch" is specified for the -X option of the command

Transaction logs are collected at the end of the backup, so it is necessary to ensure that transaction logs that occur during backup are not deleted from the primary server. Therefore, allow for a sufficient value for the wal\_keep\_segments parameter in postgresql.conf.

- When the -X option is omitted or "stream" is specified for the -X option of the command

Transaction logs are streamed, so when Mirroring Controller is running on the primary server, the connection is changed to a synchronous standby server on detection of a streaming replication connection using this command. Therefore, if a job has started on the primary server, the primary server will be impacted, therefore execute this command after stopping only the Mirroring Controller process on the primary server.



Refer to "Hot Standby" in the PostgreSQL Documentation for information on the standby signal file.

4. Set the parameters shown in the table below in the postgresql.conf file.

Table 2.5 Parameters

Parameter	Content specified	Remarks
synchronous_standby_names	'primaryServerName'	Required after switching the primary server and then changing the original primary server to the new standby server.
		Enclose the name in single quotation marks (').
		Do not change this parameter while Mirroring Controller is running.
		Do not specify multiple names to this parameter as the Mirroring Controller can manage only one standby server

## 2.5.3 Starting Mirroring Controller on the Standby Server

This section explains how to start Mirroring Controller on the standby server.

When the arbitration server is used for automatic degradation, start the Mirroring Controller arbitration process on the arbitration server in advance.

1. After ensuring that the Mirroring Controller process of the primary server has started, start Mirroring Controller on the standby server.

#### Enabling automatic switch/disconnection

As the instance administrator user, execute the mc\_ctl command in start mode with the -f option specified. This action enables automatic switch/disconnection.

If you start Mirroring Controller and the instance without specifying the -f option, automatic switch/disconnection will not be enabled. To enable both, start Mirroring Controller and then execute the mc\_ctl command in enable-failover mode or restart Mirroring Controller with the -f option specified.

Example)

\$ mc\_ctl start -M /mcdir/inst1

#### Disabling automatic switch/disconnection

As the instance administrator user, execute the mc\_ctl command in start mode with the -F option specified.

Example)

\$ mc\_ctl start -M /mcdir/inst1 -F

2. Check the status of the Mirroring Controller process.

As the instance administrator user, execute the mc\_ctl command in status mode. Ensure that "mirroring status" is switchable.

Example)

\$ mc\_ctl status -M /mcdir/inst1



- When the arbitration server is used for automatic degradation, the time required for the database server to connect to the arbitration server is added on. Therefore, Mirroring Controller startup may take longer than when the arbitration server is not used.
- If the parameter for heartbeat monitoring of operating systems or servers set by the abitration server is greater than parameter for heartbeat monitoring of operating systems and servers of the Mirroring Controller, the Mirroring Controller may fail to start. In this case, check the contents of the message notification and review the parameters for heartbeat monitoring of operating systems or servers for the arbitration server or Mirroring Controller.
- If the heartbeat\_error\_action parameter in *serverIdentifier*.conf is set to "message", even if automatic switch/disconnection is enabled and Mirroring Controller is started, only message output is performed when a heartbeat abnormality is detected during heartbeat monitoring of operating systems and servers switch/disconnection is not performed.

- Mirroring Controller startup usually fails if the standby server is mistakenly started as the primary server or if the old primary server is not recovered after the switch and is then mistakenly started as the primary server. However, if the admin network is disconnected, then startup does not fail, and both servers may become primary servers. Therefore, ensure that the admin network is connected before starting Mirroring Controller.



- The mc\_ctl command fails if the Mirroring Controller arbitration process has not been started on the arbitration server when the arbitration server is used for automatic degradation. However, if the Mirroring Controller arbitration process cannot be started in advance, it can be started by specifying the --async-connect-arbitration in the mc\_ctl command.

•••••

- After Mirroring Controller is started, automatic switch/disconnection can be enabled or disabled using the enable-failover or disable-failover mode of the mc\_ctl command.

## 2.6 Creating a User Exit for a Database Server

This section explains how to create a user exit for a database server.

The user command types explained below can be used as user exits. These commands are called by Mirroring Controller management processes.

The user can create user exits as required.

Specify the user commands that were created for the parameters in the server configuration file of the database server. Refer to "A.4.1 Server Configuration File for the Database Servers" for information on these parameters.

#### **User command types**

- Fencing command

This user command performs fencing if Mirroring Controller performs arbitration processing and determines that a database server is unstable.

- Arbitration command

This user command performs arbitration processing in lieu of the arbitration server when there is no arbitration server.

- State transition commands

These user commands are called when Mirroring Controller performs state transition of a database server. It includes the following types:

- Post-switch command

This user command is called after a promotion from standby server to primary server.

- Pre-detach command

This user command is called before the standby server is disconnected from a cluster system.

If the pre-detach command is specified on both the primary server and standby server, it is called first on the standby server and then on the primary server.

If the settings are configured to forcibly stop the instance on the standby server when the standby server is disconnected, the predetach command is called on the standby server and then the instance on the standby server is stopped.

- Post-attach command

This user command is called after the standby server has been attached to a cluster system.

If the post-attach command is specified on both the primary server and standby server, it is called first on the primary server and then on the standby server.



When the arbitration server is used for automatic degradation and the requirements can be satisfied using the fencing command on the arbitration server only, the fencing command on the database server is not required. In addition, if the requirements can be satisfied using the fencing command on the database server only, create a fencing command on the arbitration server for termination processing only (without implementation).

Table 2.6 Availability of user commands, and database server calling the command

	Operation when a heartbeat abnormality is detected using operating system or server heartbeat monitoring				Database server calling the command	
User command	Message output	Automatic degradation by arbitration server	Automatic degradation by arbitration command	Unconditio nal automatic degradatio n	Primary server	Standby server
Fencing command	Y (*1)	Y (*2)	R	N	Y	Y
Arbitration command	N	N	R	N	Y	Y
Post-switch command	Y	Y	Y	Y	Y	N
Pre-detach command	Y	Y	Y	Y	Y	Y (*3)
Post-attach command	Y	Y	Y	Y	Y	Y (*3)

- R: Required
- Y: Can be used
- N: Cannot be used
- \*1: Called only when the mc\_ctl command is used to execute forced switching or forced disconnection.
- \*2: Creation of a fencing command on a database server is optional, but it must be created on the arbitration server.
- \*3: If message output or unconditional automatic degradation is selected, this command is called only from the primary server.



Refer to "Appendix C User Commands" for information on the interface for each user command.

## 2.7 Confirming the Streaming Replication Status

Before performing the setup of the database multiplexing mode, ensure that the prerequisite streaming replication feature has been set up correctly.

Perform the following procedure:

1. On the primary server, ensure that single-row searches can be performed using the pg\_stat\_replication statistics view.

An example output of the psql command is shown below.

Example)

```
client_hostname |
client_port
               55098
backend_start
             2020-03-23 11:17:49.628793+09
backend_xmin
               streaming
sent_lsn
               0/3000060
write_lsn
               0/3000060
               0/3000060
flush_lsn
               0/3000060
replay_lsn
write_lag
flush_lag
replay_lag
sync_priority
                | 1
sync_state
                 sync
reply_time
                2020-03-23 11:23:27.703366+09
```

#### 2. Confirm the search results of step 1.

Ensure that the connection established with the intended standby server is in synchronous mode.

Table 2.7 Items to be checked

Item	Required value
application_name	Value specified for synchronous_standby_names parameter in the postgresql.conf file of the primary server.
client_addr	IP address of the standby server.
state	"streaming".
sync_state	"sync".



- Refer to "The Statistics Collector" in "Server Administration" in the PostgreSQL Documentation for information on the pg\_stat\_replication statistics view.
- Note that the pg\_stat\_replication statistics view may change in the future.

## 2.8 Checking the Connection Status

This section explains how to check the connection status from a database server or the arbitration server.

## 2.8.1 Checking the Connection Status on a Database Server

This section explains how to use a database server to check the connection status of the Mirroring Controller arbitration process and the Mirroring Controller process on the primary server and standby server.

Perform the following procedure:

 $1. \ \ On the primary server and standby server, execute the mc\_ctl command in status mode with the \textit{--arbiter} option specified.$ 

```
Example)
```

The mc\_ctl command is executed with the --arbiter option specified, and the status is output.

2. On the primary server and standby server, check the result displayed by executing the mc\_ctl command in status mode in step 1.

Items to be checked

Check that the output status" is "online".



Refer to the Reference for information on the mc\_ctl command.

## 2.8.2 Checking the Connection Status on the Arbitration Server

This section explains how to use the arbitration server to check the connection status of the Mirroring Controller arbitration process and the Mirroring Controller process on the primary server and standby server.

Perform the following procedure:

1. Execute the mc arb command in status mode on the arbitration server.

The example below executes the mc\_arb command, and shows the status.

#### Linux

Example)

\$ mc_arb st	atus -M /mcarb_c	dir/arbiter1
server_id	host	status
server1 server2	192.0.3.100 192.0.3.110	online online

#### Windows

Example)

> mc_arb st	atus -M D:\mcark	o_dir\arbiter1
server_id	host	status
server1 server2	192.0.3.100 192.0.3.110	online online

2. Check the result displayed by executing the  $mc_{arb}$  command in step 1.

Items to be checked

Check that the output status is "online" on both lines.



Refer to the Reference for information on the mc\_arb command.

## 2.9 Creating Applications

This section explains how to create applications using database multiplexing, and points that should be noted when you create the applications.

## 2.9.1 Application Connection Server Settings

If database multiplexing is used and a failover occurs, it will be necessary to switch the application connection server. Accordingly, use the application connection switch feature to create applications.



Refer to "Application Connection Switch Feature" in the Application Development Guide for details.

## 2.10 Checking the Behavior

To check if the environment setup was performed correctly, start the application and then check the behavior of the switch and rebuild.

## 2.11 Tuning

This section explains how to tune database multiplexing mode.

## 2.11.1 Tuning to Stabilize the Database Multiplexing Mode

When large amounts of data are updated, the write-to load for the database will become great, and the multiplexing state may become unstable.

Accordingly, by editing the parameters below in the postgresql.conf file, a stable multiplexing state can be maintained. Refer to "Estimating Transaction Log Space Requirements" in the Installation and Setup Guide for Server for information on transaction log space requirements.

#### Table 2.8 Parameters

Parameter	Content
wal_keep_segments	Refer to "2.4.2 Creating, Setting, and Registering the Primary Server Instance" for details.
max_wal_size	The transaction log is written out according to the checkpoint trigger.
	If a transaction log with the capacity of the value specified in this parameter is generated, the checkpoint will be executed.
	If a large value is specified in this parameter, the time required for crash recovery will increase.
	If a small value is specified in this parameter, many checkpoints will be generated, which will affect the performance of the applications that connect to the primary server.

## 2.11.2 Tuning to Stabilize Queries on the Standby Server

Queries made using reference jobs on the standby server may be canceled by jobs executed on the primary server.

To reduce the possibility of a job being canceled, specify as large a value as possible for the max\_standby\_archive\_delay parameter in the postgresql.conf file.



#### See

- Refer to "Handling Query Conflicts" in the PostgreSQL Documentation for details.
- Refer to "Standby Servers" in the PostgreSQL Documentation for details on the max\_standby\_archive\_delay parameter.

# 2.11.3 Tuning to Stabilize Queries on the Standby Server (when Performing Frequent Updates on the Primary Server)

If jobs are updated on the primary server regularly and frequently, it will be easy for the query made by the reference job on the standby server to be canceled. In this case, edit one of the postgresql.conf file parameters shown in the table below.

#### Table 2.9 Parameters

Parameter	Description	
hot_standby_feedback	When "on" is set, the deletion (vacuum) of the data area that was deleted or updated on the primary server is suppressed.  Accordingly, the query on the standby server will not be canceled. (*1)	
vacuum_defer_cleanup_age	The deletion (vacuum) of the data area that was deleted or updated on the primary server is delayed until the specified number of transactions is processed. Accordingly, the probability that the query on the standby server will be canceled decreases.	

<sup>\*1:</sup> Because the vacuum is delayed, the data storage destination disk space of the primary server comes under pressure.

Additionally, if there is conflict between accesses and queries executed on the standby server, transaction logs indicating this conflict will be transferred.

Accordingly, specify as large a value as possible for the max\_standby\_archive\_delay parameter so that access conflicts do not occur.



- Refer to "Standby Servers" in the PostgreSQL Documentation for details on the hot\_standby\_feedback parameter.
- Refer to "Master Server" in the PostgreSQL Documentation for details on the vacuum\_defer\_cleanup\_age parameter.

## 2.11.4 Tuning for Optimization of Degradation Using Abnormality Monitoring

Mirroring Controller uses a monitoring method that outputs an error if the timeout or number of retries is exceeded when accessing resources targeted for monitoring. Setting inappropriate values in these settings may lead to misdetection or a delay in automatic degradation, so you must design these values appropriately.

For example, the following type of issue occurs if the tuning related to abnormality monitoring is not performed appropriately.

- If the timeout is too short
  - Results in redundant degradation and availability falls.
- If the timeout is too long

It takes longer for automatic degradation to be performed even when an error affecting operational continuity occurs, potentially causing downtime.

You can optimize degrading operation by editing the values for the parameters in the server configuration file described below in accordance with the system. Refer to "A.4 Server Configuration File" for information on how to edit these parameters.

### 2.11.4.1 Tuning for Abnormality Monitoring of the Operating System or Server

Tuning for abnormal monitoring of the operating system or server depends on the operation when heartbeat abnormality is detected by the heartbeat monitoring of operating systems or servers.



Refer to "1.1.1 Monitoring Using Database Multiplexing Mode" for the operation when heartbeat abnormality is detected in the the heartbeat monitoring of operating systems or servers.

## 2.11.4.1.1 Tuning Abnormality Monitoring for Operations that Use an Arbitration Server for Automatic Degeneration

In an operation that use an arbitration server for automatic degeneration, the database server is periodically monitored for abnormalities so that the Mirroring Controller arbitration process can immediately respond to an arbitration request from the Mirroring Controller. The

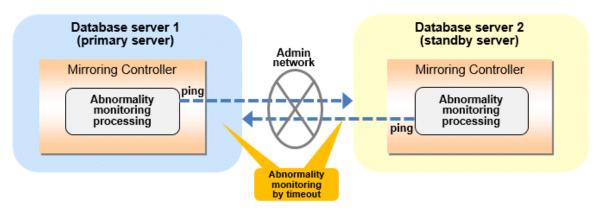
automatic degradation using the arbitration server can optimize the time from error detection to automatic degradation of the operating systems or servers by editing the following parameters.

- Parameters for the abnormality monitoring of the operating system or server in the server configuration file of the database server
- Parameters for the abnormality monitoring of the operating system or server in the arbitration configuration file
- Parameters for the arbitration processing and fencing

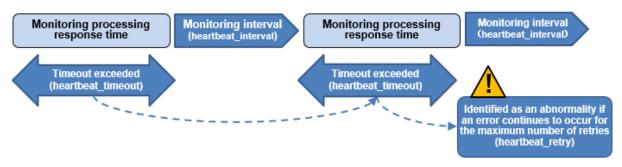
## Parameters for the abnormality monitoring of the operating system or server in the server configuration file of the database server

Table 2.10 Parameters for the abnormality monitoring of the operating system or server in the server configuration file of the database server

Parameter	Description
Abnormality monitoring interval (heartbeat_interval)	Mirroring Controller is configured so that abnormality monitoring does not place a load on the system. This parameter does not normally need to be set. (The default is 800 milliseconds.)
Abnormality monitoring timeout (heartbeat_timeout)	Take into account the time during which a load is placed continuously on the server or admin network performance. For example, it is envisaged that this parameter will be used in situations such as when performing high-load batch jobs or when a large number of online jobs occur continuously and concurrently. (The default is 1 second.)
Abnormality monitoring retries (heartbeat_retry)	This parameter can be set when needing a safety value for situations in which the value specified for heartbeat_timeout is exceeded, for example, when using systems with fluctuating loads, however, this parameter does not normally need to be set. (The default is 2 times.)



#### Flow of abnormality monitoring by timeout



The expression for calculating the time required to detect an abnormality by Mirroring Controller is shown below.

```
Abnormality detection time of Mirroring Controller = ( heartbeat_timeout(seconds) + heartbeat_interval(milliseconds) / 1000 ) x ( heartbeat_retry(number of times) + 1)
```

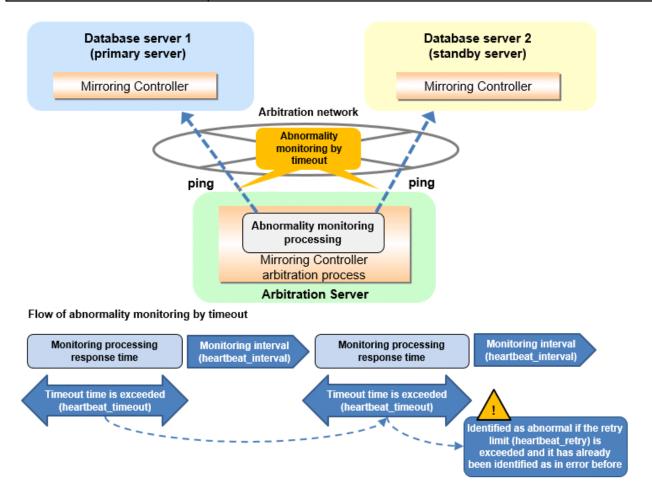
The abnormality detection time when the default value is used is shown below.

```
Abnormality detection time of Mirroring Controller = ( 1 + 800 / 1000 ) x ( 2 + 1 ) = 5.4(seconds)
```

## Parameters for the abnormality monitoring of the operating system or server in the arbitration configuration file

Table 2.11 Parameters for the abnormality monitoring of the operating system or server in the arbitration configuration file

Parameter	Description
Abnormality monitoring interval (heartbeat_interval)	Mirroring Controller arbitration process is configured so that abnormality monitoring does not place a load on the system. This parameter does not normally need to be set. (The default is the value set in heartbeat_interval in the server configuration file of the database server.) (milliseconds).
Abnormality monitoring timeout (heartbeat_timeout)	Take into account the time during which a load is placed continuously on the server and arbitration network capabilities. (The default is the value set in heartbeat_timeout in the server configuration file of the database server.) (seconds).
Abnormality monitoring retries (heartbeat_retry)	This parameter can be set when needing a safety value for situations in which the value specified for heartbeat_timeout is exceeded, for example, when using systems with fluctuating loads, however, this parameter does not normally need to be set. (The default is the value set in heartbeat_retry in the server configuration file of the database server.) (number of times)



The expression for calculating the time required to detect an abnormality by Mirroring Controller arbitration process is shown below.

Abnormality detection time of Mirroring Controller arbitration process = ( heartbeat\_timeout(seconds) + heartbeat\_interval(milliseconds) / 1000 ) x ( heartbeat\_retry(number of times) + 1)

The abnormality detection time when the default value is used is shown below.

Abnormality detection time of Mirroring Controller arbitration process = ( 1 + 800 / 1000 ) x ( 2 + 1 ) = 5.4(seconds)



The abnormality detection time of the operation for automatic degradation using the arbitration server can be calculated as follows.

Abnormality detection time = Max(Abnormality detection time by Mirroring Controller, Abnormality detection time by Mirroring Controller arbitration process)



If the heartbeat\_interval is set in the arbitration configuration file, the relationship between the parameter for operating system or server abnormality monitoring specified in the server configuration file of the database server file and the heartbeat\_interval of the arbitration configuration file must satisfy the following relational expression.

.....

Heartbeat\_interval in the arbitration configuration file (milliseconds) / 1000 ) <
( heartbeat\_timeout(seconds) + heartbeat\_interval(milliseconds) / 1000 ) \* heartbeat\_retry(number of times) + heartbeat\_timeout(seconds)</pre>

#### Parameters for the arbitration processing and fencing

Table 2.12 Parameters for the arbitration processing and fencing

Parameter	Description
Arbitration processing timeout (arbitration_timeout in the server configuration file of the database server)	Take into account the time to perform arbitration processing on the Mirroring Controller arbitration process. The value must be greater than or equal to abnormality detection time of Mirroring Controller arbitration process + fencing_command_timeout in the arbitration configuration file (seconds).
Fencing timeout (fencing_command_timeout in the arbitration configuration file)	Take into account the time to execute the fencing command (seconds).

#### Flow from the abnormality detection to the automatic degeneracy

When performing automatic degradation using the arbitration server, the flow from the abnormality detection in the operating system or server to the occurrence of automatic degeneracy and the parameters is shown below.

Flow from the abnormality detection to the automatic degeneracy	Description	Parameter	
(1) Abnormality detection	Mirroring Controller detect the database server operating system or server errors.	Parameters for the abnormal operating system or service configuration file of the	
(2) Arbitration request	Mirroring Controller that detect the operating system or server error asks the Arbitration Server	-	arbitration_timeout in the server configuration

Flow from the abnormality detection to the automatic degeneracy	Description	Para	meter
	to check the status of the other server's operating system or server.		file of the database server
(3) Arbitration processing	The Mirroring Controller arbitration process checks the status of the other server's operating system or server. However, if the result of the operating system or server abnormality monitoring by the arbitration server has been determined before the arbitration request from the Mirroring Controller of the database server, this process is not performed.	Parameters for the abnormality monitoring of the operating system or server in the arbitration configuration file	
(4) Fencing	If the Mirroring Controller arbitration process determines that the other server is an abnomaly of the operating system or server, it fences the other server and isolates it from the cluster system. If the Mirroring Controller arbitration process determines that the operating system or server status is normal, this process and the (6) are not performed.	fencing_command_ti meout in the arbitration configuration file	
(5) Return of the arbitration results	Returns the results of the arbitration to the Mirroring Controller of the database server that requested the arbitration.	-	
(6) Automatic degradation	The automatic degradation is performed. If fencing fails in (4), this procedure is not performed.	-	

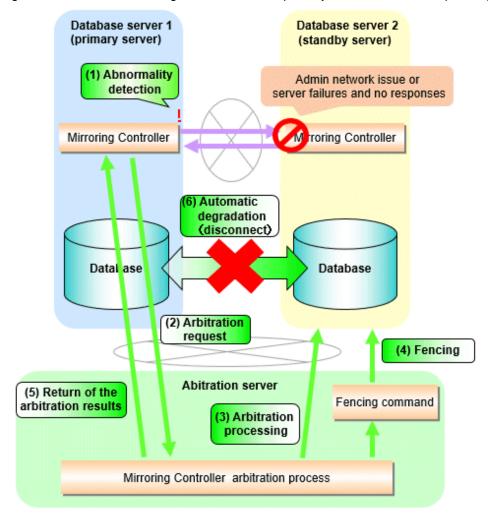
<sup>-:</sup> No associated parameters



If the fencing\_command parameter is specified in the server configuration file of the database server, the fencing command is invoked on the database server if fencing is successful on the arbitration server. In that case, add the value of the fencing\_command\_timeout parameter in the server configuration file of the database server to the estimate.

......

Figure 2.1 When the Mirroring Controller on the primary server detects an operating system or server error



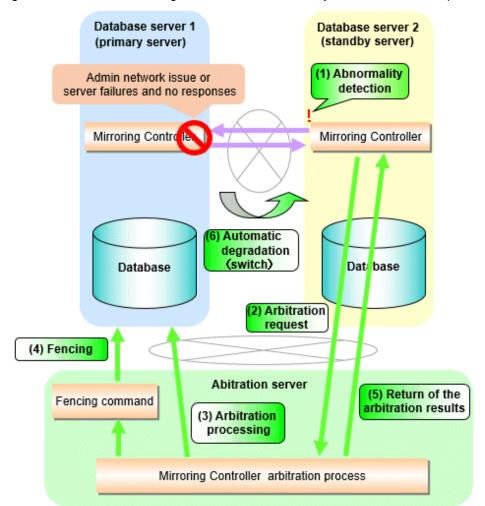


Figure 2.2 When the Mirroring Controller on the standby server detects an operating system or server error

## 2.11.4.1.2 Tuning Abnormality Monitoring for Operations that Perform Automatic Degeneration by Calling a User Exit that Determines Degeneration

In an operation that perform automatic degeneration by calling a user exit that determines degeneration, you can optimize the time from operating system or server abnormality detection to automatic degradation by editing the operating system or server abnormality monitoring parameters and parameters related to arbitration processing and fencing in the server configuration file of the database server. Refer to "Parameters for the abnormality monitoring of the operating system or server in the server configuration file of the database server" for information on the operating system or server abnormality monitoring parameters in the server configuration file of the database server.

Table 2.13 Parameters for the arbitration processing and fencing

Parameter	Description
Arbitration processing timeout (arbitration_command_timeout)	Take into account the time to execute the arbitration command(seconds).
Fencing timeout (fencing_command_timeout)	Take into account the time to execute the fencing command (seconds).

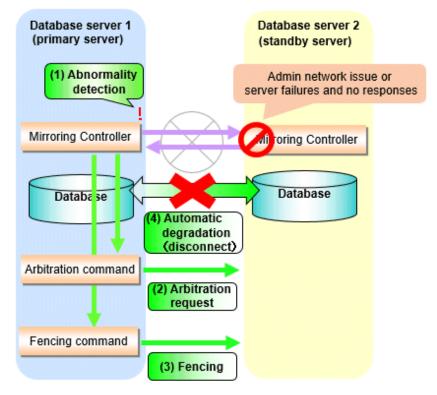
#### Flow from the abnormality detection to the automatic degeneracy

When performing automatic degradation by calling a user exit that determines degeneration, the flow from the abnormality detection in the operating system or server to the occurrence of automatic degeneracy and the parameters is shown below.

Flow from the abnormality detection to the automatic degeneracy	Description	Parameter
(1) Abnormality detection	Mirroring Controller detect the database server operating system or server errors.	Parameters for the abnormality monitoring of the operating system or server in the server configuration file of the database server
(2) Arbitration processing	An arbitration command is executed to check the status of the other server's operating system or server.	arbitration_command_timeout in the server configuration file of the database server
(3) Fencing	If the operating system or server status of the other server is abnomal in (2), it fences the other server and isolates it from the cluster system. If the operating system or server status of the other server is nomal in (2), this process and (4) are not executed.	fencing_command_timeout in the server configuration file of the database server
(4) Automatic degradation	The automatic degradation is performed. If fencing fails in (3), this procedure is not performed.	-

<sup>-:</sup> No associated parameters

Figure 2.3 When the Mirroring Controller on the primary server detects an operating system or server error



Database server 1 Database server 2 (standby server) (primary server) (1) Abnormality Admin network issue or detection server failures and no responses Mirroring Contro Mirroring Controller Database Da tabase (4) Automatic degradation (switch) Arbitration command (2) Arbitration request Fencing command (3) Fencing

Figure 2.4 When the Mirroring Controller on the standby server detects an operating system or server error

#### 2.11.4.1.3 Tuning Abnormality Monitoring for Operations that Notify Messages

In an operation that notify messages, you can optimize the abnormality detection time by editing the operating system or server abnormality monitoring parameters in the server configuration file of the database server. Refer to "Parameters for the abnormality monitoring of the operating system or server in the server configuration file of the database server" for information on the operating system or server abnormality monitoring parameters in the server configuration file of the database server. In addition, when the Mirroring Controller detects an error, it does not perform the arbitration processing, fencing, or automatic degradation, but only notification messages is performed.

## 2.11.4.1.4 Tuning Abnormality Monitoring for Operations that Perform Automatic Degenerate Unconditionally due to Heartbeat Abnormality

In an operation that perform automatic degenerate unconditionally due to heartbeat abnormality, you can optimize the time from operating system or server abnormality detection to automatic degradation by editing the operating system or server abnormality monitoring parameters in the server configuration file of the database server. Refer to "Parameters for the abnormality monitoring of the operating system or server in the server configuration file of the database server" for information on the operating system or server abnormality monitoring parameters in the server configuration file of the database server. In addition, when the Mirroring Controller detects an error, it does not perform the arbitration processing, fencing, or automatic degradation, but only automatic degenerate unconditionally is performed.



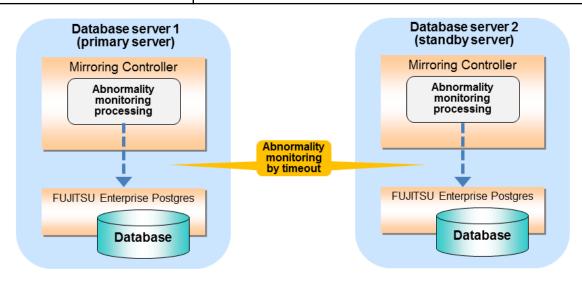
Refer to "Appendix D Notes on Performing Automatic Degradation Immediately after a Heartbeat Abnormality" for notes on the operation that perform automatic degenerate unconditionally due to heartbeat abnormality.

### 2.11.4.2 Tuning for Abnormality Monitoring of Darabase Processes

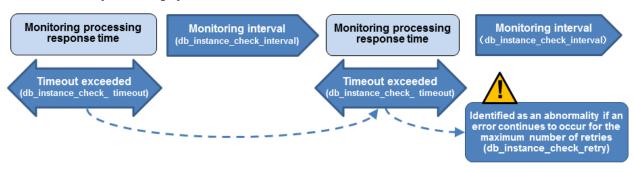
In an abnormality monitoring of database pocesses, you can optimize by editing the following parameters in the server configuration file of the database server.

Table 2.14 Parameters for abnormality monitoring of database processes

Parameter	Description
Abnormality monitoring interval (db_instance_check_interval)	Abnormality monitoring by Mirroring Controller is set so as not to place load on the system, but normally it does not need to be set. (The default is the value set in heartbeat_interval.) (milliseconds)
Timeout for abnormality monitoring of database processes (db_instance_check_timeout)	Take into account the time during which a load is placed continuously on the database. For example, it is envisaged that this parameter will be used in situations such as when performing high-load batch jobs or when a large number of online jobs occur continuously and concurrently. (The default is the value set in heartbeat_timeout.) (seconds)
Abnormality monitoring retries (db_instance_check_retry)	This parameter can be set when needing a safety value for situations in which the value specified for db_instance_check_timeout is exceeded, for example, when using systems with fluctuating loads, however, this parameter does not normally need to be set. (The default is the value set in heartbeat_retry.) (number of times)



#### Flow of abnormality monitoring by timeout



The expression for calculating the time required to detect an abnormality is shown below.

```
Abnormality detection time = ( db_instance_check_timeout(seconds) + db_instance_check_interval(milliseconds) / 1000 ) x ( db_instance_check_retry(number of times) + 1 )
```

The abnormality detection time when the default value is used is shown below.

```
Abnormality detection time = ( 1 + 800 / 1000 ) x ( 2 + 1 ) = 5.4(seconds)
```



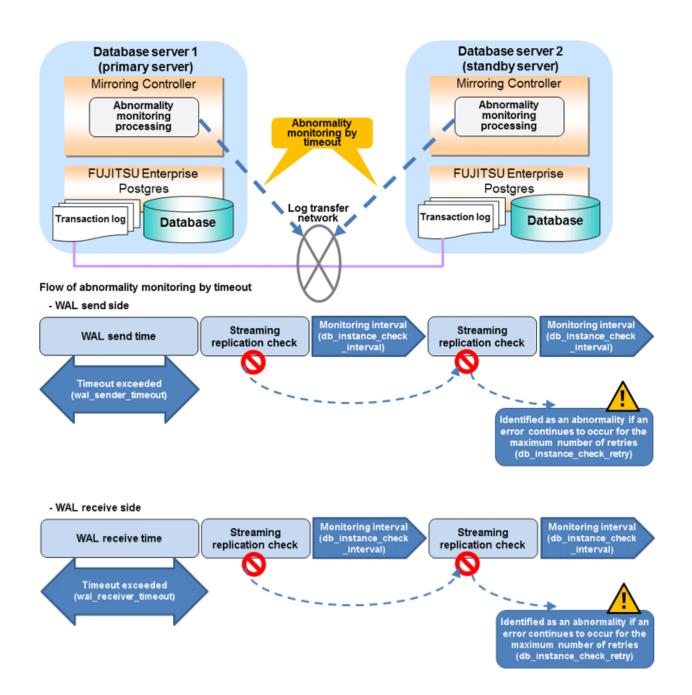
- If the db\_instance\_timeout\_action parameter in *serverIdentifier*.conf is set to "message", and the db\_instance\_check\_timeout parameter is set to a short value, a crash of the database process will be detected as "no response", and it may take time for automatic degradation to occur. Therefore, specify an appropriate timeout for db\_instance\_check\_timeout.
- If a high load on the database and an event that prevents connection to an instance occur at the same time, it is judged as abnormal without retrying monitoring.

## 2.11.4.3 Tuning for Abnormality Monitoring of Streaming Replication

In an abnormality monitoring of streaming replication, you can optimize by editing the following parameters in the server configuration file of the database server.

Table 2.15 Parameters for abnormality monitoring of streaming replication

Parameter	Description	
Abnormality monitoring interval (db_instance_check_interval)	Abnormality monitoring by Mirroring Controller is set so as not to place load on the system, but normally it does not need to be set. (The default is the value set in heartbeat_interval.) (milliseconds)	
Abnormality monitoring retries (db_instance_check_retry)	This parameter can be set when needing a safety value, such as when it is anticipated that a temporary log transfer LAN error may occur, but it does not normally need to be set. (The default is the value set in heartbeat_retry.) (number of times)	
Timeout for abnormality monitoring of streaming replication (wal_sender_timeout and wal_receiver_timeout in postgresql.conf)	Take into account the capacity and load of the log transfer network and the time during which a load is placed continuously on the database. For example, if there is a succession of data update jobs that generate a high WAL volume, you must configure the settings to avoid misdetection. (The default is 60 seconds.)	



The expression for calculating the time required to detect an abnormality is shown below.

The abnormality detection time when the default value is used is shown below.

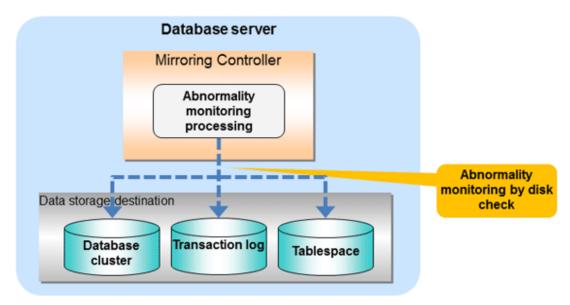
```
Abnormality detection time = 60 + (800 / 1000 \times (2 + 1))
= 62.4(seconds)
```

### 2.11.4.4 Tuning for Disk Abnormality Monitoring

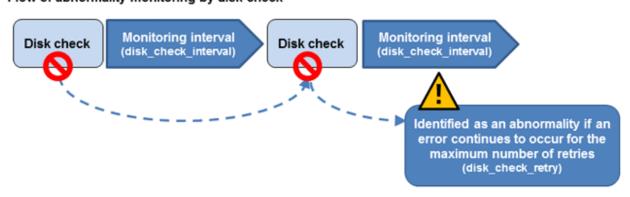
In an abnormality monitoring of the disk, you can optimize by editing the following parameters in the server configuration file of the database server.

Table 2.16 Parameters for disk abnormality monitoring

Parameter	Description	
Abnormality monitoring interval (disk_check_interval)	Abnormality monitoring by Mirroring Controller is set so as not to place load on the system, but normally it does not need to be set. (The default is the value set in heartbeat_interval.) (milliseconds)	
Abnormality monitoring retries (disk_check_retry)	This parameter can be set when needing a safety value, such as when it is anticipated that a temporary disk input/output error may occur, but normally it does not need to be set. (The default is the value set in heartbeat_retry.) (number of times)	



#### Flow of abnormality monitoring by disk check



The expression for calculating the time required to detect an abnormality is shown below.

```
Abnormality detection time = disk\_check\_interval (milliseconds) / 1000 x ( disk\_check\_retry(number of times) + 1 )
```

The abnormality detection time when the default value is used is shown below.

```
Abnormality detection time = 800 / 1000 \times (2 + 1)
= 2.4(seconds)
```



- The tuning described above impacts on the time taken from detection of a timeout until switching the primary server. Therefore, modify the values while taking into account the switch/disconnection time, using a design for which misdetection does not occur.
- Immediately selecting automatic degradation when a heartbeat abnormality occurs in operating system or server heartbeat monitoring risks causing split brain. Refer to "Appendix D Notes on Performing Automatic Degradation Immediately after a Heartbeat Abnormality" for details.



Mirroring Controller uses connections to database instances and SQL access to monitor abnormality in some resources targeted for monitoring. The connection destination database names and connection user names used for abnormality monitoring conform to the parameters in the server configuration file. The application name is "mc\_agent".

# 2.12 Setting Automatic Start and Stop of Mirroring Controller and Multiplexed Instances

Multiplexed instances and Mirroring Controller can be started and stopped automatically in line with the starting and stopping of the operating system of the database server.



To guarantee the startup sequence of Mirroring Controller on the primary and standby servers, first confirm that the primary server has started, and then start the standby servers in sequence.

The startup sequence of the Mirroring Controller process on the database server and the Mirroring Controller arbitration process on the arbitration server is not guaranteed. If the arbitration server cannot be started first, execute the mc\_ctl command in start mode with the --async-connect-arbiter option specified to start the Mirroring Controller process.

#### If using an rc script

Perform the following procedure:

1. Set the OS user account of the instance administrator user

If you logged in using the OS user account of the instance administrator user, set the environment variables required for starting the instance.

This setting is required for executing the "su -" command.

2. Edit the shell script

Copy the sample shell script, and edit the environment variables in the shell script to suit the environment.

Sample shell script

/installDir/share/rc\_mcoi.sample

3. Copy the revised shell script and set access permissions

As the OS superuser, copy the revised shell script to /etc/rc.d/init.d, and then set and register the access rights.

In the following example, the shell script was created as "rc\_mc\_inst1".

Example)

```
# cp rc_mc_instl /etc/rc.d/init.d/
# chmod 755 /etc/rc.d/init.d/rc_mc_instl
```

#### 4. Register and enable automatic start and stop

As the OS superuser, execute the chkconfig command to register and enable the script. To register and enable the script, specify the options shown in the table below.

Option	Description	
add	Registers the shell script as the startup script.	
level Sets and enables the startup script for each run level.		

In the following example, the shell script is registered, and then enabled for run levels 3 and 5.

Example)

```
# chkconfig --add rc_mc_inst1
# chkconfig --level 35 rc_mc_inst1 on
```

#### If using systemd



If you use systemd to start the Mirroring Controller and multiplexed instances, wait for time correction, network setup, and so on.

#### Perform the following procedure:

#### 1. Create a unit file

Copy the unit file sample stored in the directory below, and revise it to match the target instance.

Sample file

/installDir/share/mcoi.service.sample

#### Example)

In the following example, the installation directory is "/opt/fsepv<x>server64", and the instance name is "inst1". Note that "<x>" indicates the product version.

```
# cp /opt/fsepv<x>server64/share/mcoi.service.sample /usr/lib/systemd/system/
mcoi_inst1.service
```

Revise the underlined portions of the options below in the unit file.

Section	Option	Specified value	Description
Unit	Description	FUJITSU Enterprise Postgres MirroringController <u>instanceName</u>	Specifies the feature overview.  Specifies the name of the target instance. (*1)
Service	ExecStart	/bin/bash -c ' <u>installDir</u> /bin/mc_std start <u>installDir</u> <u>MirroringControllerManagementDi</u> <u>r' mc_ctlOption'</u>	Command to be executed when the service is started.  Specify the option you want to add when the mc_ctl command is executed without the -M option in the mc_ctl option.  Note that the content specified in this mc_ctl option is carried over from the mc_std command to the mc_ctl command. (*2)
	ExecStop	/bin/bash -c ' <i>installDin</i> /bin/mc_std stop <i>installDir</i>	Command to be executed when the service is stopped.

Section	Option	Specified value	Description
		MirroringControllerManagementDi r'mc_ctlOption	Specify the option you want to add when the mc_ctl command is executed without the -M option in the mc_ctl option.  However, to use themc-only option to stop only Mirroring Controller, you must use themc-only option at startup.  Note that the content specified in this mc_ctl option is carried over from the mc_std command to the mc_ctl command.
	User	<u>User</u>	OS user account of the instance administrator user.
	Group	Group	Group to which the instance administrator user belongs.

<sup>\*1:</sup> The instance name should be as nameThatIdentifiesTheInstance.

The naming conventions for identifying the instance are as follows:

- Up to 16 bytes
- The first character must be an ASCII alphabetic character
- The other characters must be ASCII alphanumeric characters
- \*2: When the arbitration server is used for automatic degradation, start the Mirroring Controller arbitration process on the arbitration server and then start the Mirroring Controller process on the database server. If the arbitration server cannot be started first, specify the --async-connect-arbitration to start the Mirroring Controller process.

#### 2. Enable automatic start and stop

As the OS superuser, use the systemctl command to enable automatic start and stop.

Example)

# systemctl enable mcoi\_instl.service



If automatic start and stop of Mirroring Controller has been configured using systemd, to stop Mirroring Controller, do not use the mc\_ctl command, but instead use the systemctl command as the OS superuser.

#### Example)

# systemctl stop mcoi\_instl.service

If the instance does not stop, refer to "Actions in Response to Failure to Stop an Instance" in the Operation Guide to stop the instance. Then, specify the -e option in the mc\_ctl command to forcibly stop Mirroring Controller.

#### Example)

\$ mc\_ctl stop -M /mcdir/inst1 -e

If Mirroring Controller is stopped using the mc\_ctl command, the message below is output to the system log, however there is no issue because automatic stop is executed by systemd.

#### Message

FATAL: failed to stop Mirroring Controller target server:" $\{0\}$ " (MCA00043)

# 2.13 Setting Automatic Start and Stop of the Mirroring Controller Arbitration Process

You can automatically start or stop the Mirroring Controller arbitration process when the operating system on the arbitration server is started or stopped.

#### Linux

If using an rc script

Perform the following procedure:

1. Set the operating system user.

Configure the environment variables required for starting the Mirroring Controller arbitration process.

2. Edit the shell script.

Copy the sample shell script, and edit the environment variables in the shell script to suit the environment.

Sample shell script

```
/installDir/share/rc_mcarboi.sample
```

3. Copy the revised shell script and set access permissions.

Copy the revised shell script to /etc/rc.d/init.d, and then set and register the access rights.

```
Example)
```

In the following example, the shell script was created as " $rc_mcarb_arbiterl$ ".

```
# cp rc_mcarb_arbiter1 /etc/rc.d/init.d/
# chmod 755 /etc/rc.d/init.d/rc_mcarb_arbiter1
```

4. Register and enable automatic start and stop.

As the operating system superuser, execute the chkconfig command to register and enable the script. To register and enable the script, specify the options shown in the table below.

Option	Description	
add	Registers the shell script as the startup script.	
level	Sets and enables the startup script for each run level.	

#### Example)

In the following example, the shell script is registered, and then enabled for run levels 3 and 5.

```
# chkconfig --add rc_mcarb_arbiter1
# chkconfig --level 35 rc_mcarb_arbiter1 on
```

## If using systemd



If you use systemd to start the Mirroring Controller arbitration process, wait for time correction, network setup, and so on.

Perform the following procedure:

1. Create a unit file.

Copy the unit file sample stored in the directory below, and revise it to match the target instance.

/installDir/share/mcarboi.service.sample

### Example)

In the following example, the installation directory is "/opt/fsepv<x>assistant", and the identifier of the arbitration process is "arbiterl". Note that "<x>" indicates the product version.

# cp /opt/fsepv<x>assistant/share/mcarboi.service.sample /usr/lib/systemd/system/
mcarboi\_arbiter1.service

Revise the underlined portions of the options below in the unit file.

Section	Option	Specified value Description	
Unit	Description	FUJITSU Enterprise Postgres Mirroring Controller Arbiter <arbitrationprocessid></arbitrationprocessid>	Specifies the feature overview.
			Specifies the identifier of the targeted arbitration process. (*1)
Service	ExecStart	/bin/bash - c ' <u>installDir</u> /bin/mc_arb_std start <u>installDir</u> <u>mirroringControllerArbitrationProcessMgmtDir</u> <u>mc_arbOption</u> '	Command to be executed when the service is started.
			Specify the option you want to add when the mc_arb command is executed without the -M option in the mc_arb option.
			Note that the content specified in this mc_arb option is carried over from the mc_arb_std command in "Specified value" to the mc_arb command.
	ExecStop	/bin/bash - c ' <u>installDir</u> /bin/ mc_arb_std stop <u>installDir</u> <u>mirroringControllerArbitrationProcessMgmtDir</u> <u>mc_arbOption</u> '	Command to be executed when the service is stopped.
			Specify the option you want to add when the mc_arb command is executed without the -M option in the mc_arb option.
			Note that the content specified in this mc_arb option is carried over from the mc_arb_std command in "Specified value" to the mc_arb command.
	User	<u>User</u>	Specify the account of the operating system user.

Section	Option	Specified value	Description
	Group	Group	Specify the group to which the user belongs.

<sup>\*1:</sup> The arbitration process identifier used here is a name for identifying the Mirroring Controller arbitration process.

The naming conventions for identifying the Mirroring Controller arbitration process are as follows:

- Up to 16 bytes
- The first character must be an ASCII alphabetic character
- The other characters must be ASCII alphanumeric characters
- 2. Enable automatic start and stop.

As the operating system superuser, use the systemctl command to enable automatic start and stop.

Example)

# systemctl enable mcarboi\_arbiter1.service

#### Windows

You can configure the Windows service to perform automatic start and stop.

### Configuring during setup

When registering the Mirroring Controller arbitration process as a Windows service in "2.3.1 Configuring the Arbitration Server", specify "auto" for the -S option of the register mode used with the mc\_arb command.

Example)

> mc\_arb register -M D:\mcdir\inst1 -P \*\*\*\*\*\*\* -S auto

## Changing the configuration after setup

Use the sc config command to change the configuration of the Windows service of the Mirroring Controller arbitration process.

Example)

The configuration of the registered service name "Mirroring\_Controller\_Arbiter1" is changed.

> sc config "Mirroring\_Controller\_Arbiter1" start= auto



#### See

Refer to sc command help for information on how to configure the service.



# Information

You can check the registration status in the Windows service window or by using the sc qc command.

# 2.14 Backup Operation

This section explains the backup operation for database multiplexing mode.

# 2.14.1 Backing up Database Multiplexing Mode Information

When changing the Mirroring Controller settings, in addition to backing up the database, back up the configuration file in the Mirroring Controller management directory so that the Mirroring Controller settings are not lost.

When the arbitration server is used for automatic degradation, also back up the configuration file in the Mirroring Controller arbitration process management directory.

# 2.14.2 Database Backup Operation

Using database multiplexing mode is the same as obtaining the backup data on the standby server as a safeguard against a disk failure. Note that all server disks may be corrupted due to some cause.

As a safeguard against this type of case, execute the pgx\_dmpall command on the primary server to create the backup data.

However, it is not definite as to which server runs as the primary server, so ensure that the pgx\_dmpall command is executed periodically on all servers, so that the backup data will be obtained. For example, create a script to obtain the backup data, and set it in the operation management software.



When the pgx\_dmpall command is executed on the standby server, it will not match the statuses, however the error message shown below will be output and return the value "1".

.....

If a script that ignores only this type of error is executed on all servers, the backup data of the primary server can be obtained.

Error message

ERROR: recovery is in progress (10095)



- Consider the possibility that the server that runs as the primary server may be destroyed alongside the backup data, so it is recommended to promote another server to become the primary server, and then back up the data on the new primary server without waiting for the next scheduled backup.

......

- Specify the same backup directory name for the primary and standby servers. If different backup directory names are specified, and recovery is performed using the backup data of the other server, the recovery cannot be performed correctly.



- Period backups allow shorter recovery time and reduction in disk usage. Refer to "Backing Up the Database" in the Operation Guide for details on the backup operation.
- Refer to "Chapter 4 Action Required when an Error Occurs in Database Multiplexing Mode" for details on recovery based on the backup data that was obtained using the pgx\_dmpall command.

# Chapter 3 Operations in Database Multiplexing Mode

This chapter describes the periodic operations that are performed when running database multiplexing mode.

The periodic operations are the same as the operations on a single server.



See

Refer to "Periodic Operations" in the Operation Guide for information on the periodic operations.

# 3.1 Starting and Stopping the Mirroring Controller Arbitration Process

This section describes how to start and stop the Mirroring Controller arbitration process.

# 3.1.1 Starting the Mirroring Controller Arbitration Process

#### Linux

While the Mirroring Controller arbitration process is in a stopped state, execute the mc\_arb command in start mode to start the Mirroring Controller arbitration process.

......

Example)

\$ mc\_arb start -M /mcarb\_dir/arbiter1

#### Windows

The Mirroring Controller arbitration process can be started using one of the following options:

- Using the mc\_arb command
- Starting the service on system startup

## Using the mc\_arb command

While the Mirroring Controller arbitration process is in a stopped state, execute the mc\_arb command from the command prompt to start the Mirroring Controller arbitration process.

Example)

> mc\_arb start -M D:\mcarb\_dir\arbiter1

### Starting the service on system startup

Specify automatic start when registering the Mirroring Controller arbitration process to the Windows service during setup of database multiplexing mode. Accordingly, the Mirroring Controller arbitration process service will start on startup of the operating system.



Refer to the Reference for information on how to specify the mc\_arb command.

# 3.1.2 Stopping the Mirroring Controller Arbitration Process

#### Linux

While the Mirroring Controller arbitration process is running, execute the mc\_arb command in stop mode to stop the Mirroring Controller arbitration process.

Example)

```
$ mc_arb stop -M /mcarb_dir/arbiter1
```

#### Windows

The Mirroring Controller arbitration process can be stopped using one of the following options:

- Using the mc\_arb command
- Stopping the service

## Using the mc\_arb command

While the Mirroring Controller arbitration process is running, execute the mc\_arb command in stop mode from the command prompt to stop the Mirroring Controller arbitration process.

Example)

> mc\_arb stop -M D:\mcarb\_dir\arbiter1

## Stopping the service

Select [Administrative Tools], then [Services] to open the [Services] window, and then select the Mirroring Controller service and click the [Stop] menu.



#### See

Refer to the Reference for information on how to specify the mc\_arb command.



- The arbitration server will be forcibly stopped when the service is stopped.
- Before shutting down the operating system on the arbitration server, either stop the Mirroring Controller on the primary server or standby server or shut down the operating system on the primary server or standby server.

# 3.2 Starting and Stopping Mirroring Controller

When database multiplexing mode is used, use the mc\_ctl command to start and stop the instance and Mirroring Controller at the same time.

Do not start or stop the instance by itself.

## **Starting Mirroring Controller**

While Mirroring Controller is in a stopped state, execute the mc\_ctl command in start mode to start Mirroring Controller.

Enabling automatic switch/disconnection

Execute the mc\_ctl command in start mode.

Example)

\$ mc\_ctl start -M /mcdir/inst1

When only the instance is started and stopped, the following will happen:

- When only the instance is started

Features such as automatic switch and automatic disconnection will not work until Mirroring Controller is started.

- When only the instance is stopped

Mirroring Controller determines that an error has occurred in the instance, and performs an unnecessary automatic switch.

Automatic switch may also stop working correctly in some cases.

## Disabling automatic switch/disconnection

Execute the mc\_ctl command in start mode with the -F option specified.

Example)

\$ mc\_ctl start -M /mcdir/inst1 -F

When only the instance is started and stopped, the following will happen:

- When only the instance is started

Errors indicated in "1.1 What is Database Multiplexing Mode" will not be detected until Mirroring Controller is started.

- When only the instance is stopped

Mirroring Controller determines that an error has occurred in the instance, and outputs an error to the system log.



- To start the Mirroring Controller process only, execute the mc\_ctl command in start mode with the --mc-only option specified.
- After Mirroring Controller is started, automatic switch/disconnection can be enabled or disabled using the enable-failover or disable-failover mode of the mc\_ctl command.
- When the arbitration server is used for automatic degradation, the Mirroring Controller process startup fails on the database server if the Mirroring Controller arbitration process has not been started on the arbitration server in advance. However, even if the Mirroring Controller arbitration process cannot be started in advance, the Mirroring Controller process can be started by specifying the --async-connect-arbitration in the mc\_ctl command.



- When the arbitration server is used for automatic degradation, the database server must connect to the arbitration server, and as a result, Mirroring Controller startup may take longer.

.....

- Mirroring Controller startup usually fails if the standby server is mistakenly started as the primary server or if the old primary server is not recovered after the switch and is then mistakenly started as the primary server. However, if the admin network is disconnected, then startup does not fail, and both servers may become primary servers. Therefore, ensure that the admin network is connected before starting Mirroring Controller.

### **Stopping Mirroring Controller**

While Mirroring Controller is running, execute the mc\_ctl command in stop mode to stop Mirroring Controller process.

Example)

\$ mc\_ctl stop -M /mcdir/inst1



To stop the Mirroring Controller process only, execute the mc\_ctl command in stop mode with the --mc-only option specified.



To prevent an unintended automatic switch, before shutting down the operating system on the primary server, you must stop the Mirroring Controller, or shut down the operating system on the standby server.



Refer to the Reference for information on how to specify the mc\_ctl command.

# 3.3 Checking the Database Multiplexing Mode Status

# 3.3.1 Checking the Status of the Database Server

This section describes how to check the status of the database server.

Check the multiplexed database status by executing the mc\_ctl command in status mode.

Additionally, errors can be detected by monitoring the Mirroring Controller messages. If the status or messages are monitored periodically, you can react quickly following an automatic switch failure.

## Checking the status of the multiplexing database

When the mc\_ctl command is executed, the details of the multiplexing configuration, information about whether switch is possible following the error, and location and details of the error that caused the switch or disconnection are displayed.

After starting database multiplexing mode, execute the mc\_ctl command in status mode to check the multiplexing status.

An example of the status displayed when the mc\_ctl command is executed is shown below.

Example)

```
$ mc ctl status -M /mcdir/inst1
mirroring status
switchable
server_id host_role
                                            host_status db_proc_status
          primary
                           192.0.2.100
server1
                                            normal
                                                         normal
                                                                                 normal
          standby
                           192.0.2.110
server2
                                            normal
                                                         normal
                                                                                 normal
```

## Checking the status of connection to the Mirroring Controller arbitration process

When the arbitration server is used for automatic degradation, the status of the connection to the Mirroring Controller arbitration process can be checked by specifying the --arbiter option. If the output status is "online", it indicates that an arbitration request can be made from the database server to the arbitration server. When the arbitration server is used for automatic degradation, regularly execute the command in status mode with the --arbiter option specified and check that the output status is "online".

```
Example)
```

The mc\_ctl command is executed with the --arbiter option specified, and the status is output.

```
$ mc_ctl status --arbiter -M /mcdir/inst1
```

```
arbiter_id host status
------
arbiter 192.0.3.120 online
```

## Checking the status of data synchronization

Additionally, by referencing the pg\_stat\_replication statistics view on the primary server, the data synchronization status can be confirmed. However, when creating the monitoring program, note that the content of pg\_stat\_replication may be changed in the future.

The following example shows that the locations of the transaction log after it is sent and received (sent\_lsn, replay\_lsn) match, and that they are fully synchronized.

Example)

```
postgres=# select * from pg_stat_replication;
-[ RECORD 1 ]---+---
pid
             | 10651
usesysid
             | 10
          fsep
usename
application_name | standby
client_addr | 192.0.2.210
client_hostname
             | 55098
client_port
backend_start | 2020-03-23 11:17:49.628793+09
backend_xmin
state
              streaming
sent_lsn
              0/3000060
write_lsn
              0/3000060
flush_lsn
              0/3000060
replay_lsn
              0/3000060
write_lag
flush_lag
replay_lag
sync_priority
              | 1
sync_state
              sync
reply_time
              2020-03-23 11:23:27.703366+09
```



- Refer to "mc\_ctl" in Reference for information on the command.
- Refer to "Notes on Application Compatibility" in the Application Development Guide for information on retaining application compatibility.

.....

- Refer to "The Statistics Collector" in "Server Administration" in the PostgreSQL Documentation for details on pg\_stat\_replication.

# 3.3.2 Checking the Status of the Arbitration Server

This section describes how to check the status of the arbitration server.

The status of the connection between the Mirroring Controller arbitration process and primary server/standby server can be checked by executing the mc\_arb command in status mode.

The example below executes the mc\_arb command, and shows the status.

Linux

Example)

server1	192.0.3.100 online
server2	192.0.3.110 online

#### Windows

Example)

# 3.4 Manually Switching the Primary Server

The primary server cannot be switched automatically in the following case:

- If automatic switch/disconnection is disabled
- If output of messages is selected for heartbeat abnormalities during heartbeat monitoring of the operating system or server and the operating system/server crashes or becomes unresponsive

In this case, to manually switch the primary server, execute the mc\_ctl command in switch mode on either the primary server or the standby server.

Example)

\$ mc\_ctl switch -M /mcdir/inst1



If automatic switch/disconnection is enabled, it is possible to perform switch of primary server at any time.

# 3.5 Manually Disconnecting the Standby Server

The procedure to perform disconnection of the standby server differs depending on whether the automatic switch/disconnection is enabled or disabled.

If automatic switch/disconnection is enabled

Execute the mc\_ctl command in stop mode on the standby server.

Example)

```
$ mc_ctl stop -M /mcdir/inst1
```

If automatic switch/disconnection is disabled

1. Execute the mc\_ctl command in stop mode on the standby server.

Example)

```
$ mc_ctl stop -M /mcdir/inst1
```

- 2. Comment out the synchronous\_standby\_names parameter in the postgresql.conf file on the primary server.
- 3. Execute the pg\_ctl command in reload mode on the primary server.

Example)

```
$ pg_ctl reload -D /database/inst1
```



If automatic start and stop of Mirroring Controller has been configured using systemd, do not use the mc\_ctl command, but instead use the systemctl command. Refer to "2.12 Setting Automatic Start and Stop of Mirroring Controller and Multiplexed Instances" for details.

# 3.6 Action Required when a Heartbeat Abnormality is Detected

The message below is output when a heartbeat abnormality is detected during heartbeat monitoring of operating systems or servers:

detected an error on the monitored object "server(server identifier name)": no response:ping timeout (MCA00019)

If the heartbeat\_error\_action parameter in *serverIdentifier*.conf is set to "message", even if automatic switch/disconnection is enabled and Mirroring Controller is started, automatic switch/disconnection is not performed when a heartbeat abnormality is detected. Therefore, user action will be necessary.

This section explains the action required when the heartbeat\_error\_action parameter is set to "message" and a heartbeat abnormality is detected.

- 1. Identify the cause of the heartbeat abnormality. The possible causes are below:
  - The remote operating system or server crashed or is unresponsive
  - An admin network issue occurred
- 2. Address the cause identified in step 1.
  - The remote operating system or server crashed or is unresponsive
     Manually perform switch or disconnection using the mc\_ctl command.
  - An admin network issue occurred

Refer to "Chapter 4 Action Required when an Error Occurs in Database Multiplexing Mode", and recover the database multiplexing system.

# 3.7 Monitoring Mirroring Controller Messages

The messages that are output by Mirroring Controller are output to both the database server and the arbitration server. If the automatic switch fails, for example, an important message related to the continuation of the operation may be output, so ensure that the system log messages are monitored.

If the arbitration server is used for automatic degradation, monitor messages on both the database server and the arbitration server.

Message output destination on the database server

Messages are output to the system log.

Message output destination on the arbitration server

Linux

Messages are output to the system log.

Windows

Messages are output to the event log.



- To monitor message types considered to be important, an operating system setting must be configured beforehand. Refer to the operating system manuals, check if the message is of a message type that is monitored to be output to the system log, and configure the setting if required.

- If the heartbeat\_error\_action parameter in *serverIdentifier*.conf is set to "message", only message output is performed when a heartbeat abnormality is detected during heartbeat monitoring of operating systems and servers - automatic switch/disconnection is not performed. Therefore users need to monitor the messages. Refer to "3.6 Action Required when a Heartbeat Abnormality is Detected" for details.

Display format on the database server

programName[processId]: messageType:messageText (messageNumber)

Specify the program name in the syslog\_ident parameter of the serverIdentifier.conf file of the database server.

The message types output by Mirroring Controller, their severity, and their corresponding value in the system log are shown in the table below.

Table 3.1 Message type, severity, and corresponding value in the system log

Message type	Severity	Meaning	System log
INFO	Information	Provides information that does not fall under LOG or NOTICE. INFO	
LOG		Provides information recognized as a particularly important event in tracing the operation history.  (Example: Automatic switch is complete)	
NOTICE	Notice	Outputs information that takes into account the user instructions within the program in response to an executed or automatically executed process.	
WARNING	Warning	Provides a warning, for example it will soon be impossible to maintain the multiplexing state.  WARNII  WARNII	
ERROR	Error	Reports that an error other than FATAL or PANIC has occurred.	ERROR
FATAL		Reports that an abnormality was detected in multiplexed database systems requiring recovery of the system, and also the content and cause of the abnormality.	CRIT
PANIC		Reports that an abnormality was detected in all multiplexed database systems requiring immediate recovery of the system, and also the content and cause of the abnormality.	ALERT

The message severity has the following meanings:

- Information

Informational status. A message that was reported by the system is displayed. No action is required.

- Notice

Informational status, but a message that should be noted is displayed. If necessary, take the actions described in the "Action" section of the message.

- Warning

No error has occurred, but the user is requested to check, and take action. Take the actions described in the "Action" section of the message.

- Error

An error has occurred. Take the actions described in the "Action" section of the message.

## Display format on the arbitration server

Linux

programName[processId]: messageType: messageText (messageNumber)

Specify the program name in the syslog\_ident parameter of the arbitration.conf file of the arbitration server.

### Windows

```
eventSourceName[processId]: messageType: messageText (messageNumber)
```

Specify the event source name in the event\_source parameter of the arbitration.conf file of the arbitration server.

The message types output by Mirroring Controller, their severity, and their corresponding value in the output destination log are shown in the table below.

Table 3.2 Message type, severity, and corresponding value in the output destination log

Message type	Severity	Meaning	System log (Linux)	Event log (Windows)
INFO	Information	Provides information not categorized as LOG or NOTICE.	INFO	INFORMATION
LOG		Provides information recognized as a particularly important event in tracing the operation history.  (Example: Automatic switch is complete)		
NOTICE	Notice	Outputs information that takes into account the user instructions within the program in response to an executed or automatically executed process.	NOTICE	
WARNING	Warning	Provides a warning, for example it will soon be impossible to perform the arbitration process.	WARNING	WARNING
ERROR	Error	Reports that an error other than FATAL or PANIC has occurred.	ERROR	ERROR
FATAL		Reports that an abnormality was detected in the arbitration server requiring recovery of the system, and also the content and cause of the abnormality.	CRIT	
PANIC		Reports that an abnormality was detected in the arbitration server requiring immediate recovery of the system, and also the content and cause of the abnormality.	ALERT	

The message severity has the following meanings:

- Information

Informational status. A message that was reported by the system is displayed. No action is required.

- Notice

Informational status, but a message that should be noted is displayed. If necessary, take the actions described in the "Action" section of the message.

- Warning

No error has occurred, but the user is requested to check, and take action. Take the actions described in the "Action" section of the message.

- Error

An error has occurred. Take the actions described in the "Action" section of the message.

# 3.8 Server Maintenance

To perform maintenance tasks such as periodic server inspections and the application of updates for software products including the operating system, you must perform a planned stop of the server, and then perform the maintenance.

# 3.8.1 Rolling Updates

In database multiplexing mode, rolling updates, that perform the maintenance for the servers that comprise the cluster system, can be performed while jobs continue.

First, perform the maintenance for the standby server, and then switch the standby server to the primary server. Then, perform the maintenance for the original primary server that was switched to the standby server. This enables maintenance to be performed while jobs continue.

Note that arbitration server maintenance can be performed without affecting database server operation, so it is not necessary to consider rolling update.



If the downtime due to the maintenance of the standby server is expected to be long, refer to "Standby server downtime" in "3.9.1 Changes Required when the Standby Server is Stopped".

The flow of a rolling update is shown below.

System administrator Application Maintenance tasks on the standby serv Standby server Primary server Database Database Change the server role to 'primary server' Application Standby server Primary server Database Database System administrator Application Standby server Primary server Data Database Maintenance tasks on the new standby server Application System administrator Standby server Primary server Database Database Failback to the Application Operation from the primary server (only application

Figure 3.1 Performing a Rolling Update

Perform the following procedure as shown in the above figure:

Primary server

Database

# Standby server maintenance tasks

1. To perform the maintenance on the standby server, stop Mirroring Controller.

Example)

if required)

\$ mc\_ctl stop -M /mcdir/inst1

Standby server

Database

System status transition

server

Operation to each

2. Ensure that Mirroring Controller has completely stopped.

If the multiplexed instances and Mirroring Controller have been configured on the standby server to start and stop automatically when the operating system of the database server is started or stopped, cancel the setting to start and stop automatically.



#### See

Refer to "2.12 Setting Automatic Start and Stop of Mirroring Controller and Multiplexed Instances" for information on how to configure the multiplexed instances and Mirroring Controller to start and stop automatically when the operating system of the database server start and stops.

If using an rc script

As the OS superuser, execute the chkconfig command to disable automatic start and stop.

In the example below, the shell script for automatic start and stop was created as "rc\_mc\_inst1" and is disabled for run levels 3 and 5.

Example)

```
# chkconfig --level 35 rc_mc_inst1 off
```

#### If using systemd

As the OS superuser, execute the systemctl command to disable automatic start and stop.

The example below disables automatic start and stop of "mcoi\_inst1.service".

Example)

```
# systemctl disable mcoi_instl.service
```

- 3. Perform maintenance tasks.
- 4. Create a copy of the primary server instance on the standby server.

Execute the pg\_basebackup command to create data in the standby server by synchronizing with the primary server.

Example)

```
$ pg_basebackup -D /database/inst1 -X fetch --waldir=/transaction/inst1 --progress --verbose -R
--dbname='application_name=standbyServerName' -h primaryServerHostName -p
primaryServerPortNumber
```



#### See

The procedure for copying the primary server instance to the standby server is the same as the procedure for setting up the standby server.

.....

Refer to "2.5.2 Creating, Setting, and Registering the Standby Server Instance", and then perform the recovery.

5. Check the settings for automatic start and stop of the multiplexed instances and Mirroring Controller.

If the multiplexed instances and Mirroring Controller were configured in step 2 to not start and stop automatically when the operating system of the database server starts and stops, then change the settings back. This step can be skipped if automatic start and stop are not required.

If using an rc script

As the OS superuser, execute the chkconfig command to enable the shell script.

In the example below, the shell script for automatic start and stop was created as "rc\_mc\_inst1" and is enabled for run levels 3 and 5.

Example)

```
# chkconfig --level 35 rc_mc_inst1 on
```

#### If using systemd

As the OS superuser, execute the systemctl command to enable automatic start and stop.

The example below disables automatic start and stop of "mcoi\_inst1.service".

Example)

```
# systemctl enable mcoi_inst1.service
```

6. Start (rebuild) Mirroring Controller on the standby server.

This operation is required when determining the maintenance tasks on the standby server.

Enabling automatic switch/disconnection

As the instance administrator user, execute the mc\_ctl command in start mode.

Example)

```
$ mc_ctl start -M /mcdir/inst1
```

#### Disabling automatic switch/disconnection

As the instance administrator user, execute the mc\_ctl command in start mode with the -F option specified.

Example)

```
$ mc_ctl start -M /mcdir/inst1 -F
```



After Mirroring Controller is started, automatic switch/disconnection can be enabled or disabled using the enable-failover or disable-failover mode of the mc\_ctl command.

## Switching to the primary server

To perform the maintenance on the primary server, execute the mc\_ctl command in the switch mode on the primary server or the standby server.

Example)

```
$ mc_ctl switch -M /mcdir/inst1
```

When the switch is complete, the synchronous\_standby\_names parameter in the postgresql.conf file of the new primary server will be commented as follows:

Example)

```
#synchronous_standby_names = 'primary'
```

## New standby server maintenance tasks

1. Stop the Mirroring Controller.

On the new standby server (the primary server before the switch), execute the mc\_ctl command in stop mode.

If automatic start and stop of Mirroring Controller has been configured using systemd, do not use the mc\_ctl command, but instead use the systemctl command. Refer to "2.12 Setting Automatic Start and Stop of Mirroring Controller and Multiplexed Instances" for details.

Example)

```
$ mc_ctl stop -M /mcdir/inst1
```

2. Ensure that Mirroring Controller has completely stopped.

If the multiplexed instances and Mirroring Controller have been configured on the new standby server to start and stop automatically when the operating system of the database server is started or stopped, cancel the setting to start and stop automatically now.



#### See

Refer to "2.12 Setting Automatic Start and Stop of Mirroring Controller and Multiplexed Instances" for information on how to configure the multiplexed instances and Mirroring Controller to start and stop automatically when the operating system of the database server starts and stops.

#### If using an rc script

As the OS superuser, execute the chkconfig command to disable automatic start and stop.

In the example below, the shell script for automatic start and stop was created as "rc\_mc\_inst1" and is disabled for run levels 3 and 5.

Example)

```
# chkconfig --level 35 rc_mc_inst1 off
```

## If using systemd

As the OS superuser, execute the systemctl command to disable automatic start and stop.

The example below disables automatic start and stop of "mcoi\_inst1.service".

Example)

```
# systemctl disable mcoi_inst1.service
```

- 3. Perform the maintenance on the new standby server that was stopped.
- 4. Create a copy of the new primary server instance on the new standby server.

Execute the pg\_basebackup command to create data in the new standby server by synchronizing with the new primary server.

Example)

```
$ pg_basebackup -D /database/inst1 -X fetch --waldir=/transaction/inst1 --progress --verbose -R
--dbname='application_name=standbyServerName' -h primaryServerHostName -p
primaryServerPortNumber
```



## See

The procedure for copying the primary server instance to the standby server is the same as the procedure for setting up the standby server.

Refer to "2.5.2 Creating, Setting, and Registering the Standby Server Instance", and then perform the recovery.

5. Check the settings for automatic start and stop of the multiplexed instances and Mirroring Controller.

If the multiplexed instances and Mirroring Controller were configured in step 2 to not start and stop automatically when the operating system of the database server starts and stops, then change the settings back. This step can be skipped if automatic start and stop are not required.

If using an rc script

As the OS superuser, execute the chkconfig command to enable the shell script.

In the example below, the shell script for automatic start and stop was created as "rc\_mc\_inst1" and is enabled for run levels 3 and 5.

Example)

# chkconfig --level 35 rc\_mc\_inst1 on

### If using systemd

As the OS superuser, execute the systemctl command to enable automatic start and stop.

The example below disables automatic start and stop of "mcoi\_inst1.service".

Example)

# systemctl enable mcoi\_inst1.service

6. After the maintenance is complete, edit the following parameters in the postgresql.conf file of the standby server as required.

Copying an instance results in the value of the synchronous\_standby\_names parameter becoming the specified value on the primary server. Therefore, correct it to the specified value on the standby server. If the parameter was commented out, then you must uncomment it.

7. On the standby server, start (rebuild) Mirroring Controller.

Enabling automatic switch/disconnection

As the instance administrator user, execute the mc\_ctl command in start mode.

Example)

\$ mc\_ctl start -M /mcdir/inst1

# Disabling automatic switch/disconnection

As the instance administrator user, execute the mc\_ctl command in start mode with the -F option specified.

Example)

\$ mc\_ctl start -M /mcdir/inst1 -F



After Mirroring Controller is started, automatic switch/disconnection can be enabled or disabled using the enable-failover or disable-failover mode of the mc\_ctl command.

## Failback of the Primary Server

Revert the primary server and standby server to the original server configuration. Do this to execute the main job on the previous primary server. Refer to "4.1.1.3 Failback of the Primary Server" for details.



Obtain a backup as soon as this task is complete.

# 3.8.2 Stopping for Maintenance

Perform this procedure to stop all servers for periodic inspections, for example. On the server on which Mirroring Controller is running, execute the mc\_ctl command in stop mode to stop the instance and Mirroring Controller.

If automatic start and stop of Mirroring Controller has been configured using systemd, do not use the mc\_ctl command, but instead use the systemctl command. Refer to "2.12 Setting Automatic Start and Stop of Mirroring Controller and Multiplexed Instances" for details.

After that, on the server where the Mirroring Controller arbitration process is running, execute the mc\_arb command in stop mode to stop the Mirroring Controller arbitration process.

# **Stopping Mirroring Controller**

Example)

```
$ mc_ctl stop -M /mcdir/inst1 -a
```

## **Stopping the Mirroring Controller arbitration process**

Linux

Example)

```
$ mc_arb stop -M /mcarb_dir/arbiter1
```

#### Windows

Example)

```
> mc_arb stop -M D:\mcarb_dir\arbiter1
```

# 3.8.3 Arbitration Server Maintenance

Arbitration server maintenance can be performed without affecting database server operation.

Follow the procedure below to perform arbitration server maintenance.

1. Execute the mc\_arb command in stop mode to forcibly stop the Mirroring Controller arbitration process.

Linux

Example)

```
$ mc_arb stop -M /mcarb_dir/arbiter1 -e
```

## Windows

Example)

```
> mc_arb stop -M D:\mcarb_dir\arbiter1 -e
```

- 2. Perform maintenance tasks.
- 3. Execute the mc\_arb command in start mode to restart the Mirroring Controller arbitration process.

Linux

Example)

```
$ mc_arb start -M /mcarb_dir/arbiter1
```

## Windows

Example)

```
> mc_arb start -M D:\mcarb_dir\arbiter1
```

4. Execute the mc\_arb command in status mode to check that the arbitration server is connected to the database server.

The example below executes the mc\_arb command, and shows the status.

Linux

Example)

```
$ mc_arb status -M /mcarb_dir/arbiter1
```

server_id	host	status
server1	192.0.3.100	online
server2	192.0.3.110	online

#### Windows

Example)

> mc_arb st	atus -M D:\mcark	_dir\arbiter1
server_id	host	status
server1 server2	192.0.3.100 192.0.3.110	online online

5. Check the command output.

Items to be checked

Check that the output status is "online" on both lines.

# 3.9 Changes in Operation

The following changes in operation may be required:

- Changes required when the standby server is stopped
- Changing from single server mode to database multiplexing mode
- Changing from database multiplexing mode to single server mode
- Changing to database multiplexing mode when the arbitration server is used for automatic degradation
- Changing parameters
- Uninstalling in the database multiplexing mode

# 3.9.1 Changes Required when the Standby Server is Stopped

## Operation when the standby server is stopped

Before performing maintenance for the primary server instance when the standby server has been stopped, stop Mirroring Controller on the primary server, comment out the synchronous\_standby\_names parameter in the postgresql.conf file of the primary server, and then execute the pg\_ctl command in reload mode.

If this operation is not performed, operations performed on the primary server for the instance will remain in a wait state.



Refer to "pg\_ctl" in Reference for information on the command.

## Standby server downtime

If you specified the synchronous\_standby\_names parameter of the postgresql.conf file and then the standby server instance is stopped, consider the points below.

- The wal\_sender\_timeout parameter in the postgresql.conf file

If the standby server is stopped after the timeout set in this parameter was exceeded, an error stating that the transaction log could not be received may be output to the primary server system log, and all transaction logs that should be transferred to the standby server are accumulated.

- The wal\_keep\_segments parameter in the postgresql.conf file

If a transaction log that exceeds the value set in this parameter was generated while the standby server was stopped, the transaction log may be deleted.

Additionally, setting this parameter requires consideration regarding stabilization of the database multiplexing mode. Refer to "2.11.1 Tuning to Stabilize the Database Multiplexing Mode" for details.



The standby server must be rebuilt if the pending transaction log to be transferred to the standby server is lost when the standby server is started after the maintenance task is complete.

Take the action advised in the recovery operation that starts from "4.1.1.1.3 Identify cause of error and perform recovery" through to "4.1.1.2 Rebuild the Standby Server".

# 3.9.2 Changing from Single Server Mode to Database Multiplexing Mode

The procedure for switching single server mode to database multiplexing mode for the purposes of high reliability and load distribution of the system is explained below.

This procedure is equivalent to the setup procedure explained in "Chapter 2 Setting Up Database Multiplexing Mode".



## If the data storage destination directory name is not comprised of ASCII characters

Stop the application job and then migrate to a directory with a name that uses only ASCII characters:

- 1. Stop the database instance on the primary server.
- 2. Change the name of the data storage destination directory to one that uses only ASCII characters.



When encrypting the storage data, refer to "Database Multiplexing Mode" in the Operation Guide, and then perform the setup for encryption on the primary and standby servers.

.....

1. Install on the arbitration server

Perform this step only if the arbitration server is used for automatic degradation.

Install the Server Assistant on the server where the Mirroring Controller arbitration process is started.

Refer to "Installation" in the Installation and Setup Guide for Server Assistant for information on how to install the Server Assistant.

2. Install on the standby server

Install FUJITSU Enterprise Postgres on the server to be started as the standby server.

Refer to "Installation" in the Installation and Setup Guide for Server for information on how to install FUJITSU Enterprise Postgres.

Use ASCII characters in the data storage destination directory.

3. Stop the application jobs

Stop the application jobs to be connected to the primary server.

4. Change the primary server settings

To allow connections from the server to be started as the standby server, configure the settings in step 2 and thereafter of "2.4.2 Creating, Setting, and Registering the Primary Server Instance" on the primary server.

5. Set up the arbitration server

Refer to "2.3 Setting Up the Arbitration Server" for details.

Perform this step only if the arbitration server is used for automatic degradation.

6. Set up database multiplexing mode on the primary server

Refer to "2.4.1 Setting Up Database Multiplexing Mode on the Primary Server" for details.

7. Set up database multiplexing mode on the standby server

Refer to "2.5.1 Setting Up Database Multiplexing Mode on the Standby Server" for details.

8. Create the standby server instance and start it

Refer to "2.5.2 Creating, Setting, and Registering the Standby Server Instance" for details.

After the above steps are completed, refer to the remaining explanations in "Chapter 2 Setting Up Database Multiplexing Mode" and ensure that the required settings and operations are completed.

# 3.9.3 Changing from Database Multiplexing Mode to Single Server Mode

The procedure for stopping database multiplexing mode and changing to single server mode is explained below.

Some tasks must be performed on the database server, and others must be performed on the arbitration server.

The tasks on the arbitration server are required only if the arbitration server is used for automatic degradation.

#### Tasks on the database server

1. Determine the server for which the instance is to be stopped, and switch this server

Determine the server that is to be excluded as the database multiplexing mode target, and for which the instance is to be stopped.

If the server for which the instance is to be stopped is the primary server, execute the mc\_ctl command in the switch mode to switch the standby server to the primary server.

The standby server after the switch is complete will be the server for which the instance is to be stopped.

If the server for which the instance is to be stopped is the standby server, there is no need to perform the switch operation.

Example)

2. Stop Mirroring Controller and the instance, and delete the file resources

On the server that was determined in step 1, execute the mc\_ctl command in stop mode to stop Mirroring Controller and the instance.

If automatic start and stop of Mirroring Controller has been configured using systemd, do not use the mc\_ctl command, but instead use the systemctl command. Refer to "2.12 Setting Automatic Start and Stop of Mirroring Controller and Multiplexed Instances" for details.

Example)

```
$ mc_ctl stop -M /mcdir/inst1
```

\$ mc ctl switch -M /mcdir/inst1

Then, delete the following file resources:

- Data storage destination directory
- Mirroring Controller management directory

Example)

```
$ rm -rf /database/inst1
$ rm -rf /mcdir/inst1
```



Refer to "Security-Related Notes" in the Operation Guide for details on deleting the data securely.

3. Stop the application jobs

Stop the application jobs to be connected to the primary server.

4. Stop Mirroring Controller and the instance on the primary server

Execute the mc\_ctl command in stop mode on the primary server.

If automatic start and stop of Mirroring Controller has been configured using systemd, do not use the mc\_ctl command, but instead use the systemctl command. Refer to "2.12 Setting Automatic Start and Stop of Mirroring Controller and Multiplexed Instances" for details.

Example)

\$ mc\_ctl stop -M /mcdir/inst1

5. Delete the database multiplexing mode settings that were configured for the primary server instance.

Reset the postgresql.conf file parameters to their values before the database multiplexing operation was set.

Delete the file resources from the Mirroring Controller management directory.

If the backup operation was performed, delete the following resources:

- Mirroring Controller management directory backup data obtained in database multiplexing mode
- Instance backup data obtained in database multiplexing mode

Additionally, if the primary\_conninfo parameter is set in the postgresql.auto.conf file, execute the ALTER SYSTEM RESET statement to delete the setting.

Example)

An example execution of the psql command is shown below.

postgres=# ALTER SYSTEM RESET primary\_conninfo;

After these actions are performed, ensure that the backup data is collected when starting the single operation.



#### See

- Refer to "Security-Related Notes" in the Operation Guide for details on deleting the data securely.
- Refer to "2.14 Backup Operation" for details on the backup operation.
- Refer to "Appendix A Parameters" for details on the postgresql.conf file parameters.



In the above procedure, if the postgresql.conf file of the single primary server can be changed by reloading the file, the operation mode can be changed without stopping the application job.

In that case, execute the mc\_ctl command in stop mode with the --mc-only option specified to stop only Mirroring Controller in relation to stopping the primary server.

### Tasks on the arbitration server

Linux

1. Execute the mc\_arb command in stop mode to stop the Mirroring Controller arbitration process.

Example)

```
$ mc_arb stop -M /mcarb_dir/arbiter1
```

2. Delete the Mirroring Controller arbitration process management directory.

Example)

```
$ rm -rf /mcarb_dir/arbiter1
```

#### Windows

1. Execute the mc\_arb command in stop mode to stop the Mirroring Controller arbitration process.

Example)

```
> mc_arb stop -M D:\mcarb_dir\arbiter1
```

2. Unregister the Mirroring Controller arbitration process from the Windows service.

Execute the mc\_arb command in unregister mode to unregister the Mirroring Controller arbitration process from the Windows service.

Example)

```
> mc_arb unregister -M D:\mcarb_dir\arbiter1
```

3. Delete registrations related to the event log

If error logs are output to the event log in "2.2.2.1 Preparing to Output Error Logs to the Event Log (Windows)", delete the registered event source name for each instance.

Example)

```
> regsvr32 /u /i:"Mirroring Controller arbtier1" "c:\Program Files\Fujitsu
\fsepv<x>assistant64\lib\mcarbevent.dll"
```

Note that "<x>" indicates the product version.

4. Delete the Mirroring Controller arbitration process management directory.

Example)

```
> rmdir /S /Q D:\mcarb_dir\arbiter1
```

# 3.9.4 Changing to Database Multiplexing Mode when the Arbitration Server is Used for Automatic Degradation

This section provides the procedure to change to database multiplexing mode using the Mirroring Controller only on the database server when the arbitration server is used for automatic degradation.

Some tasks must be performed on the database server, and others must be performed on the arbitration server.

## Tasks on the arbitration server

1. Set up the arbitration server.

Refer to "2.3 Setting Up the Arbitration Server" for information on how to set up the arbitration server.

#### Tasks on the database server

1. On the server where Mirroring Controller is running, execute the mc\_ctl command in stop mode to stop Mirroring Controller on the primary server and standby server.

Example)

```
$ mc_ctl stop -M /mcdir/instl -a --mc-only
```

2. Edit the network.conf file of the primary server and standby server to add the information of the arbitration server.

Refer to "A.3 Network Configuration File" for details.

The definition example of the network.conf file of the primary server is shown below:

```
Example)
```

The IDs of the primary server and standby server are set to "server1" and "server2", and their port numbers are set to "27540" and "27541". The ID of the server of the Mirroring Controller arbitration process is set to "arbiter", and its port number is set to "27541".

```
server1 192.0.2.100,192.0.3.100 27540,27541 server
server2 192.0.2.110,192.0.3.110 27540,27541 server
arbiter 192.0.3.120 27541 arbiter
```



- Ensure that the port numbers set for the primary server, standby server, and arbitration server do not conflict with other software. Also do not configure the same segment for the admin network and arbitration network.

- If the server type is "server", two IP addresses or host names, and two port numbers need to be specified in the following order:
  - IP address or host name of the database server used as the admin network
  - IP address or host name of the database server used as the arbitration network
  - Port number of the database server used as the admin network
  - Port number of the database server used as the arbitration network
- If the server type is "arbiter", specify the IP address or host name set for the my\_address parameter and the port number set for the port parameter in arbitration.conf.
- 3. Edit the *serverIdentifier*.conf file of the primary server and standby server to add parameters required for the operation where the arbitration server is used for automatic degradation.

Refer to "A.4.1 Server Configuration File for the Database Servers" for information on the parameters required when the arbitration server is used for automatic degradation.

4. On the primary server and standby server, execute the mc\_ctl command in start mode to start the Mirroring Controller process.

Example)

```
$ mc_ctl start -M /mcdir/inst1 --mc-only
```

#### Common tasks

1. Check the connection status from the database server or arbitration server.

Refer to "2.8 Checking the Connection Status" for details.

# 3.9.5 Changing Parameters

Stop Mirroring Controller before editing the Mirroring Controller server configuration file and network configuration file.

If the Mirroring Controller process crashes or becomes unresponsive, restart is performed automatically by the Mirroring Controller monitoring process, and the configuration file is reloaded. Therefore, if the configuration file was being edited, unintended behavior will occur.

# 3.9.6 Uninstalling in Database Multiplexing Mode

This section explains how to uninstall FUJITSU Enterprise Postgres on a server using database multiplexing mode.

Some tasks must be performed on the database server, and others must be performed on the arbitration server.

The tasks on the arbitration server are required only if the arbitration server is used for automatic degradation.

### Tasks on the database server

- Stop the multiplexed instances and Mirroring Controller
   Refer to "3.2 Starting and Stopping Mirroring Controller" for information on how to stop the instance.
- Uninstall FUJITSU Enterprise Postgres
   Refer to "Uninstallation" in the Installation and Setup Guide for Server for information on how to uninstall FUJITSU Enterprise Postgres.

### Tasks on the arbitration server

Refer to "Uninstallation" in the Installation and Setup Guide for Server Assistant, and uninstall the Server Assistant.

# Chapter 4 Action Required when an Error Occurs in Database Multiplexing Mode

This chapter describes the action required if an error occurs in database multiplexing mode.

In database multiplexing mode, when an error is detected, the switch or disconnection of the standby server is performed automatically, so that only the primary server starts degrading. In this case, the recovery tasks will be required for the standby server on which the switch or disconnection was performed.

Other possible cases are as follows:

- When automatic switch fails
- When automatic disconnection fails
- When all servers or instances were stopped

# 4.1 Action Required when Server Degradation Occurs

If the server has started degrading, the recovery tasks will vary depending on whether the cause was the switch (failover or switchover), or the disconnection.

Execute the mc\_ctl command in status mode, or refer to the system log, and check if the cause of the server degradation was the switch or the disconnection.

In the example below, the mc\_ctl command is executed in status mode.

If a switch has occurred, "switched" (the switch is complete and the server is in a degrading state) is displayed for "mirroring status".

Example)

```
$ mc_ctl status -M /mcdir/inst1
mirroring status
-----switched
:
```

If a disconnection has occurred, "not-switchable" (disconnection was performed so the server cannot be switched) is displayed for "mirroring status".

Example)

```
$ mc_ctl status -M /mcdir/instl
mirroring status
-----
not-switchable
:
```



If Mirroring Controller detects any errors on the server on which operations are continuing during recovery to database multiplexing mode from a degrading operation state, perform the procedure in "4.1.3 Addressing Errors During Degrading Operation", and then recover to database multiplexing mode.

# 4.1.1 Operations when the Server has Started Degrading after a Switch has Occurred

This section explains the operations when the server has started degrading after a switch has occurred.



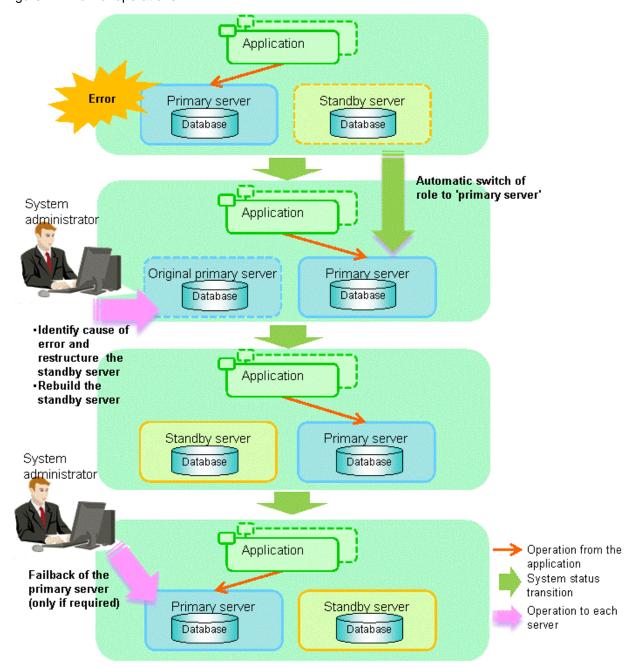
- After a switch has occurred as a result of an abnormality on the primary server, the database will not have a multiplexed configuration until the standby server is rebuilt. Remove the cause of the error as quickly as possible, and then rebuild the standby server.
- If the reference job was executed on the standby server, and the servers are switched because an error occurred on the primary server, the load is concentrated on the new primary server. Accordingly, pause the reference job on the original standby server, rebuild the original primary server as the new standby server, and then resume the reference job for the new standby server.
- If the instance on the new primary server is stopped before the original primary server where the error occurred is rebuilt as the new standby server, a split brain occurs at startup from the instance on the original primary server. Therefore, start the instance on the new primary server before rebuilding the standby server.

If the switch occurred and the server has started degrading, perform the following operations to recover the standby server and revert it to its original state:

- Identify Cause of Error and Restore the Standby Server
- Rebuild the Standby Server
- Failback of the Primary Server (only if required)

The flow of these operations is shown in the figure below.

Figure 4.1 Flow of operations



## 4.1.1.1 Identify Cause of Error and Restore the Standby Server

Perform the recovery according to the following procedure:

- 1. Stop Mirroring Controller
- 2. Recovery of the Mirroring Controller management directory
- 3. Identify cause of error and perform recovery

## 4.1.1.1.1 Stop Mirroring Controller

Execute the mc\_ctl command in stop mode for the original primary server on which the error occurred.

If automatic start and stop of Mirroring Controller has been configured using systemd, do not use the mc\_ctl command, but instead use the systemctl command. Refer to "2.12 Setting Automatic Start and Stop of Mirroring Controller and Multiplexed Instances" for details.

\$ mc\_ctl stop -M /mcdir/inst1

This also stops the instance that is required to perform the recovery.



If the instance does not stop, refer to "Actions in Response to Failure to Stop an Instance" in the Operation Guide, and then stop the instance.

Then, specify the -e option in the above command to forcibly stop Mirroring Controller.

## 4.1.1.1.2 Recovery of the Mirroring Controller management directory

Copy the files in the Mirroring Controller management directory from the backup data, and then perform the recovery.

## 4.1.1.3 Identify cause of error and perform recovery

Refer to the system log of the primary server and the standby server to identify the cause of the error, and then perform recovery.

The following commands can be used to recover a standby server. Select depending on the recovery and the situation.

- pg\_basebackup

Creates a copy of all resources of the primary server instance.

- pg\_rewind

Creates a copy of only the updated files on the new primary server. For this reason, if this command is used to incorporate a new standby server, recovery time can be shortened. To use this command to build the original primary server as a new standby server, at least one of the following must be met:

- a. Checksums were enabled when an instance was created and the original primary server was stopped normally, or
- b. The original primary server was stopped normally and the wal\_log\_hints parameter of postgresql.conf was enabled when an instance was started.

Additionally, full\_page\_writes must be enabled, which is its default value.

If it is required to start the original primary server again to then stop it normally, both servers will temporarily operate as primary servers. If data is updated on the original primary server, data corruption will occur, so take measures such as changing settings in pg\_hba.conf to disallow client connections.

When executing the pg\_rewind command, the timeline ID of the source server and the target server must be different. However, immediately after promotion of the new primary server, it may not be possible to execute the pg\_rewind command due to the application processing of unapplied update transaction logs and subsequent update processing of timeline IDs using checkpoint processing not being completed. Therefore, execute the pg\_rewind command after unapplied update transaction logs are no longer on the new primary server and updating of timeline IDs is complete.



- Refer to "pg\_basebackup" in "Reference" in the PostgreSQL Documentation for information on the pg\_basebackup command.

.....

 $- \ Refer to \ "pg\_rewind" \ in \ "Reference" \ in the \ PostgreSQL \ Documentation \ for \ information \ on \ the \ pg\_rewind \ command.$ 

The example below executes the pg\_rewind command to perform recovery by synchronizing data on the original primary server with the new primary server.

1. Use the pg\_ctl command to start the original primary server instance.

Example)

\$ pg\_ctl start -D /database/inst1

2. Use the pg\_ctl command to stop the original primary server instance.

Example)

```
$ pg_ctl stop -D /database/inst1
```

3. Wait for the application of unapplied update transaction logs on the new primary server.

Execute the SQL below on the new primary server, and wait until the result is false.

```
# select pg_is_in_recovery();
```

Example)

```
$ psql -h hostNameOfNewPrimaryServer -p portNumOfNewPrimaryServer -d dbName -c "select
pg_is_in_recovery();"
```

Any database can be connected to.



If the pg\_rewind command is executed immediately after promotion of the new primary server, the processing in steps 3 and 4 is required. If update-type SQL can be executed on the new primary server and checkpoint processing is executed after promotion, the processing in steps 3 and 4 will not be necessary.

4. Update the timeline ID.

Execute checkpoint processing, and update the timeline ID.

```
$ psql -h hostNameOfNewPrimaryServer -p portNumOfNewPrimaryServer -d dbName -c "checkpoint;"
```

Any database can be connected to.

5. Create a copy of the new primary server instance in the original primary server (new standby server).

Execute the pg\_rewind command to synchronize the new standby server data with the new primary server.

Example)

```
\label{local_policy} $$ pg_rewind -D /database/inst1 --source-server='user=userName host=newPrimaryServerHostName port=newPrimaryServerPortNumber dbname=dbName'
```

6. Specify parameters in the postgresql.conf file of the original primary server (new standby server).

Set the parameters required for the standby server in postgresql.conf.

Refer to "Table 2.5 Parameters" for information on the parameters to set in postgresql.conf.

7. Create the standby signal file on the original primary server (new standby server).

The file will be used to connect to the new primary server for streaming replication.

8. Use the pg\_ctl command to start the original primary(new standby server) server instance.

Example)

```
$ pg_ctl start -D /database/inst1
```

9. Use the ALTER SYSTEM SET statement to change the connection string specified for the primary\_conninfo parameter in postgresql.auto.conf file using as shown below.

Table 4.1 Changes to the connection string

Keyword	Content specified	Remarks
host	hostNameOfNewPrimaryServer	Specify informations for connecting to the
port	portNumOfNewPrimaryServer	primary server in the form of a PostgreSQL connection string.
application_name	newstandbyServerName	TostgresQL connection string.

Example)

An example execution of the psql command is shown below.

postgres=# ALTER SYSTEM SET primary\_conninfo='host=hostNameOfNewPrimaryServer
port=portNumOfNewPrimaryServer application\_name=newstandbyServerName';



- If you need to set a connection string other than host, port and application\_name, include it in the setting of the primary\_conninfo parameter.

.....

- The primary\_conninfo parameter should not be set in the postgresql.conf file, but only in the postgresql.auto.conf file using the ALTER SYSTEM SET statement.
- 10. Use the pg\_ctl command to stop the original primary server (new standby server) instance.

Example)

\$ pg\_ctl stop -D /database/inst1



### See

- Refer to "Hot Standby" in the PostgreSQL Documentation for details on the standby.signal file.
- Refer to "Setting Up a Standby Server" in the PostgreSQL Documentation for details on the primary\_conninfo.



A new timeline is branched for the new primary server due to promotion, so 'latest' needs to be specified for the recovery\_target\_timeline parameter so that the old primary server (new standby server) follows the new primary server.

# 4.1.1.2 Rebuild the Standby Server

The starting of the recovered original primary server as the standby server is referred to as the "standby server rebuild".

On the original primary server, start Mirroring Controller and the instance.

Enabling automatic switch/disconnection

As the instance administrator user, execute the mc\_ctl command in start mode.

Example)

\$ mc\_ctl start -M /mcdir/inst1

## Disabling automatic switch/disconnection

As the instance administrator user, execute the mc\_ctl command in start mode with the -F option specified.

Example)

\$ mc\_ctl start -M /mcdir/inst1 -F



After Mirroring Controller is started, automatic switch/disconnection can be enabled or disabled using the enable-failover or disable-failover mode of the mc\_ctl command.

# 4.1.1.3 Failback of the Primary Server

To revert the primary server and standby server to the original server configuration after rebuilding the standby server, perform failback for the primary server.

Do this to execute the main job on the previous primary server.

Perform the following procedure:

1. Failback of the primary server

Execute the mc\_ctl command in switch mode on the primary server or the standby server.

Example)

```
$ mc_ctl switch -M /mcdir/inst1
```

After executing the mc\_ctl command in switch mode, the status will be as follows:

Example)

2. Stop the original primary server

On the original primary server, execute the mc\_ctl command in stop mode to stop Mirroring Controller and the instance.

If automatic start and stop of Mirroring Controller has been configured using systemd, do not use the mc\_ctl command, but instead use the systemctl command. Refer to "2.12 Setting Automatic Start and Stop of Mirroring Controller and Multiplexed Instances" for details.

Example)

```
$ mc_ctl stop -M /mcdir/inst1
```

3. Create a copy of the new primary server instance in the original primary server (new standby server)

Execute the pg\_basebackup command to create data in the new standby server by synchronizing with the new primary server.

Example)

```
$ pg_basebackup -D /database/instl -X fetch --waldir=/transaction/instl --progress --verbose -R
--dbname='application_name=standbyServerName' -h primaryServerHostName -p
primaryServerPortNumber
```



The procedure for copying the new primary server instance to the new standby server is the same as the procedure for setting up the new standby server.

Refer to "2.5.2 Creating, Setting, and Registering the Standby Server Instance", and then perform the recovery.

4. Rebuild the standby server

On the standby server, start Mirroring Controller and the instance.

Enabling automatic switch/disconnection

As the instance administrator user, execute the mc\_ctl command in start mode.

Example)

\$ mc\_ctl start -M /mcdir/inst1

Disabling automatic switch/disconnection

As the instance administrator user, execute the mc\_ctl command in start mode with the -F option specified.

Example)

\$ mc\_ctl start -M /mcdir/inst1 -F



After Mirroring Controller is started, automatic switch/disconnection can be enabled or disabled using the enable-failover or disable-failover mode of the mc\_ctl command.

# 4.1.2 Operations when the Server has Started Degrading after a Disconnection has Occurred

This section explains the operations when the server has started degrading after a disconnection has occurred.



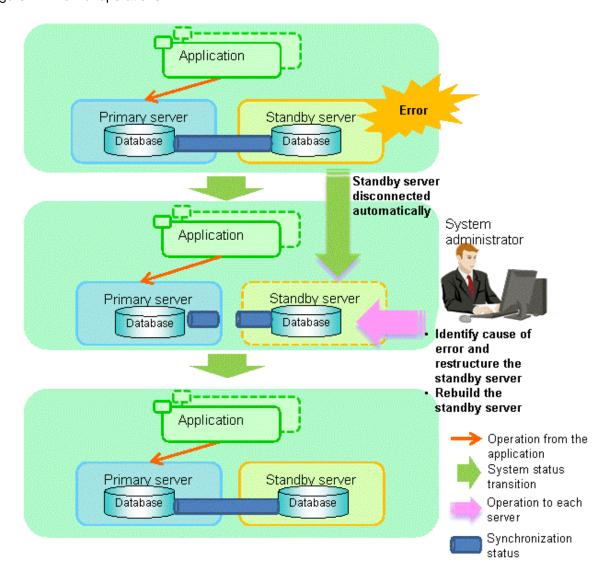
After a disconnection has occurred as a result of an abnormality on the standby server, the database will not have a multiplexed configuration until the standby server is rebuilt. Remove the cause of the error as quickly as possible, and then rebuild the standby server.

If the disconnection occurred and the server has started degrading, perform the following operations to recover the standby server and revert it to its original state:

- Identify Cause of Error and Restore the Standby Server
- Rebuild the Standby Server

The flow of these operations is shown in the figure below.

Figure 4.2 Flow of operations



# 4.1.2.1 Identify Cause of Error and Restore the Standby Server

Perform the recovery according to the following procedure:

- 1. Stop Mirroring Controller
- 2. Recovery of the Mirroring Controller management directory
- 3. Identify cause of error and perform recovery

# 4.1.2.1.1 Stop Mirroring Controller

Execute the mc\_ctl command in stop mode for the standby server on which the error occurred.

If automatic start and stop of Mirroring Controller has been configured using systemd, do not use the mc\_ctl command, but instead use the systemctl command. Refer to "2.12 Setting Automatic Start and Stop of Mirroring Controller and Multiplexed Instances" for details.

Example)

\$ mc\_ctl stop -M /mcdir/instl

This also stops the instance that is required to perform the recovery.



If the instance does not stop, refer to "Actions in Response to Failure to Stop an Instance" in the Operation Guide, and then stop the instance.

Then, specify the -e option in the above command to forcibly stop Mirroring Controller.

### 4.1.2.1.2 Recovery of the Mirroring Controller management directory

Copy the files in the Mirroring Controller management directory from the backup data, and then perform the recovery.

### 4.1.2.1.3 Identify cause of error and perform recovery

Refer to the system logs of the primary server and the standby server to identify the cause of the error, and then perform recovery.

Execute the pg\_basebackup command to perform recovery by synchronizing data in the primary server with the standby server.

Example)

\$ pg\_basebackup -D /database/inst1 -X fetch --waldir=/transaction/inst1 --progress --verbose -R --dbname='application\_name=standbyServerName' -h primaryServerHostName -p primaryServerPortNumber



See

This recovery procedure is the same as the procedure for setting up the standby server.

Refer to "2.5.2 Creating, Setting, and Registering the Standby Server Instance", and then perform the recovery.

### 4.1.2.2 Rebuild the Standby Server

Start the Mirroring Controller and the instance of the standby server, and rebuild the standby server.

Enabling automatic switch/disconnection

As the instance administrator user, execute the mc\_ctl command in start mode.

Example)

\$ mc\_ctl start -M /mcdir/inst1

#### Disabling automatic switch/disconnection

As the instance administrator user, execute the mc\_ctl command in start mode with the -F option specified.

Example)

\$ mc\_ctl start -M /mcdir/inst1 -F



**Point** 

After Mirroring Controller is started, automatic switch/disconnection can be enabled or disabled using the enable-failover or disable-failover mode of the mc\_ctl command.

## 4.1.3 Addressing Errors During Degrading Operation

This section explains how to address errors that may occur on the server on which operation is continuing during degrading operation triggered by a switch or disconnection.

#### If needing to recover from backup data

If it is necessary to recover the database using backup data due to data becoming corrupted from disk failure or user operation error, refer to the following for information on recovery to database multiplexing mode:

- Action Required when All Database Servers or Instances Stopped
- Recovering from an Incorrect User Operation

#### If a temporary error occurs

If a temporary error occurs, such as due to a high load on the server or insufficient system resources, remove the cause of the error and restart Mirroring Controller, and then refer to the following for details on recovery to database multiplexing mode:

- Operations when the Server has Started Degrading after a Switch has Occurred
- Operations when the Server has Started Degrading after a Disconnection has Occurred



Refer to "3.2 Starting and Stopping Mirroring Controller" for information on restarting Mirroring Controller.

# 4.2 Action Required when Automatic Switch Fails

If the system behavior is unstable, for example there are insufficient temporary system resources, the Mirroring Controller automatic switch may fail.

Perform the switch manually using one of the following methods:

- Refer to the procedures in "3.4 Manually Switching the Primary Server".
- In the standby server, execute the mc\_ctl command in switch mode with the -force option specified to forcibly perform the switch.

  Example)

\$ mc\_ctl switch -M /mcdir/instl --force



- Even if connection cannot be established between database servers, it is possible to fence the primary server and forcibly switch by executing the mc\_ctl command in switch mode with the --force option specified.

- The primary server is not fenced in the cases below, so stop Mirroring Controller and instances of the primary server database in advance:
  - The --no-fencing option is specified when performing forced switch.
  - The heartbeat\_error\_action parameter in *serverIdentifier*.conf is set to "message" and the fencing command is not configured to be used (the fencing\_command parameter is omitted in *serverIdentifier*.conf).
  - The heartbeat\_error\_action parameter in serverIdentifier.conf is set to "fallback".



See

#### Recovery to database multiplexing mode

Refer to "4.1.1.2 Rebuild the Standby Server" and "4.1.1.3 Failback of the Primary Server" for information on recovery to database multiplexing mode.

# 4.3 Action Required when Automatic Disconnection Fails

If the system behavior is unstable, for example there are insufficient system resources such as available memory or free disk space, automatic disconnection using Mirroring Controller may not be possible.

Perform the disconnection manually using one of the following methods:

- Refer to the procedures in "3.5 Manually Disconnecting the Standby Server".
- In the primary server, execute the mc\_ctl command in detach mode to perform forced disconnection.

Example)

\$ mc\_ctl detach -M /mcdir/inst1



- Even if connection cannot be established between database servers, it is possible to fence the standby server and forcibly disconnect by executing the mc\_ctl command in detach mode.
- In the cases below, stop Mirroring Controller and instances of the standby server database in advance so that the standby server is not fenced:
  - The --no-fencing option is specified when performing forced disconnection.
  - The heartbeat\_error\_action parameter in *serverIdentifier*.conf is set to "message" and the fencing command is not configured to be used (the fencing\_command parameter is omitted in *serverIdentifier.conf*).
  - The heartbeat\_error\_action parameter in *serverIdentifier*.conf is set to "fallback".



Recovery to database multiplexing mode

Refer to "4.1.2.2 Rebuild the Standby Server" for information on recovery to database multiplexing mode.

# 4.4 Action Required when All Database Servers or Instances Stopped

This section explains what happens when all database servers or instances on the database server have stopped, so jobs cannot continue.



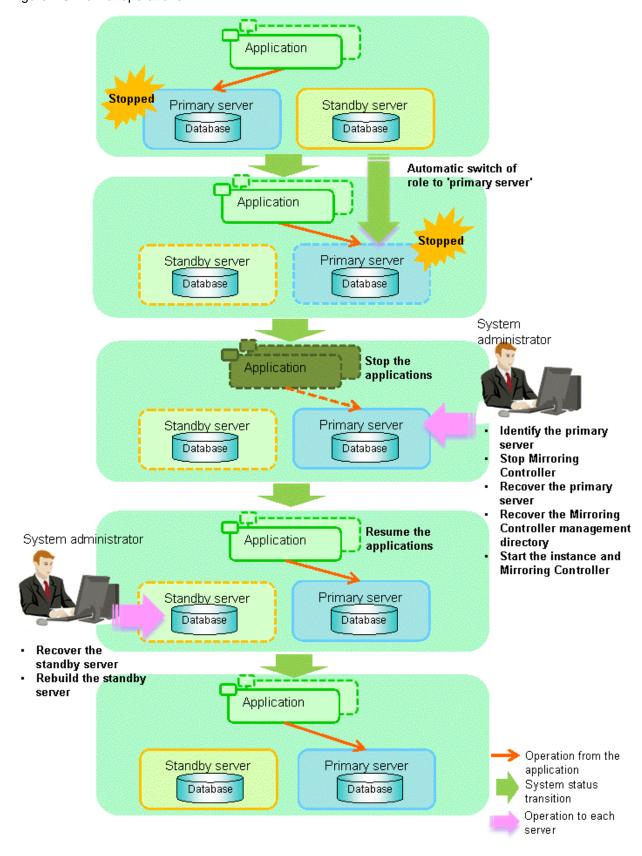
See

### Recovery to database multiplexing mode

Refer to "4.1.1.2 Rebuild the Standby Server" and "4.1.1.3 Failback of the Primary Server" for information on recovery to database multiplexing mode.

The flow of these recovery operations is shown in the figure below.

Figure 4.3 Flow of operations



Perform the following procedure:

Stop the applications
 Stop the applications from running.

#### 2. Identify the primary server

Use one of the following methods to identify the primary server that was running before the servers or instances stopped:

- Refer to the system log on each server and identify the server where the following message was output.

Message:

```
MirroringControllerOpen[30017]: LOG: promotion processing completed (MCA00062)
```

 On each server, execute the mc\_ctl command in status mode to search the servers for which "none(inactivated primary)" is displayed.

#### 3. Stop Mirroring Controller on the primary server

Execute the mc\_ctl command in stop mode on the primary server.

If automatic start and stop of Mirroring Controller has been configured using systemd, do not use the mc\_ctl command, but instead use the systemctl command. Refer to "2.12 Setting Automatic Start and Stop of Mirroring Controller and Multiplexed Instances" for details.

Example)

\$ mc\_ctl stop -M /mcdir/inst1



#### Forcibly stopping Mirroring Controller

If Mirroring Controller does not stop, specify the -e option in the stop mode of the mc\_ctl command and then execute the command.

Example)

#### 4. Recover the primary server

First, refer to "Actions when an Error Occurs" in the Operation Guide, and then identify the cause of the error and perform recovery.

Next, recover the primary server using the recovery method that uses the pgx\_rcvall command based on the backup data.

If the backup operation was performed using the pgx\_dmpall command based on the instructions in "2.14.2 Database Backup Operation", perform the following procedure for the recovery:

- a. Perform the following operations on both the primary server and the standby server, and check the server containing the backup data and the archive log that show the latest date.
  - Execute the pgx\_rcvall command with the -l option specified and identify the backup data that shows the latest date.
  - Identify the archive log that shows the latest date, as shown below.

Example)

```
$ ls -ltr backupDataStorageDir/*_wal
```

- b. If the latest backup data exists on the standby server, copy (\*1) the backup data and overwrite (\*2) it to each backup storage destination directory on the primary server.
- c. If the latest archive log and transaction log file exist on the standby server, copy (\*1) the archive log and overwrite (\*2) it to the backup storage destination directory on the primary server.
- d. Execute the pgx\_rcvall command on the primary server, specifying the backup storage destination directory of the primary server.



\*1: The backup data may contain a symbolic link, so copy the backup data so that the symbolic link is not converted to an ordinary file (with the tar command, for example).

\*2: If you can save a copy of the backup storage destination directory, do so without overwriting it.



#### See

Refer to "Actions when an Error Occurs" in the Operation Guide for information on the pgx\_rcvall command.

5. Recover the Mirroring Controller management directory

Copy the files in the Mirroring Controller management directory from the backup data, and then perform the recovery.

6. Start the primary server instance and Mirroring Controller

Enabling automatic switch/disconnection

As the instance administrator user, execute the mc\_ctl command in start mode.

Example)

\$ mc\_ctl start -M /mcdir/inst1

#### Disabling automatic switch/disconnection

As the instance administrator user, execute the mc\_ctl command in start mode with the -F option specified.

Example)

\$ mc\_ctl start -M /mcdir/inst1 -F



### **Point**

After Mirroring Controller is started, automatic switch/disconnection can be enabled or disabled using the enable-failover or disable-failover mode of the mc\_ctl command.

7. Resume applications

Resume the applications.

8. Stop Mirroring Controller on the standby server

Execute the mc\_ctl command in stop mode on the standby server.

If automatic start and stop of Mirroring Controller has been configured using systemd, do not use the mc\_ctl command, but instead use the systemctl command. Refer to "2.12 Setting Automatic Start and Stop of Mirroring Controller and Multiplexed Instances" for details.

Example)

\$ mc\_ctl stop -M /mcdir/inst1

9. Recover the standby server

Refer to "2.5.2 Creating, Setting, and Registering the Standby Server Instance", and then recover (set up) the standby server from the primary server.

10. Rebuild the standby server

On the standby server, start Mirroring Controller and the instance.

### Enabling automatic switch/disconnection

As the instance administrator user, execute the mc\_ctl command in start mode.

Example)

\$ mc\_ctl start -M /mcdir/inst1

#### Disabling automatic switch/disconnection

As the instance administrator user, execute the mc\_ctl command in start mode with the -F option specified.

Example)

\$ mc\_ctl start -M /mcdir/inst1 -F



After Mirroring Controller is started, automatic switch/disconnection can be enabled or disabled using the enable-failover or disable-failover mode of the mc\_ctl command.

# 4.5 Recovering from an Incorrect User Operation

This section describes how to recover an instance when data has been corrupted due to incorrect user operation.

For example, when data has been corrupted due to incorrect user operation, such as data being unintentionally changed or deleted by an application or command, it is necessary to restore the original data on the primary server and resynchronize with the standby server.

Use the following procedure to perform recovery.

1. Identify the primary server

Execute the mc\_ctl command in status mode on each server, and search for a server for which "primary" or "none(inactivated primary)" is displayed.

2. Stop the applications and commands that caused the incorrect operation to occur

Stop applications and commands that are running on the primary server. This will minimize the impact caused by the incorrect data.

Also, if any applications used for reference by the standby server are running, stop them too.

3. Stop the instance and Mirroring Controller

Stop the instance and Mirroring Controller on both the primary server and standby server.

Example)

\$ mc\_ctl stop -a -M /mcdir/inst1

4. Recover the database on the primary server

Recover the database using the recovery method in which the pgx\_rcvall command uses the backup data to recover the database to a restore point prior to the time when the incorrect operation was performed.



See

Refer to "Recovering from an Incorrect User Operation" in the Operation Guide for information on using the pgx\_rcvall command to recover the database to a restore point, and then perform only the database recovery procedure while the instance is in a stop state.

5. Start the instance and Mirroring Controller

Start the instance and Mirroring Controller on the primary server.

Enabling automatic switch/disconnection

As the instance administrator user, execute the mc\_ctl command in start mode.

Example)

\$ mc\_ctl start -M /mcdir/inst1

### Disabling automatic switch/disconnection

As the instance administrator user, execute the mc\_ctl command in start mode with the -F option specified.

Example)

\$ mc\_ctl start -M /mcdir/inst1 -F



### Point

After Mirroring Controller is started, automatic switch/disconnection can be enabled or disabled using the enable-failover or disablefailover mode of the mc\_ctl command.

.....

6. Build the new standby server

Refer to "2.5 Setting Up the Standby Server" for information on building (setting up) a standby server from the primary server.

# Chapter 5 Managing Mirroring Controller Using WebAdmin

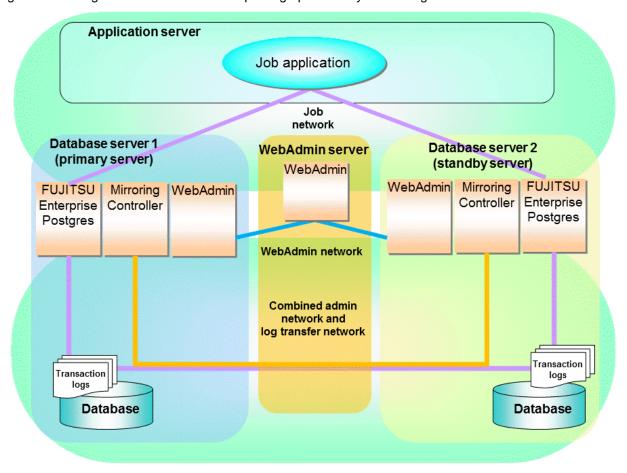
This chapter describes how to set up and manage Mirroring Controller in a streaming replication cluster using WebAdmin.

Mirroring Controller can be used to monitor a streaming replication cluster and perform automatic switching or disconnect synchronous replication when there is an error.

WebAdmin can be used to set up Mirroring Controller in an existing replication cluster. Mirroring Controller can be set up for either synchronous standby instances or asynchronous standby instances.

The configuration of the database multiplexing system built using WebAdmin is shown below:

Figure 5.1 Configuration of database multiplexing operation system using WebAdmin





- If Mirroring Controller is set up to the replication cluster using WebAdmin, the network with the host name (or IP address) specified in [Host name] will be used as the admin network and the log transfer network.
- To use a network other than the job network as the log transfer network, before building the replication cluster specify a host name other than the job network one in [Host name].



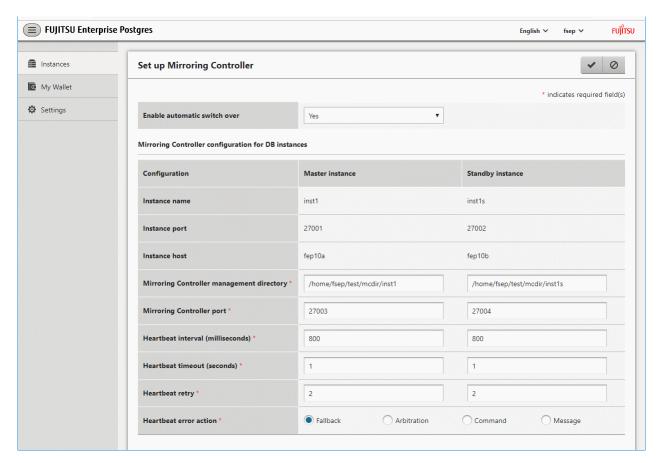
If you set up the arbitration server using WebAdmin, install WebAdmin on the arbitration server.

### **5.1 Mirroring Controller Setup**

Perform the following procedure to set up Mirroring Controller in a streaming replication cluster.

- 1. In the [Instances] tab, select the standby instance on which Mirroring Controller needs to be set up.
- 2. Click
- 3. Enter the information for the Mirroring Controller to be set up.

In the example below, Mirroring Controller is being set up for the replication cluster having master instance "inst1" and standby instance "inst1s".



The instance name, host address and port of the master and standby instances are displayed for easy reference.

Enter the following items on master instance and on standby instance fields for Mirroring Controller setup, as shown in the above screenshot:

- [Enable automatic switch over]: Toggles the automatic switch/disconnection functionality. Select "Yes". The default is "No".
- [Mirroring Controller management directory]: Directory where the Mirroring Controller configuration files will be stored. When the [Mirroring Controller management directory] is entered, WebAdmin will search the Mirroring Controller configuration files in the entered directory based on the [Data storage path] of the corresponding DB instance. If Mirroring Controller configuration files are found, the Mirroring Controller fields will be auto filled.
- [Mirroring Controller port]: Port number of Mirroring Controller. Note that if the Windows firewall feature is enabled, you must enable the port number of Mirroring Controller. Refer to "E.2 Windows Firewall Settings" for details.
- [Heartbeat interval (milliseconds)]: Number of milliseconds between two consecutive heartbeat checks. The default is "800".
- [Heartbeat timeout (seconds)]: Number of seconds for the heartbeat timeout. The default is "1".
- [Heartbeat retry]: Number of retries for heartbeat monitoring, before failover occurs. The default is "2".

- [Heartbeat error action]: Operation when a heartbeat abnormality is detected. The default is "Fallback".

When using FUJITSU Enterprise Postgres 10 and 11 instances created with previous versions, the instances will be in compatibility mode, and the "Fallback" is preselected and cannot be changed in the [Heartbeat error action] for Mirroring Controller setup.

When setting up Mirroring Controller for FUJITSU Enterprise Postgres 9.5 and 9.6 instances, the [Heartbeat error action] is not supported and therefore is not displayed.

When the [Heartbeat error action] is set to "Arbitration", the following extra items are displayed:

- [Arbitration network IP address]: IP address of the arbitration network.
- [Mirroring Controller Arbitration port]: Port number of Mirroring Controller for communicating with the arbitration server.

The [Arbitration server configuration] section is also displayed with the following items. The [Arbitration server configuration] will not be auto filled.

- [Location]: Location of the arbitration server. "Local" or "Remote" can be selected depending on your configuration.

If the arbitration server and WebAdmin server are located on the same server, you can select "Local" and the following items are displayed:

- [Arbitration management directory]: Directory where the arbitration server configuration files will be stored.
- [Arbitration server host or IP address]: Host name or IP address of the arbitration server.
- [Arbitration process port]: Port number for the arbitration process.
- [Fencing command]: Full path of the fencing command that fences a database server when an abnormality is detected.

If "Remote" is set for the item, the items below are displayed in addition to the above items.

- In the [Arbitration server configuration] section, [Operating system credential] is displayed where you can enter the following information:

[User name]: User name to access the arbitration server.

[Password]: Password to access the arbitration server.

- In the [Remote WebAdmin for Arbitration server] section, the following items are displayed:

[Remote WebAdmin address]: IP address of the remote WebAdmin installed on the arbitration server.

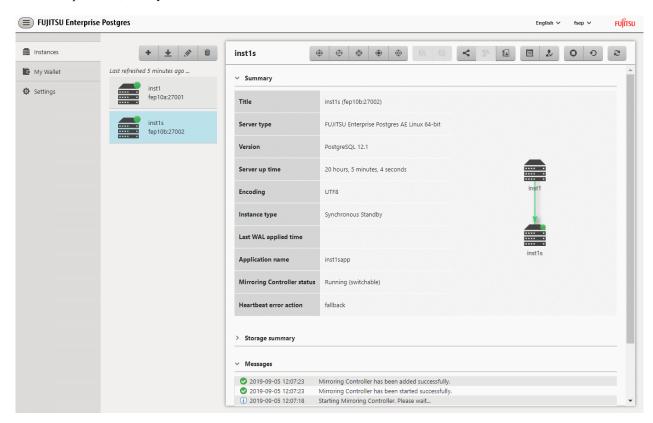
[Remote WebAdmin port]: Port number for the WebAdmin installed on the arbitration server.

When the [Heartbeat error action] is set to "Command", the following extra items are displayed:

- [Arbitration command]: Full path of the arbitration command to be executed when an abnormality is detected.
- [Fencing command]: Full path of the fencing command that fences a database server when an abnormality is detected.
- 4. Click to set up Mirroring Controller.

5. Upon successful completion, Mirroring Controller will be started on master and standby instances.

In the [Instances] tab, select standby instance. The page below is displayed, in which you can check the Mirroring Controller status. In the example below, standby instance "inst1s" is used.



After the Mirroring Controller has been set up, ([Edit Mirroring Controller] button) and ([Mirroring Controller Configuration] button) are available. These buttons are displayed only when FUJITSU Enterprise Postgres 10, 11 and 12 instances are created with FUJITSU Enterprise Postgres WebAdmin 12.

For FUJITSU Enterprise Postgres 9.5 and 9.6 instances, the [Heartbeat error action] will not be displayed.

When the [Heartbeat error action] is "Arbitration", the following information is displayed: whether the arbitration status is "online" or "offline", the arbitration server IP address and the arbitration process port.



Operating system credential (User name, Password) should not contain hazardous characters. Refer to "Appendix F WebAdmin Disallow User Inputs Containing Hazardous Characters".

# 5.2 Edit Mirroring Controller Setup

Settings made in "5.1 Mirroring Controller Setup" can be updated in either the master instance or a standby instance using WebAdmin.

Perform the following procedure to edit Mirroring Controller configuration:

- 1. In the [Instances] tab, select the instance for which the Mirroring Controller configuration is to be edited.
- 2. Click .
- 3. Enter the information for the Mirroring Controller to be updated. Refer to "5.1 Mirroring Controller Setup".
- 4. Click to update the Mirroring Controller.

5. Upon successful completion, Mirroring Controller will be started on master and standby instances.

Editing and saving the [Edit Mirroring Controller] page will reset all other settings that are not listed on this page to default values.

# 5.3 Mirroring Controller Configuration

The information related to Mirroring Controller monitoring and control (refer to "A.4.1 Server Configuration File for the Database Servers") and the information related to arbitration and control of the Mirroring Controller arbitration process (refer to "A.4.2 Arbitration Configuration File") can be set using WebAdmin. You can view and update the configuration on either the master instance or the standby instance.

Perform the following procedure:

- 1. In the [Instances] tab, select the instance for the Mirroring Controller configuration you want to view.
- 2. Click to view the Mirroring Controller configuration.
- 3. Click to show the editing page for the Mirroring Controller configuration. The Mirroring Controller configurations defined during [Mirroring Controller Setup] are read-only on this page. Refer to "5.1 Mirroring Controller Setup".

Additionally, refer to the "Appendix A Parameters" for information about the settings and the corresponding parameter names.

The items common to all [Heartbeat error action] are:

- Target DB
- Core file path
- Syslog facility
- Syslog identity
- Remote call timeout (milliseconds)
- Agent alive timeout (seconds)
- DB instance check interval (milliseconds)
- DB instance check timeout (seconds)
- DB instance check retry
- DB instance timeout action
- Disk check interval (milliseconds)
- Disk check retry
- Tablespace directory error action
- Post-switch command
- Post-promote command

(Post-promote command is replaced in FUJITSU Enterprise Postgres 12. The Post-promote command is still valid and will be displayed when it is used in the server configuration file of Mirroring Controller.)

- Post-attach command
- Pre-detach command
- State transition command timeout (seconds)
- Check synchronous standby names validation

When the [Heartbeat error action] is set to "Arbitration", the following extra items are displayed:

- Arbitration timeout (seconds)
- Arbiter alive interval (milliseconds)

- Arbiter alive retry
- Arbiter alive timeout (seconds)
- Arbiter connect interval (milliseconds)
- Arbiter connect timeout (seconds)
- Fencing command
- Fencing command timeout (seconds)
- Shutdown detached synchronous standby

When the [Heartbeat error action] is set to "Arbitration", the [Arbitration server configuration] section is displayed with the following items:

- Core file path
- Syslog facility
- Syslog identity
- Fencing command timeout (seconds)
- Heartbeat interval (milliseconds)
- Heartbeat timeout (seconds)
- Heartbeat retry

When the [Heartbeat error action] is set to "Command", the following extra items are available:

- Fencing command timeout (seconds)
- Arbitration command timeout (seconds)
- Shutdown detached synchronous standby

When the [Heartbeat error action] is set to "Message", the following extra items are available:

- Fencing command
- Fencing command timeout (seconds)

In addition, the following configurations are provided:

- DB instance JDBC connection SSL parameters
- DB instance libpq connection SSL parameters
- 4. Click v to update the Mirroring Controller configurations.

## 5.4 Stopping Mirroring Controller

Mirroring Controller can be stopped either in master instance or in standby instance using WebAdmin.

Perform the following procedure to stop Mirroring Controller.

- 1. In the [Instances] tab, select the instance where to stop Mirroring Controller.
- 2. Click
- 3. In the confirmation dialog box, click [Yes].

Mirroring Controller will be stopped on the selected instance. The Mirroring Controller status will be updated, and a confirmation message entry will be displayed in the [Message] section.

# 5.5 Starting Mirroring Controller

Mirroring Controller can be started either in master instance or in standby instance using WebAdmin.

Perform the following procedure to start Mirroring Controller.

- 1. In the [Instances] tab, select the instance where to start Mirroring Controller.
- 2. Click
- 3. In the confirmation dialog box, select the desired failover mode, and then click [Yes].

Mirroring Controller will be started on the selected instance. The Mirroring Controller status will be updated, and a confirmation message entry will be displayed in the [Message] section.

## 5.6 Disabling Failover Mode

Disabling failover mode in Mirroring Controller disables automatic switch/disconnection between master and standby instances.

Perform the following procedure to disable failover mode.

- 1. In the [Instances] tab, select the instance.
- 2. Click
- 3. In the confirmation dialog box, click [Yes].

Failover mode will be disabled in Mirroring Controller. The Mirroring Controller status will be updated and a confirmation message entry will be displayed in the [Message] section.

# 5.7 Enabling Failover Mode

Enabling failover mode in Mirroring Controller enables automatic switch/disconnection between master and standby instances.

Perform the following procedure to enable failover.

- 1. In the [Instances] tab, select the instance.
- 2. Click
- 3. In the confirmation dialog box, click [Yes].

Failover mode will be enabled in Mirroring Controller. The Mirroring Controller status will be updated and a confirmation message entry will be displayed in the [Message] section.

## 5.8 Deleting Mirroring Controller Setup

Deleting Mirroring Controller setup removes its setup from master and standby instances.

- 1. In the [Instances] tab, select the instance.
- 2. Click 🛞
- 3. In the confirmation dialog box, click [Yes].

Mirroring Controller setup will be removed from the cluster. The cluster status will be updated and a confirmation message entry will be displayed in the [Message] section.

For the instances in FUJITSU Enterprise Postgres 12, WebAdmin does not delete the Mirroring Controller management directory and the configuration files.

## 5.9 Status Update after Failover

When Mirroring Controller performs a failover, standby instance will be promoted to standalone instance. The Mirroring Controller setup will be removed from both standby and master instances.

The following scenario describes one of the ways in which failover can be triggered, and the results achieved by the use of Mirroring Controller in WebAdmin.

- 1. In the [Instances] tab, select the master instance "inst1".
- 2. Click .
- 3. In the confirmation dialog box, the warning "This instance is being monitored by Mirroring Controller. Stopping the instance may result in cluster failover." is displayed.
- 4. Choose the stop mode and click [Yes].

In the server, the following takes place:

- a. The master instance is stopped.
- b. Failover is triggered in Mirroring Controller.
- c. The Mirroring Controller setup is removed from both master and standby instances
- d. Standby instance is promoted to standalone.
- 5. When the instance is refreshed in WebAdmin, the latest status of the instances will be displayed.



When failover is performed, the Mirroring Controller setup is removed from both master and standby instances. Therefore, to manage the Mirroring Controller using WebAdmin again, create the standby instance and set up Mirroring Controller.

Refer to "Creating a Standby Instance" in the Operation Guide for details.

Refer to "5.1 Mirroring Controller Setup" for details.

# 5.10 Action Required when an Error Occurs in the Combined Admin Network and Log Transfer Network

Communication errors may temporarily occur in the network used as the admin network and log transfer network due to reasons such as high load on the server or insufficient system resources. Because of this, there is a risk of causing a split-brain situation by mistake even though the server has no issues.

Split brain is a phenomenon in which both servers temporarily operate as primary servers, causing data updates to be performed on both servers.

#### How to detect split brain using WebAdmin

If the conditions below are met, split brain may occur. Refer to "Split-brain detection method" and "How to recover from a split-brain" in "Appendix D Notes on Performing Automatic Degradation Immediately after a Heartbeat Abnormality" and take the actions described.

- 1. A standby instance is selected in the [Instances] tab, and
- 2. "Standalone" is displayed in [Instance type], and
- 3. A master instance is selected in the [Instances] tab, and
- 4. "Standalone" is displayed in [Instance type].



The admin network is important because Mirroring Controllers use it to confirm the status of each server.

The log transfer network is also important to maintain the data freshness.

Therefore, use network configurations resistant to faults for these networks by using the network redundancy channel bonding feature provided by the operating system or network driver vendor.

# 5.11 Performing Automatic Degradation Using the Arbitration Server

If database multiplexing is performed using WebAdmin, it is also possible to perform automatic degradation using the arbitration server. In such cases, it is necessary to perform tasks on the database server and the arbitration server after setting up Mirroring Controller in WebAdmin.

#### Tasks on the arbitration server

Perform setup of the arbitration server using Mirroring Controller commands.

1. Set up the arbitration server.

Refer to "2.3 Setting Up the Arbitration Server" in "Chapter 2 Setting Up Database Multiplexing Mode" for information on how to set up the arbitration server.

#### Tasks on the database server

Change some of the settings after setting up Mirroring Controller in WebAdmin.

1. Set up Mirroring Controller in WebAdmin.

Refer to "5.1 Mirroring Controller Setup" for details.

2. Use WebAdmin to stop Mirroring Controller on the master and standby instances.

Refer to "5.4 Stopping Mirroring Controller" for details.

3. Edit the network configuration file of the master and standby instances, and add the arbitration server information.

The network configuration file is network.conf, which exists in the Mirroring Controller management directory specified during Mirroring Controller setup. Refer to "A.3 Network Configuration File" for details.

A definition example of network.conf is shown below.

#### Example:

The port number of the database server to be used as the arbitration network is set to "27541". The ID of the server of the Mirroring Controller arbitration process is set to "arbiter", and its port number is set to "27541".

```
dbsvm27500 192.0.2.100,192.0.3.100 27540,27541 server dbsvs27500 192.0.2.110,192.0.3.110 27540,27541 server arbiter 192.0.3.120 27541 arbiter
```



- Ensure that the port numbers set for the database server and the arbitration server do not conflict with other software. In addition, do not configure the same segment for the admin network and the arbitration network.
- If the server type is "server", two IP addresses or host names, and two port numbers need to be specified in the following order:
  - IP address or host name of the database server used as the admin network
  - IP address or host name of the database server used as the arbitration network
  - Port number of the database server used as the admin network
  - Port number of the database server used as the arbitration network

- If the server type is "arbiter", specify the IP address or host name set for the my\_address parameter and the port number set for the port parameter in arbitration.conf of the arbitration server.
- 4. Edit the server configuration file of the master and standby instances, and add the parameters required for automatic degradation using the arbitration server.

The server configuration file is *instanceName*.conf or *instancePort*.conf, which exists in the Mirroring Controller management directory specified during Mirroring Controller setup.

To perform automatic degradation using the arbitration server, set the heartbeat\_error\_action parameter to "arbitration".

Refer to "A.4.1 Server Configuration File for the Database Servers" for information on other parameters.

5. Use WebAdmin to start Mirroring Controller on the master and standby instances.

Refer to "5.5 Starting Mirroring Controller" for details.

#### Common tasks

1. Use the Mirroring Controller command to check the connection status from the database server or the arbitration server.

Refer to "2.8 Checking the Connection Status" for information on how to check the connection status.

# Appendix A Parameters

This appendix describes the configuration files and parameters required by the database multiplexing mode.



Refer to "Server Configuration" in the PostgreSQL Documentation for information on the postgresql.conf file.

# A.1 Parameters Set on the Primary Server

The content for the parameters set in the postgresql.conf file of the primary server is shown in the table below.

Table A.1 postgresgl.conf file

Parameter	Value set	Explanation
wal_level	replica or logical	Specify the output level for the transaction log.
		Specify "logical" when logical decoding is also to be used.
max_wal_senders	2 or more	Specify "2" when building a Mirroring Controller cluster system.
		When additionally connecting asynchronous standby servers to the cluster system, add the number of simultaneous connections from these standby servers.
synchronous_standby_names	'standbyServerName'	Use single quotation marks (') to enclose the name that will identify the standby server. Any name can be specified.
		Do not change this parameter while Mirroring Controller is running.
		Do not specify multiple names to this parameter as the Mirroring Controller can manage only one standby server.
hot_standby	on	Specify whether queries can be run on the standby server.
		Specify "on".
wal_keep_segments	Number of file segments	If a delay exceeding the value set in this parameter occurs, the WAL segment required later by the primary server may be deleted.
		Additionally, if you stop a standby server (for maintenance, for example), consider the stop time and set a value that will not cause the WAL segment to be deleted.
		Refer to "Estimating Transaction Log Space Requirements" in the Installation and Setup Guide for Server for information on estimating the WAL segment.
wal_log_hints	on	When using the pg_rewind command to recover a standby server, specify this parameter or enable checksums when executing the initdb command.
wal_sender_timeout	Timeout (milliseconds)	Specify the time period after which it is determined that the receiver process (walreceiver) of the transaction log is in an abnormal state on the primary server.
		The specified value must be larger than the value set for the wal_receiver_status_interval parameter set in the postgresql.conf file of the standby server.

Parameter	Value set	Explanation
		By aligning this value with the value for the database process heartbeat monitoring time, you can unify the time after which it is determined that an error has occurred.
wal_receiver_timeout	Timeout (milliseconds)	Specify the time period after which it is determined that an error has occurred when the transaction log was received on the standby server.
		By aligning this value with the value for the database process heartbeat monitoring time, you can unify the time after which it is determined that an error has occurred.
archive_mode	on	Specify the archive log mode.
archive_command	'installDir/bin/pgx_walcopy.cmd "%p" "backupDataStorageDestinationDirectory/archived_wal/%f"	Specify the command and storage destination to save the transaction log.
backup_destination	Backup data storage destination directory	Specify the name of directory where to store the backup data.
		Set the permissions so that only the instance administrator user can access the specified directory.
		Specify the same full path on all servers, so that the backup data of other servers can be used to perform recovery.
listen_addresses	Primary server IP address, host name, or "*"	Specify the IP address or host name of the primary server. Specify the IP address or corresponding host name that will be used to connect to the log transfer network.
		The content specified is also used to allow connections from client applications.
		To receive the connection and the transaction log from any client or standby server, specify "*".
		Refer to "Connections and Authentication" in the PostgreSQL Documentation for details.
max_connections	Number of simultaneous client connections to the instance + superuser_reserved_connections	The value specified is also used to restrict the number of connections from client applications and the number of connections for the management of instances.
	value	Refer to "When an Instance was Created with the initdb Command" in the Installation and Setup Guide for Server, and "Connections and Authentication" in the PostgreSQL Documentation, for details.
superuser_reserved_connection s	Add the number of simultaneous executions of mc_ctl status (*1) + 2	Specify the number of connections reserved for connections from database superusers.
		Add the number of connections from Mirroring Controller processes. Also reflect the added value in the max_connections parameter.
restart_after_crash	off	If "on" is specified, or the default value is used for this parameter, behavior equivalent to restarting FUJITSU Enterprise Postgres, including crash recovery, will be performed when some server processes end abnormally.
		However, when database multiplexing monitoring is used, a failover will occur after an error is detected when some server processes end abnormally, and the restart of those server processes is forcibly stopped. Specify "off" to

Parameter	Value set	Explanation
		prevent behavior such as this from occurring for no apparent reason.
synchronous_commit	on or remote_apply	Specify up to what position WAL send is to be performed before transaction commit processing returns a normal termination response to a client.  The recommended value is "on" or "remote_apply" to prevent data loss caused by operating system or server down immediately after a switch or switch.
recovery_target_timeline	latest	Specify "latest" so that the new standby server (original primary server) will follow the new primary server when a switch occurs.  This parameter is required when the original primary server is incorporated as a new standby server after the primary server is switched.

<sup>\*1:</sup> Number of simultaneous executions of the mc\_ctl command in the status mode.

# A.2 Parameters Set on the Standby Server

This section explains the content of the file and parameters set on the standby server. After editing postgresql.conf file, start the instance. The content for the parameters specified in postgresql.conf file is shown in the table below.

Table A.2 postgresql.conf file

Parameter	Value set	Explanation
wal_level	replica or logical	Specify the output level for the transaction log.
		Specify "logical" when logical decoding is also to be used.
max_wal_senders	2 or more	Specify "2" when building a Mirroring Controller cluster system.
		When additionally connecting asynchronous standby servers to the cluster system, add the number of simultaneous connections from these standby servers.
synchronous_standby_names	'primaryServerName'	Use single quotation marks (') to enclose the name that will identify the primary server. Any name can be specified.
		This name will be required to rebuild the original primary server as the new standby server after the primary server was switched.
		Do not change this parameter while Mirroring Controller is running.
		Do not specify multiple names to this parameter as the Mirroring Controller can manage only one standby server.
hot_standby	on	Specify whether queries can be run on the standby server.
		Specify "on".
wal_keep_segments	Number of file segments	If a delay exceeding the value set in this parameter occurs, the WAL segment required later by the standby server may be deleted by the primary server.

Parameter	Value set	Explanation
		Additionally, if you stop a standby server (for maintenance, for example), consider the stop time and set a value that will not cause the WAL segment to be deleted.
		Refer to "Estimating Transaction Log Space Requirements" in the Installation and Setup Guide for Server for information on estimating the WAL segment.
wal_log_hints	on	When using the pg_rewind command to recover a standby server, specify this parameter or enable checksums when executing the initdb command.
wal_sender_timeout	Timeout (milliseconds)	Specify the time period after which it is determined that the receiver process (walreceiver) of the transaction log is in an abnormal state on the primary server.
		The specified value must be larger than the value set for the wal_receiver_status_interval parameter set in the postgresql.conf file of the standby server.
		By aligning this value with the value for the database process heartbeat monitoring time, you can unify the time after which it is determined that an error has occurred.
wal_receiver_timeout	Timeout (milliseconds)	Specify the time period after which it is determined that an error has occurred when the transaction log was received on the standby server.
		By aligning this value with the value for the database process heartbeat monitoring time, you can unify the time after which it is determined that an error has occurred.
backup_destination	Backup data storage destination	Specify the name of the backup data storage directory.
	directory	Set the permissions so that only the instance administrator user can access the specified directory.
		Specify the same full path on all servers so that the backup data of other servers can be used to perform recovery.
archive_mode	on	Specify the archive log mode.
archive_command	'installDir/bin/pgx_walcopy.cmd "%p" "backupDataStorageDestinationD irectory/archived_wal/%f"	Specify the command and storage destination to save the transaction log.
listen_addresses	Standby server IP address, host name, or "*"	Specify the IP address or host name of the standby server. Specify the IP address or corresponding host name that will be used to connect to the log transfer network.
		The content specified is also used to allow connections from client applications.
		To receive the connection and the transaction log from any client or standby server, specify "*".
		Refer to "Connections and Authentication" in the PostgreSQL Documentation for details.

Parameter	Value set	Explanation
max_connections	Number of simultaneous client connections to the instance + superuser_reserved_connections	The value specified is also used to restrict the number of connections from client applications and the number of connections for the management of instances.
	value	Refer to "When an Instance was Created with the initdb Command" in the Installation and Setup Guide for Server, and "Connections and Authentication" in the PostgreSQL Documentation, for details.
superuser_reserved_connections	Add the number of simultaneous executions of mc_ctl status (*1)	Specify the number of connections reserved for connections from database superusers.
	+2	Add the number of connections from Mirroring Controller processes. Also reflect the added value in the max_connections parameter.
restart_after_crash	off	If "on" is specified, or the default value is used for this parameter, behavior equivalent to restarting FUJITSU Enterprise Postgres, including crash recovery, will be performed when some server processes end abnormally.
		However, when database multiplexing monitoring is used, a failover will occur after an error is detected when some server processes end abnormally, and the restart of those server processes is forcibly stopped. Specify "off" to prevent behavior such as this from occurring for no apparent reason.
synchronous_commit	on or remote_apply	Specify up to what position WAL send is to be performed before transaction commit processing returns a normal termination response to a client.
		The recommended value is "on" or "remote_apply" to prevent data loss caused by operating system or server down immediately after a switch or switch.
primary_conninfo	'streamingReplication ConnectionDestinationInfo'	Use single quotation marks (') to enclose the connection destination information of the streaming replication.
		The default value of this parameter is automatically set to postgresql.auto.conf in the procedure to run pg _ basebackup for instance setup.
recovery_target_timeline	latest	Specify "latest" so that the new standby server (original primary server) will follow the new primary server when a switch occurs.
		This parameter is required when the original primary server is incorporated as a new standby server after the primary server is switched.

# A.3 Network Configuration File

This section explains the network configuration file (network.conf) to be defined individually for the database servers and the arbitration server. Define the same content on the primary server and standby server.

For database multiplexing mode, define the network configuration for the following in network.conf.

- Integration between Mirroring Controller processes
- Integration between a Mirroring Controller process and the Mirroring Controller arbitration process

### Items to be defined in network.conf

#### Format:

```
serverIdentifier hostName[,hostName] portNum[,portNum] [serverType]
Or,
serverIdentifier ipAddr[,ipAddr] portNum[,portNum] [serverType]
```

Specify the server identifier, IP address or host name, port number, and server type, using a space as the delimiter.

The items are explained in the table below.

#### Table A.3 network.conf file

Item	Description	
serverIdentifier	Specify any identifier for the server. The maximum length is 64 bytes. Use ASCII characters excluding spaces and number signs (#) to specify this parameter.	
ipAddrOrHostName	Specify the IP address or its corresponding host name that will connect to the admin network that performs communication between the database servers, and to the arbitration network that performs communication between a database server and the arbitration server.	
	When specifying two IP addresses or host names delimited by a comma, do not insert a space after the comma.	
	Use ASCII characters excluding spaces to specify the host name.	
portNum	A port number cannot be specified if it exceeds the range 0 to 65535. Ensure that the port number does not conflict with other software. Do not specify an ephemeral port that may temporarily be assigned by another program. Note that the value specified in this parameter must also be set in the services file.	
	When specifying two port numbers delimited by a comma, do not insert a space after the comma.	
<i>serverType</i>	Specify "server" for a database server ("server" can be omitted), or "arbiter" for the arbitration server.	

### Content to be defined on the database servers

This section explains the network.conf content to be defined on the database servers.

The content to be defined depends on the operation settings at the time a heartbeat abnormality is detected.

When automatic degradation by the arbitration server is selected

- Specify definitions related to the admin network and arbitration network.
- Specify the IP address or host name and port number according to the server type (database server or arbitration server) as shown in the table below.

Server type	IP address or host name		Port nu	mber
	First	Second	First	Second
server	IP address or host name used as the admin network	IP address or host name used as the arbitration network (*1)	Port number used as the admin network	Port number used as the arbitration network (*1)
arbiter	IP address or host name of the arbitration server  Specify the same value as that specified in the my_address parameter of arbitration.conf on the arbitration server.	Not required	Port number on the arbitration server  Specify the same value as that specified in the port parameter of arbitration.conf on the arbitration server.	Not required

<sup>\*1:</sup> This value can be omitted from definitions not related to the local server. If it is omitted, network.conf must be created on both the primary server and standby server.

### Example)

IPv4

```
server1 192.0.2.100,192.0.3.100 27540,27541 server
server2 192.0.2.110,192.0.3.110 27540,27541 server
arbiter 192.0.3.120 27541 arbiter
```

IPv6

```
server1 2001:258:8404:1217:250:56ff:fea7:559f,2001:258:8404:1217:250:56ff:fea8:559f
27540,27541 server
server2 2001:258:8404:1217:250:56ff:fea7:55a0,2001:258:8404:1217:250:56ff:fea8:55a0
27540,27541 server
arbiter 2001:258:8404:1217:250:56ff:fea8:55a0 27541 arbiter
```

When operation other than automatic degradation by the arbitration server is selected

- Specify definitions related to the admin network.
- Define the same content on the primary server and standby server.
- Define lines for database servers only.
- Specify only one IP address or host name and port number.

IP address or host name		Port number	
First	Second	First	Second
IP address or host name to be used as the admin network	Not required	Port number used as the admin network	Not required

#### Example)

The literal space represents a space.

IPv4

```
server1 192.0.2.100 27540
server2 192.0.2.110 27540
```

IPv6

```
server1 2001:258:8404:1217:250:56ff:fea7:559f 27540
server2 2001:258:8404:1217:250:56ff:fea7:55a0 27540
```

#### Content to be defined on the arbitration server

This section explains the network.conf content to be defined on the arbitration server.

- Specify definitions related to the arbitration network.
- Define lines for database servers only.
- For the IP address or host name, specify the same value as the second IP address or host name specified in the database server line in network.conf of the database server.

- For the port number, specify the same value as the second port number specified in the database server line in network.conf of the database server.

#### Example)

The literal space represents a space.

IPv4

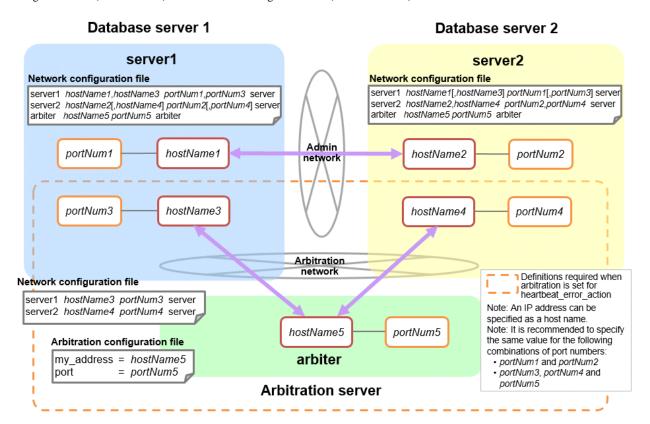
```
server1 192.0.3.100 27541
server2 192.0.3.110 27541
```

IPv6

```
server1 2001:258:8404:1217:250:56ff:fea8:559f 27541
server2 2001:258:8404:1217:250:56ff:fea8:55a0 27541
```

### Relationship between network-related definitions

Refer to the diagram below for the relationship between the host names and IP addresses or port numbers specified in the network configuration file (network.conf) and arbitration configuration file (arbitration.conf).



# A.4 Server Configuration File

# A.4.1 Server Configuration File for the Database Servers

Define the information related to Mirroring Controller monitoring and control in the serverIdentifier.conf file. The maximum length of the server identifier is 64 bytes. Use ASCII characters excluding spaces to specify this parameter.

If the primary server and standby server environments are different, define content that is different, according to the environment.

Table A.4 serverldentifier.conf file

Parameter	Value set	Explanation
db_instance	'dataStorageDestinationDir' [Example]	Specify using single quotation marks (') to enclose the data storage destination directory used to identify the monitoring target instance.
	db_instance = '/database1/inst1'	Use ASCII characters to specify this parameter.
target_db	postgres or template1	Specify the name of the database to be connected to the database instance.
		The default is "postgres".
db_instance_username	'usernameToConnectToDbInstance'	Specify the username to connect to the database instance.
		Use ASCII characters to specify this parameter.
		Specify this parameter if the database administrator user is different from the operating system user who starts Mirroring Controller.
		Enclose the username of the database superuser in single quotation marks (').
		The maximum length of the username is 63 bytes.
		The default is the operating system user who starts Mirroring Controller.
db_instance_password	'passwordOfInstanceAdminUser'	Specify the password used when Mirroring Controller connects to a database instance, enclosed in single quotation marks (').
		Use ASCII characters to specify this parameter.
		If password authentication is performed, you must specify this parameter in the settings used when Mirroring Controller connects to a database instance.
		If you specify this parameter when password authentication is not performed, the parameter will be ignored.
		If the specified value of this parameter includes ' or  write \' or \ respectively.
enable_hash_in_password	on or off	Specify on to treat the # in the db_instance_password specification as a password character, or off to treat it as a comment.
		The default is "off".
core_file_path	'coreFileOutputDir'	Specify the directory to which the core file is to be output, enclosed in single quotation marks (').
		Use ASCII characters to specify this parameter.
		If this parameter is omitted, it will be assumed that the Mirroring Controller management directory was specified.
syslog_facility	Specify LOCAL0, LOCAL1, LOCAL2, LOCAL3, LOCAL4,	When the import of logs to the syslog is enabled, the value of this parameter will be used for "facility" of the syslog.
	LOCAL5, LOCAL6, or LOCAL7.	The default is "LOCAL0".
syslog_ident (*1)	'programName'	Specify using single quotation marks (') to enclose the program name used to identify the Mirroring Controller message in the system log.

Parameter	Value set	Explanation
		Use ASCII characters excluding spaces to specify this parameter.
		The default is 'MirroringControllerOpen'.
remote_call_timeout	Admin communication timeout	Specify the timeout value (milliseconds) of the Mirroring Controller agent process for communication between servers.
		Specify a value between 0 and 2147483647 to be less than the operation system TCP connection timeout (*2).
		The value 0 indicates that there is no timeout limit.
		The default is 70000 milliseconds (70 seconds).
agent_alive_timeout	Timeout for Mirroring Controller process heartbeat monitoring	If there is no response for at least the number of seconds specified, the Mirroring Controller process is restarted.
	(seconds)	Specify 0 or a value between 2 and 2147483647. The value 0 indicates that there is no timeout limit.
		The default is 0 seconds.
heartbeat_error_action	Operation when a heartbeat abnormality is detected using	arbitration: Perform automatic degradation using the arbitration server.
	operating system or server heartbeat monitoring	command: Call a user exit to determine degradation, and perform automatic degradation if required.
		message: Notify messages.
		fallback: Perform automatic degradation unconditionally.
		The default is "arbitration".
		Set the same value on the primary server and standby server.
heartbeat_interval	Interval time for abnormality monitoring during heartbeat	Abnormality monitoring of the operating system or server is performed at the interval specified in heartbeat_interval.
	monitoring of the operating system or server (milliseconds)	If an error is detected, operation will conform to the value specified for heartbeat_error_action. If "arbitration" is specified in heartbeat_error_action, the error detection time during monitoring of the operating system or server becomes longer than when the arbitration server is not used, by up to the value specified for arbitration_timeout.
		Specify a value between 1 and 2147483647.
		The specified value is used as the default for db_instance_check_interval and disk_check_interval.
		The default is 800 milliseconds.
heartbeat_timeout	Timeout for abnormality monitoring during heartbeat monitoring of the operating system or server (seconds)	If there is no response for at least the number of seconds specified, it will be assumed that an error has occurred that requires the primary server to be switched, or the standby server to be disconnected.
		If an error is detected, operation will conform to the value specified for heartbeat_error_action. If "arbitration" is specified in heartbeat_error_action, the error detection time during monitoring of the operating system or server becomes longer than when the arbitration server is not used, by up to the value specified for arbitration_timeout.

Parameter	Value set	Explanation
		Specify a value between 1 and 2147483647.
		The specified value is used as the default for db_instance_check_timeout.
		The default is 1 second.
heartbeat_retry	Number of retries for abnormality monitoring during heartbeat monitoring of the operating system or server (number of times)	Specify the number of retries to be performed when an error has been detected that requires the primary server to be switched, or the standby server to be disconnected. If an error is detected in succession more than the specified number of times, switch or disconnection will be performed.
		If an error is detected, operation will conform to the value specified for heartbeat_error_action. If "arbitration" is specified in heartbeat_error_action, the error detection time during monitoring of the operating system or server becomes longer than when the arbitration server is not used, by up to the value specified for arbitration_timeout.
		Specify a value between 0 and 2147483647.
		The specified value is used as the default for db_instance_check_retry and disk_check_retry.
		The default is 2 times.
db_instance_check_interval	Database process heartbeat monitoring interval (milliseconds)	Heartbeat monitoring of the database process is performed at the interval specified in db_instance_check_interval.
		This parameter setting is also used for abnormality monitoring of streaming replication.
		Specify a value between 1 and 2147483647.
		The default is the value set for heartbeat_interval.
db_instance_check_timeout	Database process heartbeat monitoring timeout (seconds)	If there is no response for at least the number of seconds specified, it will be assumed that an error has occurred that requires the primary server to be switched, or the standby server to be disconnected.
		Specify a value between 1 and 2147483647.
		The default is the value set for heartbeat_timeout.
db_instance_check_retry	Number of retries for database process heartbeat monitoring (number of times)	Specify the number of retries to be performed when an error has been detected that requires the primary server to be switched, or the standby server to be disconnected. If an error is detected in succession more than the specified number of times, switch or disconnection will be performed. However, if it detects that the database process is down, it will immediately switch or disconnect regardless of the setting of this parameter.
		This parameter setting is also used for abnormality monitoring of streaming replication.
		Specify a value between 0 and 2147483647.
		The default number of retries is the value set for heartbeat_retry.
db_instance_timeout_action	none, message, or failover	Specify the behavior for no-response monitoring of the instance.

Parameter	Value set	Explanation
		none: Do not perform no-response monitoring.
		message: Notify messages if an error is detected during no- response monitoring.
		failover: Perform automatic degradation if an error is detected during no-response monitoring.
		The default is "failover".
disk_check_interval	Interval time for disk abnormality monitoring (milliseconds)	Abnormality monitoring of disk failure is performed at the interval specified in disk_check_interval. If the file cannot be created, it will be assumed that an error has occurred that requires the primary server to be switched, or the standby server to be disconnected.
		Specify a value between 1 and 2147483647.
		The default is the value set for heartbeat_interval.
disk_check_retry	Number of retries for disk abnormality monitoring (number of times)	Specify the number of retries to be performed when an error has been detected that requires the primary server to be switched, or the standby server to be disconnected.
		If an error is detected in succession more than the specified number of times, switch or disconnection will be performed.
		Specify a value between 0 and 2147483647.
		The default number of retries is the value set for heartbeat_retry.
tablespace_directory_error_ action	message or failover	Specify the behavior to be implemented if an error is detected in the tablespace storage directory.
		message: Notify messages.
		failover: Perform automatic degradation.
		The default is "failover".
arbiter_alive_interval	Interval time for monitoring connection to the Mirroring Controller	A heartbeat is sent to the Mirroring Controller arbitration process at the specified interval.
	arbitration process (milliseconds)	Specify a value between 1 and 2147483647.
		The default is 16000 milliseconds.
		This parameter does not need to be set for operation that does not use the arbitration server.
arbiter_alive_timeout	Timeout for monitoring connection to the Mirroring Controller arbitration process (seconds)	If the heartbeat does not respond within the specified number of seconds, the Mirroring Controller arbitration process is determined to have been disconnected, a message is output, and reconnection is attempted.
		Specify a value between 1 and 2147483647.
		The default is 20 seconds.
		This parameter does not need to be set for operation that does not use the arbitration server.
arbiter_alive_retry	Number of retries for monitoring connection to the Mirroring Controller arbitration process (number of times)	Specify the number of heartbeat retries to be performed if an error is detected in the heartbeat to the Mirroring Controller arbitration process. If the heartbeat does not respond within the specified number of retries, the

Parameter	Value set	Explanation
		Mirroring Controller arbitration process is determined to have been disconnected.
		Specify a value between 0 and 2147483647.
		The default is 0 times.
		This parameter does not need to be set for operation that does not use the arbitration server.
arbiter_connect_interval	Attempt interval for connection to the Mirroring Controller arbitration process (milliseconds)	Reconnection is attempted at the specified interval if connection fails at startup of the Mirroring Controller process or if the Mirroring Controller arbitration process is disconnected.
		Specify a value between 1 and 2147483647.
		The default is 16000 milliseconds.
		This parameter does not need to be set for operation that does not use the arbitration server.
arbiter_connect_timeout	Timeout for connection to the Mirroring Controller arbitration process (seconds)	If reconnection at startup of the Mirroring Controller process or after disconnection of the Mirroring Controller arbitration process does not succeed within the specified number of seconds, connection to the Mirroring Controller arbitration process is determined to have failed and reconnection is attempted.
		Specify a value between 1 and 2147483647.
		The default is 20 seconds.
		This parameter does not need to be set for operation that does not use the arbitration server.
fencing_command	'fencingCmdFilePath'  [Setting example]  fencing_command = '/mc/fencing_dir/	Specify the full path of the fencing command that fences a database server where an error is determined to have occurred.
	execute_fencing.sh'	Enclose the path in single quotation marks (').
		Specify the path using less than 1024 bytes.
		This parameter must be specified when "command" is set for heartbeat_error_action.
fencing_command_timeout	Fencing command timeout (seconds)	If the command does not respond within the specified number of seconds, fencing is determined to have failed and a signal (SIGTERM) is sent to the fencing command execution process.
		Specify a value between 1 and 2147483647.
		The default is 20 seconds.
arbitration_timeout	Arbitration processing timeout in the Mirroring Controller arbitration process (seconds)	The specified value must be at least equal to the value of fencing_command_timeout in the arbitration configuration file, which is the heartbeat monitoring time of the operating system or server.
		If there is no response for at least the number of seconds specified, the primary server will not be switched and the standby server will not be disconnected. Therefore, perform degradation manually.
		If the heartbeat_interval, heartbeat_timeout, and heartbeat_retry values are specified in arbitration.conf for

Parameter	Value set	Explanation
		the arbitration server, use the arbitration server values to design arbitration_timeout.
		Specify a value between 1 and 2147483647.
		The default is 30 seconds.  This parameter does not need to be set for operation that does not use the arbitration server.
arbitration_command	'arbitrationCmdFilePath'  [Setting example]  arbitration_command = '/mc/ arbitration_dir/ execute_arbitration_command.sh'	Specify the full path of the arbitration command to be executed when an abnormality is detected during heartbeat monitoring of the operating system or server. Enclose the path in single quotation marks (').  Specify the path using less than 1024 bytes.  This parameter must be specified when "command" is set for heartbeat_error_action.
arbitration_command_timeo ut	Arbitration command timeout (seconds)	If the arbitration command does not respond within the specified number of seconds, it is determined that execution of the arbitration command has failed and a signal (SIGTERM) is sent to the arbitration command execution process.
		Specify a value between 1 and 2147483647.
		The default is 30 seconds.
		This parameter can be specified only when "command" is set for heartbeat_error_action.
shutdown_detached_synchr onous_standby	on or off	Specify whether to forcibly stop the instance on the standby server when the standby server is disconnected.
		on: Stop the instance.
		off: Do not stop the instance.
		If "on" is specified and the pre-detach command was created, the pre-detach command is executed and then the instance is stopped.
		The default is "off".
post_switch_command	'postSwitchCmdFilePath' [Setting example]	Specify the full path of the command to be called by Mirroring Controller after a new primary server is promoted during a failover of the primary server.
	post_switch_command = '/mc/ status_change/	Enclose the path in single quotation marks (').
	execute_post_switch.sh'	Specify the path using less than 1024 bytes.
post_attach_command	'postAttachCmdFilePath' [Setting example]	Specify the full path of the command to be called by Mirroring Controller after the standby server is attached to the cluster system.
	post_attach_command = '/mc/ status_change/ execute_post_attach.sh'	Enclose the path in single quotation marks (').  Specify the path using less than 1024 bytes.
pre_detach_command	'preDetachCmdFilePath'  [Setting example]  pre_detach_command = '/mc/ status_change/execute_pre_detach.sh'	Specify the full path of the command to be called by Mirroring Controller before the standby server is disconnected from the cluster system.  Enclose the path in single quotation marks (').  Specify the path using less than 1024 bytes.

Parameter	Value set	Explanation
status_change_command_ti meout	State transition command timeout (seconds)	Specify the timeout value of the post-switch command, post-attach command, and pre-detach command. If the command does not respond within the specified number of seconds, a signal (SIGTERM) is sent to the execution process of the status change command.  Specify a timeout between 1 and 2147483647.
		The default is 20 seconds.
check_synchronous_standb y_names_validation	on or off	Specify whether Mirroring Controller is to periodically check during operations whether the synchronous_standby_names parameter in postgresql.conf was changed by an incorrect user operation.
		However, it is not recommended to enable this parameter, because performing this check causes Mirroring Controller to use the CPU of the database server redundantly and execute SQL statements at high frequency.
		This parameter is compatible with operations in FUJITSU Enterprise Postgres 9.6 or earlier.
		The default is "off".
db_instance_ext_pq_connin fo	'libpqConnectionSSLParmToConnect ToDbinstance'	Specify, in key-value form, the connection parameter for libpq that Mirroring Controller adds when connecting to a database. The connection parameters you can specify are those related to SSL. Use ASCII characters to specify this parameter.
		The connection parameter specified in this parameter must also be specified in the db_instance_ext_jdbc_conninfo.
db_instance_ext_jdbc_conni nfo	'JDBCConnectionSSLParmToConne ctToDbinstance'	Specify, in URI form, the connection parameter for JDBC that Mirroring Controller adds when connecting to a database. The connection parameters you can specify are those related to SSL. Use ASCII characters to specify this parameter.
		The connection parameter specified in this parameter must also be specified in the db_instance_ext_pq_conninfo.

<sup>\*1:</sup> By specifying the syslog\_ident parameter of the postgresql.conf file, the Mirroring Controller output content can be referenced transparently, so log reference is easy.

The availability of some parameters depends on the value set for the heartbeat\_error\_action parameter that sets the operation to be performed if heartbeat monitoring of the operating system or server detects a heartbeat abnormality.

Table A.5 Parameter availability depending on the value set for the heartbeat\_error\_action parameter

Parameter	Value set			
Falametei	arbitration	command	message	fallback
arbiter_alive_interval	Y	N	N	N
arbiter_alive_timeout	Y	N	N	N
arbiter_alive_retry	Y	N	N	N
arbiter_connect_interval	Y	N	N	N

<sup>\*2:</sup> The operating system TCP connection timeout period is determined by the kernel parameter tcp\_syn\_retries. The remote\_call\_timeout parameter must be set to a value that is shorter than the timeout period for the operating system TCP connection timeout, so change either parameter as necessary.

Parameter	Value set			
Faiametei	arbitration	command	message	fallback
arbiter_connect_timeout	Y	N	N	N
arbitration_timeout	Y	N	N	N
arbitration_command	N	R	N	N
arbitration_command_timeout	N	Y	N	N
fencing_command	Y	R	Y	N
fencing_command_timeout	Y	Y	Y	N
shutdown_detached_synchronous_standby	Y	Y	N	N

R: Required

Y: Can be specified

N: Cannot be specified

# A.4.2 Arbitration Configuration File

In arbitration.conf, define the information related to arbitration and control of the Mirroring Controller arbitration process.

### Linux

Table A.6 arbitration.conf file (Linux)

Parameter	Value set	Description
port	Port number of the Mirroring Controller arbitration process	The specified value must not exceed the range 0 to 65535. Ensure that the port number does not conflict with other software. Do not specify an ephemeral port that may temporarily be assigned by another program.  For the port number of the arbitration server to be specified in network.conf on the database server, specify the same value as the port number specified
		in this parameter.
my_address	'ipAddrOrHostNameThatAcceptsConn ectionFromMirroringControllerProces sOnDbServer' [Setting example] my_address = '192.0.3.120'	For the IP address or host name of the arbitration server to be specified in network.conf on the database server, specify the same value as the IP address or host name specified in this parameter.  IPv4 and IPv6 addresses can be specified.
		Specify the IP address or host name, enclosed in single quotation marks (').
core_file_path	'coreFileOutputDir'	Specify the directory to which the core file is to be output, enclosed in single quotation marks ('). Use ASCII characters to specify this parameter.
		If this parameter is omitted, it will be assumed that the Mirroring Controller arbitration process management directory was specified.
syslog_facility	Specify LOCAL0, LOCAL1, LOCAL2, LOCAL3, LOCAL4, LOCAL5, LOCAL6, or LOCAL7.	When the import of logs to the syslog is enabled, the value of this parameter will be used for "facility" of the syslog.
		The default is "LOCAL0".

Parameter	Value set	Description
syslog_ident	'programName'	Specify using single quotation marks (') to enclose the program name used to identify the Mirroring Controller arbitration process message in the system log. Use ASCII characters excluding spaces to specify this parameter.
		The default is 'MirroringControllerArbiter'.
fencing_command	'fencingCmdFilePath'  [Setting example]  fencing_command = '/arbiter/ fencing_dir/execute_fencing.sh'	Specify the full path of the fencing command that fences a database server where an error is determined to have occurred.  Enclose the path in single quotation marks (').  Specify the path using less than 1024 bytes.
fencing_command_timeout	Fencing command timeout (seconds)	If the command does not respond within the specified number of seconds, fencing is determined to have failed and a signal (SIGTERM) is sent to the fencing command execution process.  Specify a value between 1 and 2147483647.  The default is 20 seconds.
heartbeat_interval(*1)	Interval time for heartbeat monitoring of the operating system or server (milliseconds)	The heartbeat monitoring of the database server is checked at the specified interval and arbitration is performed.
		Specify a value between 1 and 2147483647.
		The default is the value specified in serverIdentifier.conf of the database server.
		Specify this parameter to perform optimization taking into account differences in the line load to the admin network and the reduction in the time it takes to degrade.
heartbeat_timeout	Timeout for heartbeat monitoring of the operating system or server (seconds)	If there is no response for at least the number of seconds specified, it will be assumed that an error has occurred that requires the primary server or standby server to be fenced.
		Specify a value between 1 and 2147483647.
		The default is the value specified in <i>serverIdentifier</i> .conf of the database server.
		Specify this parameter to perform optimization taking into account differences in the line load to the admin network and the reduction in the time it takes to degrade.
heartbeat_retry	Number of retries for heartbeat monitoring of the operating system or server (number of times)	Specify the number of retries to be performed when an error has been detected that requires the primary server or standby server to be fenced.
		If an error is detected in succession more than the specified number of times, fencing will be performed.
		Specify a value between 0 and 2147483647.
		The default is the value specified in serverIdentifier.conf of the database server.

Parameter	Value set	Description
		Specify this parameter to perform optimization taking into account differences in the line load to the admin network and the reduction in the time it takes to degrade.

<sup>\*1:</sup>Refer to "2.11.4 Tuning for Optimization of Degradation Using Abnormality Monitoring" for information on the tuning parameters for operating system or server abnormality monitoring when using an arbitration server.

### Windows

Table A.7 arbitration.conf file (Windows)

Parameter	Value set	Description
port	Port number of the Mirroring Controller arbitration process	The specified value must not exceed the range 0 to 65535. Ensure that the port number does not conflict with other software. Do not specify an ephemeral port that may temporarily be assigned by another program.
		For the port number of the arbitration server to be specified in network.conf on the database server, specify the same value as the port number specified in this parameter.
my_address	'ipAddrOrHostNameThatAcceptsConn ectionFromMirroringControllerProces sOnDbServer' [Setting example] my_address = '192.0.3.120'	For the IP address or host name of the arbitration server to be specified in network.conf on the database server, specify the same value as the IP address or host name specified in this parameter.  IPv4 and IPv6 addresses can be specified.  Specify the IP address or host name, enclosed in single quotation marks (').
core_file_path	'coreFileOutputDir'	Specify the directory to which the core file is to be output, enclosed in single quotation marks ('). Use ASCII characters, and specify "\\" as the path delimiter.
		If this parameter is omitted, it will be assumed that the Mirroring Controller arbitration process management directory was specified.
service_name	'registeredServiceNameOfMirroringCo ntrollerArbitrationProcess'	Specify the Mirroring Controller arbitration process service name to be registered as a Windows service, enclosed in single quotation marks ('). Use ASCII characters excluding forward slash (/) and backslash (\) to specify this parameter.
		The service name is up to 124 bytes.
event_source	'eventSourceName'	The default is 'MirroringControllerArbiter'.  Specify the event source name to be used to identify the Mirroring Controller arbitration process message in the event log, enclosed in single quotation marks ('). Use ASCII characters to specify this parameter.
		The maximum length of the event source name is 255 bytes.
		The default is 'MirroringControllerArbiter'.

Parameter	Value set	Description
fencing_command	'fencingCmdFilePath' [Setting example]	Specify the full path of the fencing command that fences a database server where an error is determined to have occurred.
	fencing_command = 'c:\\arbiter\ \fencing_dir\\execute_fencing.bat'	Specify "\\" as the delimiter.
		Enclose the path in single quotation marks (').
		Specify the path using less than 260 bytes.
		Any multibyte characters must use the same encoding as the operating system.
fencing_command_timeout	Fencing command timeout (seconds)	If the command does not respond within the specified number of seconds, fencing is determined to have failed and a signal (SIGTERM) is sent to the fencing command execution process.
		Specify a value between 1 and 2147483647.
		The default is 20 seconds.
heartbeat_interval(*1)	Interval time for heartbeat monitoring of the operating system or server (milliseconds)	The heartbeat monitoring of the database server is checked at the specified interval and arbitration is performed.
		Specify a value between 1 and 2147483647.
		The default is the value specified in serverIdentifier.conf of the database server.
		Specify this parameter to perform optimization taking into account differences in the line load to the admin network and the reduction in the time it takes to degrade.
heartbeat_timeout	Timeout for heartbeat monitoring of the operating system or server (seconds)	If there is no response for at least the number of seconds specified, it will be assumed that an error has occurred that requires the primary server or standby server to be fenced.
		Specify a value between 1 and 2147483647.
		The default is the value specified in serverIdentifier.conf of the database server.
		Specify this parameter to perform optimization taking into account differences in the line load to the admin network and the reduction in the time it takes to degrade.
heartbeat_retry	Number of retries for heartbeat monitoring of the operating system or server (number of times)	Specify the number of retries to be performed when an error has been detected that requires the primary server or standby server to be fenced.
		If an error is detected in succession more than the specified number of times, fencing will be performed.
		Specify a value between 0 and 2147483647.
		The default is the value specified in serverIdentifier.conf of the database server.
		Specify this parameter to perform optimization taking into account differences in the line load to

Parameter	Value set	Description
		the admin network and the reduction in the time it takes to degrade.

<sup>\*1:</sup>Refer to "2.11.4 Tuning for Optimization of Degradation Using Abnormality Monitoring" for information on the tuning parameters for operating system or server abnormality monitoring when using an arbitration server.

## Appendix B Supplementary Information on Building the Primary Server and Standby Server on the Same Server

The primary server and standby server can be pseudo-configured on the same server for system testing, for example. Out of consideration for performance and reliability, do not use this type of configuration for any other purposes. For this reason, do not use this type of configuration in a production environment.

Note that the setup and operations is the same as if the primary and standby servers are built on different servers.

This appendix provides supplementary information explaining how to configure the primary server and standby server on the same server.



Even if automatic degradation by an arbitration server is set when the primary server and standby server are configured on the same server, there will be no effect of it.

## **B.1 Backup Data Storage Destination Directory**

It is not a problem if the same backup data storage destination directory is used on the primary server and standby server.

## B.2 How to Execute the mc\_ctl Command

When executing the mc\_ctl command, specify the server identifier in the --local-server option in order to identify the operation destination server.

Below is an example of starting Mirroring Controller of the server "server1" defined in the network.conf file. For mc\_ctl command operations using another mode, also specify the --local-server option.

Define two server identifiers for the same IP address with different port numbers in the network.conf file.

Example)

```
server1 192.0.2.100 27540
server2 192.0.2.100 27541
```

Ensure that the port numbers of both primary server and standby server do not conflict with any other software.

Enabling automatic switch/disconnection

Start Mirroring Controller of the server "server1":

Example)

```
$ mc_ctl start -M /mcdir/inst1 --local-server server1
```

Stop Mirroring Controller of the server "server1":

Example)

```
$ mc_ctl stop -M /mcdir/inst1 --local-server server1
```

Disabling automatic switch/disconnection

Start Mirroring Controller of the server "server1":

Example)

```
$ mc_ctl start -M /mcdir/inst1 -F --local-server server1
```

Stop Mirroring Controller of the server "server1":

Example)

\$ mc\_ctl stop -M /mcdir/inst1 --local-server server1



#### If using an rc script

Append the --local-server option to the mc\_ctl command call defined in the automatic start shell script.

#### If using systemd

Add the --local-server option to the mc\_ctl option specification for ExecStart and ExecStop of the unit file for systemd.

Refer to "2.12 Setting Automatic Start and Stop of Mirroring Controller and Multiplexed Instances" for details.

## Appendix C User Commands

This appendix describes three categories of commands:

- Fencing command
- Arbitration command
- State transition commands

This appendix describes each category of user command.

## C.1 Fencing Command

#### **Format**

The syntax for calling the fencing command from the Mirroring Controller process or the Mirroring Controller arbitration process is described below.

Fencing command of the database server

fencingCmd executionMode mcDegradationOper cmdServerId targetServerId primarycenter

Fencing command of the arbitration server

fencingCmd executionMode mcDegradationOper targetServerId

#### Input

Fencing command of the database server

Execution mode

monitor: Detect issues via automatic monitoring of the Mirroring Controller process

command: Mirroring Controller command execution (switch mode or detach mode of the mc\_ctl command)

Degradation operation to be performed by Mirroring Controller

switch: Switch detach: Disconnect

cmdServerId

ID of the database server that called the command

targetServerId

ID of the database server to be fenced

primarycenter

Fixed value

Fencing command of the arbitration server

Execution mode

monitor: Detect issues via automatic monitoring of the Mirroring Controller process

command: Mirroring Controller command execution (switch mode or detach mode of the mc\_ctl command)

Degradation operation to be performed by Mirroring Controller

switch: Switch detach: Disconnect

targetServerId

ID of the database server to be fenced

#### Output

#### Return value

0: Mirroring Controller will continue the degradation process.

Other than 0: Mirroring Controller will cancel the degradation process.

#### Description

Identifies the database server targeted for fencing based on the input server identifier, and implements the process that isolates it from the cluster system.

#### **Notes**

- The command is executed by the operating system user who started Mirroring Controller or the Mirroring Controller arbitration process. Therefore, if the command is to be executed by a specific operating system user, change the executing user of the command accordingly.
- The operating system user who started Mirroring Controller or the Mirroring Controller arbitration process must have execution privileges to the command. Otherwise, the degradation process will be canceled.
- From a security point of view, set the access privileges as necessary so that the fencing command is not overwritten and unauthorized operations are not performed by unintended operating system users.
- If the fencing command returns a value other than 0, Mirroring Controller will cancel the degradation process, so it is necessary for the user to check the status of the server, and switch or disconnect it manually.
- Before executing the fencing command, check if the server is already fenced, to avoid the command terminating abnormally.
- If the command times out, Mirroring Controller will stop the command, output an error message, and cancel the degradation process.

#### Information

The fencing command can be implemented by simply stopping the operating system or server. For example, if stopping the power for the database server, it is possible to use a utility to control the hardware control board in environments equipped with boards compatible with IPMI hardware standard.

#### Linux

Below is a sample script of a fencing command that powers off the database server using the IPMI tool.

Sample shell script

/installDir/share/mcarb\_execute\_fencing.sh.sample

#### Windows

Below is a sample script of a fencing command that powers off the database server using IPMIUTIL.

Sample shell script

installDir\share\mcarb\_execute\_fencing.bat.sample

#### C.2 Arbitration Command

#### **Format**

The syntax for calling the arbitration command from the Mirroring Controller process is described below.

arbitrationCmd cmdServerId targetServerId primarycenter

#### Input

cmdServerId

ID of the database server that called the command

targetServerId

ID of the database server to arbitrate

primarycenter

Fixed value

#### Output

Return value

0: The database server to arbitrate has an issue, and Mirroring Controller will continue the degradation process.

Other than 0: The database server to arbitrate is normal, and Mirroring Controller will cancel the degradation process.

#### Description

Identifies the database server to arbitrate based on the input server identifier, and checks the status of the server.

#### **Notes**

- The command is executed by the operating system user who started Mirroring Controller.
- The operating system user who started Mirroring Controller must have execution privileges to the command. Otherwise, the command will not be called, and the degradation process will be canceled.
- If the command times out, Mirroring Controller will stop the command, output an error message, and cancel the degradation process.

### C.3 State Transition Commands

State transition commands include the three types of user commands below. Any of the commands can be implemented by Mirroring Controller in conjunction with database server status transitions.

- Post-switch command
- Pre-detach command
- Post-attach command

#### C.3.1 Post-switch Command

#### **Format**

The syntax for calling the post-switch command from the Mirroring Controller process is described below.

postswitchCmd serverIdentifier primarycenter

#### Input

serverldentifier

ID of the database server (new primary server) that was switched

primarycenter

Fixed value

#### Output

Return value

None

#### Notes

- The command is executed by the operating system user who started Mirroring Controller.
- The operating system user who started Mirroring Controller must have execution privileges to the command. Otherwise, the command will not be called.
- If the command times out, Mirroring Controller will stop the command, output an error message, and cancel the process.

#### C.3.2 Pre-detach Command

#### **Format**

The syntax for calling the pre-detach command from the Mirroring Controller process is described below.

predetachCmd cmdServerId serverRole targetServerId primarycenter

#### Input

#### cmdServerId

ID of the database server that called the command

#### Server role

Role of the database server that called the command

primary: Primary standby: Standby

#### targetServerId

ID of the standby server to be disconnected from the cluster system

primarycenter

Fixed value

#### Output

Return value

None

#### **Notes**

- The command is executed by the operating system user who started Mirroring Controller.
- The operating system user who started Mirroring Controller must have execution privileges to the command. Otherwise, the command will not be called, however, Mirroring Controller will output an error message and continue the process.
- If the command times out, Mirroring Controller will stop the command, output an error message, and cancel the process.

#### C.3.3 Post-attach Command

#### **Format**

The syntax for calling the post-attach command from the Mirroring Controller process is described below.

postattachCmd cmdServerId serverRole targetServerId primarycenter

#### Input

#### cmdServerId

ID of the database server that called the command

#### Server role

Role of the database server that called the command

primary: Primary standby: Standby

#### targetServerId

ID of the standby server to be attached to the cluster system

primarycenter

Fixed value

#### Output

Return value

None

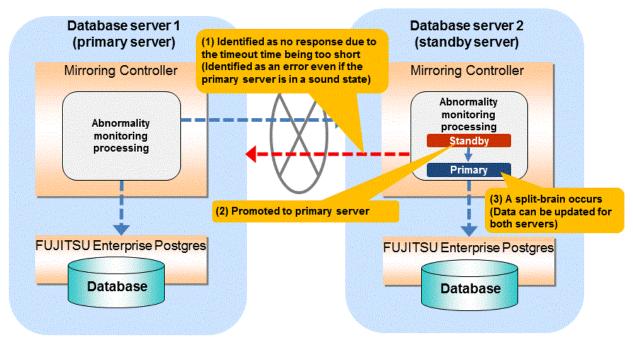
#### **Notes**

- The command is executed by the operating system user who started Mirroring Controller.
- The operating system user who started Mirroring Controller must have execution privileges to the command. Otherwise, the command will not be called.
- If the command times out, Mirroring Controller will stop the command, output an error message, and cancel the process.

## Appendix D Notes on Performing Automatic Degradation Immediately after a Heartbeat Abnormality

The type of issue below occurs if automatic degradation is performed unconditionally after an issue is detected during heartbeat monitoring of an operating system or server, and heartbeat monitoring was not properly tuned.

#### • If the timeout time is too short



#### If the timeout time is too long Database server 1 Database server 2 (primary server) (standby server) Mirroring Controller Mirroring Controller (1) The operating system or server crashes or is unresponsive Abnormality ity itoring monitoring p cessi processing Failure occurs (2) Errors cannot be detected in realtime due to the timeout time being too long FUJITSU Enterprise Postgres FUJITSU Enterprise Postgres **Database Database** (3) Operations cannot Application server continue until the server is switched Job application

#### Notes on monitoring when the operating system or server crashes or is unresponsive

As illustrated in the diagram above, timeout is used to monitor whether the operating system or server crashes or is unresponsive. Therefore, if tuning has not been performed correctly, there is a risk of a split-brain mistakenly occurring even if the server is in a sound state.

Split-brain is a phenomenon in which both servers temporarily operate as primary servers, causing data updates to be performed on both servers.

#### Split-brain detection method

It can be confirmed that split-brain occurs under the following conditions:

- 1. When the mc\_ctl command is executed in status mode on both servers, the "host\_role" of both servers is output as "primary", and
- 2. The following message is output to the system log of one of the servers:

```
promotion processing completed (MCA00062)
```

#### How to recover from a split-brain

Use the procedure described below. Note that the new primary server is the server that was confirmed in step 2 of the aforementioned detection method.

- 1. Stop all applications that are running on the old and new primary servers.
- Investigate and recover the database.Investigate the update results that have not been reflected to the new primary server from the database of the old primary server, and apply to the new primary server as necessary.
- 3. Stop the old primary server instance and the Mirroring Controller.

- 4. Resume the applications that were stopped in step 1.
- Recover the old primary server.
   While referring to "2.5 Setting Up the Standby Server", build (set up) the old primary server as the new standby server, from the new primary server.

## Appendix E Supplementary Procedure on Configuring for Operation in Database Multiplexing Mode

This appendix explains a supplementary procedure on the configuration required for operation in database multiplexing mode.

## **E.1 Security Policy Settings**

This section explains how to configure the security settings to enable an operating system user account designated as an instance administrator user to log on as a service.

1. Displaying the [Local Security Policy] window

In Windows, select [Administrative Tools], and then click [Local Security Policy].

- 2. Setting up security
  - 1. In the [Local Security Policy] window, select [Security Settings], select [Local Policies], and then click [User Rights Assignment].
  - 2. Under [Policy] in the [User Rights Assignment] window, double-click [Log on as a service].
  - 3. In the [Log on as a service Properties] window, set the following:
    - a. Select the [Local Security Setting] tab.
    - b. On the [Local Security Setting] tab, click [Add User or Group].
    - c. In the [Select Users or Groups] window, enter the operating system user account of the instance administrator user in [Enter the object names to select].
    - d. Click [OK].
  - 4. In the [Log on as a service Properties] window, click [OK].
  - 5. From the [Local Security Policy] tree, click [Local Policies], and then double-click [Security Options].
  - Scroll down and double-click [User Account Control: Behavior of the elevation prompt for administrators in Admin Approval Mode].
  - 7. From the drop-down menu, select the "Elevate without prompting" in the [Local Security Setting] tab.
  - 8. Click [OK].

## **E.2 Windows Firewall Settings**

This section explains how to enable the port number used by Mirroring Controller, if the Windows firewall feature is enabled.

#### Windows Server(R) 2019:

- 1. In the [Windows Defender Firewall] window, click [Advanced settings] on the left side of the window.
- 2. In the [Windows Defender Firewall with Advanced Security] window, click [Inbound Rules] on the left side of the window.
- 3. Click [New Rule] on the right side of the window.
- 4. In the [New Inbound Rule Wizard] window, select [Port], and then click [Next].
- 5. Select [TCP] and [Specific local ports], then specify the port number defined in the network definition file, and then click [Next].
- 6. Select [Allow the connection], and then click [Next].
- 7. Select the profiles for which this rule applies, and then click [Next].
- 8. In [Name], specify the desired name, and then click [Finish].
- 9. In the [Windows Defender Firewall with Advanced Security] window, check if the added rule is enabled under [Inbound Rules] in the center of the window.

#### Other than above:

- 1. In the [Windows Firewall] window, click [Advanced settings] on the left side of the window.
- 2. In the [Windows Firewall with Advanced Security] window, click [Inbound Rules] on the left side of the window.
- 3. Click [New Rule] on the right side of the window.
- 4. In the [New Inbound Rule Wizard] window, select [Port], and then click [Next].
- 5. Select [TCP] and [Specific local ports], then specify the port number defined in the network definition file, and then click [Next].
- 6. Select [Allow the connection], and then click [Next].
- 7. Select the profiles for which this rule applies, and then click [Next].
- 8. In [Name], specify the desired name, and then click [Finish].
- 9. In the [Windows Firewall with Advanced Security] window, check if the added rule is enabled under [Inbound Rules] in the center of the window.

## Appendix F WebAdmin Disallow User Inputs Containing Hazardous Characters

WebAdmin considers the following as hazardous characters, which are not allowed in user inputs.

| (pipe sign)
& (ampersand sign)
; (semicolon sign)
\$ (dollar sign)
% (percent sign)
@ (at sign)
' (single apostrophe)
" (quotation mark)
\' (backslash-escaped apostrophe)
\" (backslash-escaped quotation mark)
<> (triangular parenthesis)
() (parenthesis)
+ (plus sign)
CR (Carriage return, ASCII 0x0d)
LF (Line feed, ASCII 0x0a)

, (comma sign)
\ (backslash)

## Appendix G Collecting Failure Investigation Data

If the cause of an error that occurs while building the environment or during operations is unclear, data must be collected for initial investigation.

This appendix describes how to collect data for initial investigation.

Use FJQSS (Information Collection Tool) to collect data for initial investigation.



Refer to the FJQSS manual for information on how to use FJQSS.



#### Using a database server to collect data

- When using FJQSS to collect data for initial investigation, you must set the following environment variables:
  - Environment variable required for collecting data from the database

Refer to "Collecting Failure Investigation Data" in the Operation Guide.

- MCCONTROLDIR
- Set the Mirroring Controller management directory.
   Refer to "1.2.1 Mirroring Controller Resources" for information on the Mirroring Controller management directory.

#### Using an arbitration server to collect data

#### Linux

- When using FJQSS to collect data for initial investigation, you must set the following environment variables:
  - ARBCONTROLDIR

Set the Mirroring Controller arbitration process management directory.

Refer to "1.2.1 Mirroring Controller Resources" for information on the Mirroring Controller arbitration server management directory.

- ARBUSER

Specify the operating system user who started the Mirroring Controller arbitration process.

- FSEP\_SA\_HOME

Set the Server Assistant installation directory.

#### Windows

- When using FJQSS to collect data for initial investigation, a window will be displayed for you to set the following environment variables:
  - ARBCONTROLDIR

Set the Mirroring Controller arbitration process management directory.

Refer to "1.2.1 Mirroring Controller Resources" for information on the Mirroring Controller arbitration server management directory.

## Index

[A]	[E]
Action Required when a Heartbeat Abnormality is Detected 66	Encryption of Transaction Logs Transferred to the Standby
Action Required when All Database Servers or Instances Stopped	Server1
94	
Action Required when an Error Occurs in the Database	[F]
Multiplexing Mode83	Failback of the Primary Server89
Action Required when Automatic Disconnection Fails93	Fencing Command
Action Required when Automatic Switch Fails93	[H]
Action Required when Server Degradation Occurs	How to Execute the mc_ctl Command130
Addressing Errors During Degrading Operation92	110w to Execute the Inc_cti Command
Application Connection Server Settings39	[1]
arbitration.conf file (Linux)	Identify cause of error and perform recovery86,92
arbitration.conf file (Windows)127	Identify Cause of Error and Restore the Standby Server85
Arbitration Command	Identify Cause of Error and Restructure the Standby Server 91
Arbitration Configuration File	If Performing the Referencing Job on the Synchronous Standby
Arbitration Server Maintenance	Server
Arbitration Server Process9	If Prioritizing the Main Job on the Primary Server
Arbitration Server Resources9	Installation1
Authentication of the Standby Server13	
	[M]
[B]	Manually Disconnecting the Standby Server65
Backing up Database Multiplexing Mode Information58	Manually Switching the Primary Server
Backup Data Storage Destination Directory	Matching the system times
Backup Operation	Mirroring Controller Resources
[C]	Monitoring Mirroring Controller Messages66
Changes in Operation	Monitoring Using Database Multiplexing Mode
Changes Required when the Standby Server is Stopped76	TAID.
Changing from Database Multiplexing Mode to Single Server	[N]
Mode	network.conf file
Changing from Single Server Mode to Database Multiplexing	Network Configuration File
Mode	Notes on Performing Automatic Degradation Immediately after a
Changing Parameters81	Heartbeat Abnormality
Changing to Database Multiplexing Mode when the Arbitration	Heartoeat Aonormanty
Server is Used for Automatic Degradation80	[0]
Checking the Behavior	Operations in Database Multiplexing Mode60
Checking the Connection Status37	Operations when the Server has Started Degrading after a
Checking the Connection Status on a Database Server37	Disconnection has Occurred90
Checking the Connection Status on the Arbitration Server38	Operations when the Server has Started Degrading after a Switch
Checking the Database Multiplexing Mode Status	has Occurred83
Checking the Status of the Arbitration Server64	Overview of Database Multiplexing Mode
Checking the Status of the Database Server63	
Configuring ICMP14	[P]
Configuring the Arbitration Server17	Parameters
Confirming the Streaming Replication Status36	Parameters Set on the Primary Server
Creating, Setting, and Registering the Primary Server Instance27	Parameters Set on the Standby Server112
Creating, Setting, and Registering the Standby Server Instance32	Post-attach Command
Creating Applications	Post-switch Command
Creating a User Exit for a Database Server35	postgresql.conf file
Creating a User Exit for the Arbitration Server22	Pre-detach Command
	Preparing for Setup
[D]	Preparing the Arbitration Server10
Database Backup Operation59	Preparing the Backup Disk10
Database Server Processes9	Preparing the Database Server10
Deciding on Operation when a Heartbeat Abnormality is	Preparing to Output Error Logs to the Event Log (Windows). 10

[R]
Rebuild the Standby Server
Recovering from an Incorrect User Operation98
Recovery of the Mirroring Controller management directory
86,92
Redundancy of the Admin and Log Transfer Networks10
Referencing on the Standby Server6
Rolling Updates
[S]
Security in Database Multiplexing11
Security Policy Settings
Security Policy Settings (Windows)
ServerConfiguration File
Server Configuration File for the Database Servers
server Configuration Pile for the Database Servers
Server Maintenance
Setting Automatic Start and Stop of Mirroring Controller and
Multiplexed Instances
Setting Automatic Start and Stop of the Mirroring Controller
Arbitration Process
Setting Up Database Multiplexing Mode
Setting Up Database Multiplexing Mode on the Primary Server23
Setting Up Database Multiplexing Mode on the Standby Server31
Setting Up the Arbitration Server
Setting Up the Primary Server
Setting Up the Standby Server
Setup
Starting and Stopping Mirroring Controller
Starting and Stopping the Mirroring Controller Arbitration
Process
Starting Mirroring Controller on the Primary Server30
Starting Mirroring Controller on the Standby Server34
Starting the Mirroring Controller Arbitration Process 23,60
State Transition Commands
Stop Mirroring Controller85,91
Stopping for Maintenance
Stopping the Mirroring Controller Arbitration Process60
Supplementary Information on Building the Primary Server and
Standby Server on the Same Server
Supplementary Procedure on Configuring for Operation in
Database Multiplexing Mode
System Configuration for Database Multiplexing Mode 6
[T]
Tuning for Optimization of Degrading Operation Using
Abnormality Monitoring
Tuning to Stabilize Queries on the Standby Server
Tuning to Stabilize Queries on the Standby Server (when
Performing Frequent Updates on the Primary Server)
Tuning to Stabilize the Database Multiplexing Mode39
[U]
Uninstalling in Database Multiplexing Mode
Users who perform setup and operations on the arbitration server $% \left( 1\right) =\left( 1\right) \left( 1\right) \left($
Users who perform setup and operations on the database server14

[W]
What is Database Multiplexing Mode
Windows Firewall Settings
<u> </u>



# FUJITSU Enterprise Postgres 12



# Cluster Operation Guide Database Multiplexing (Windows)

## **Preface**

#### Purpose of this document

This document describes the tasks required for using the database multiplexing feature of FUJITSU Enterprise Postgres.

#### Intended readers

This document is intended for those who set up and use the database multiplexing feature.

Readers of this document are also assumed to have general knowledge of:

- PostgreSQL
- SQL
- Windows

#### Structure of this document

This document is structured as follows:

#### Chapter 1 Overview of Database Multiplexing Mode

Provides an overview of database multiplexing mode.

#### Chapter 2 Setting Up Database Multiplexing Mode

Describes how to set up database multiplexing mode.

#### Chapter 3 Operations in Database Multiplexing Mode

Explains periodic database multiplexing mode.

#### Chapter 4 Action Required when an Error Occurs in Database Multiplexing Mode

Explains the action required when an error occurs during a database multiplexing mode.

#### Chapter 5 Managing Mirroring Controller Using WebAdmin

Explains how to set up and manage Mirroring Controller in a streaming replication cluster using WebAdmin.

#### Appendix A Parameters

Explains the configuration files and parameters required for database multiplexing mode.

#### Appendix B Supplementary Information on Building the Primary Server and Standby Server on the Same Server

Explains supplementary information on building the primary server and standby server on the same server.

#### Appendix C User Commands

Explains the user commands.

#### Appendix D Notes on Performing Automatic Degradation Immediately after a Heartbeat Abnormality

Provides notes when performing automatic degradation unconditionally after a heartbeat abnormality is detected during heartbeat monitoring of an operating system or server.

#### Appendix E Supplementary Procedure on Configuring for Operation in Database Multiplexing Mode

Explains supplementary procedure on the configuration required for operation in database multiplexing mode.

#### Appendix F WebAdmin Disallow User Inputs Containing Hazardous Characters

Explains characters not allowed in WebAdmin.

#### Appendix G Collecting Failure Investigation Data

Explains how to collect data for initial investigation.

#### **Export restrictions**

Exportation/release of this document may require necessary procedures in accordance with the regulations of your resident country and/or US export control laws.

#### Issue date and version

Edition 1.0: July 2020

#### Copyright

Copyright 2016-2020 FUJITSU LIMITED

## Contents

Chapter 1 Overview of Database Multiplexing Mode	
1.1 What is Database Multiplexing Mode	
1.1.1 Monitoring Using Database Multiplexing Mode	
1.1.2 Referencing on the Standby Server	
1.1.2.1 If Prioritizing the Main Job on the Primary Server	
1.1.2.2 If Performing the Referencing Job on the Synchronous Standby Server	(
1.2 System Configuration for Database Multiplexing Mode	
1.2.1 Mirroring Controller Resources.	8
1.2.1.1 Database Server Resources.	8
1.2.1.2 Arbitration Server Resources	9
1.2.2 Mirroring Controller Processes	9
1.2.2.1 Database Server Processes	9
1.2.2.2 Arbitration Server Process	9
1.2.3 Redundancy of the Admin and Log Transfer Networks	10
1.2.4 Notes on CPU Architecture and Products	10
1.3 Deciding on Operation when a Heartbeat Abnormality is Detected	10
1.4 Security in Database Multiplexing	
1.4.1 Authentication of the Standby Server	
1.4.2 Encryption of Transaction Logs Transferred to the Standby Server	14
Chapter 2 Setting Up Database Multiplexing Mode	15
2.1 Installation.	16
2.2 Preparing for Setup	17
2.2.1 Preparing the Database Server	17
2.2.1.1 Preparing the Backup Disk	17
2.2.1.2 Preparatory Tasks for the Output of Error Logs to the Event Log	17
2.2.1.3 Security Policy Settings	18
2.2.2 Preparing the Arbitration Server	18
2.2.2.1 Preparing to Output Error Logs to the Event Log (Windows)	18
2.2.2.2 Security Policy Settings (Windows)	19
2.3 Setting Up the Arbitration Server	19
2.3.1 Configuring the Arbitration Server	19
2.3.2 Creating a User Exit for the Arbitration Server.	24
2.3.3 Starting the Mirroring Controller Arbitration Process.	24
2.4 Setting Up the Primary Server	25
2.4.1 Setting Up Database Multiplexing Mode on the Primary Server	
2.4.2 Creating, Setting, and Registering the Primary Server Instance	
2.4.3 Starting Mirroring Controller on the Primary Server.	
2.5 Setting Up the Standby Server	
2.5.1 Setting Up Database Multiplexing Mode on the Standby Server	
2.5.2 Creating, Setting, and Registering the Standby Server Instance	
2.5.3 Starting Mirroring Controller on the Standby Server	
2.6 Creating a User Exit for a Database Server.	
2.7 Confirming the Streaming Replication Status	
2.8 Checking the Connection Status.	
2.8.1 Checking the Connection Status on a Database Server.	
2.8.2 Checking the Connection Status on the Arbitration Server.	
2.9 Creating Applications.	
2.9.1 Application Connection Server Settings.	
2.10 Checking the Behavior	
2.11 Tuning.	
2.11.1 Tuning to Stabilize the Database Multiplexing Mode	
2.11.2 Tuning to Stabilize Queries on the Standby Server	
2.11.2 Tuning to Stabilize Queries on the Standby Server (when Performing Frequent Updates on the Primary Server)	
2.11.4 Tuning for Optimization of Degradation Using Abnormality Monitoring	

2.11.4.1.1 Tuning Abnormality Monitoring for Operations that Use an Arbitration Server for Automatic Degeneration.  2.11.4.1.2 Tuning Abnormality Monitoring for Operations that Perform Automatic Degeneration by Calling a User Exist bat December 1. December 1	2.11.4.1 Tuning for Abnormality Monitoring of the Operating System or Server	
Determines Degeneration		
2.11.4.1.3 Tuning Abnormality Monitoring for Operations that Notify Messages. 2.11.4.1 Tuning Abnormality Monitoring of Operations that Perform Automatic Degenerate Unconditionally due to Heartbean Abnormality Monitoring of Darabase Processes. 3.1.1.4.2 Tuning for Abnormality Monitoring of Streaming Replication. 3.5.2.11.4.4 Tuning for Disk Abnormality Monitoring of Streaming Replication. 3.5.2.11.4.4 Tuning for Disk Abnormality Monitoring. 3.5.2.11.4.4 Tuning for Disk Abnormality Monitoring. 3.5.2.1.4.4 Extensional Start and Stop of Mirroring Controller and Multiplexed Instances. 3.5.2.1.3 Setting Automatic Start and Stop of Mirroring Controller Arbitration Process. 3.1.4 Backup Operation. 3.2.1.4 Backing up Database Multiplexing Mode Information. 3.2.1.4.1 Backing up Database Multiplexing Mode. 3.1.3 Surring and Stopping the Mirroring Controller Arbitration Process. 3.1.3 Surring and Stopping the Mirroring Controller Arbitration Process. 4.6.3.1.3 Surring the Mirroring Controller Arbitration Process. 4.6.3.2 Surring and Stopping Mirroring Controller. 4.3.2 Surring and Stopping Mirroring Controller. 4.3.3 Checking the Database Multiplexing Mode Status. 3.3 Checking the Database Multiplexing Mode Status. 3.3 Checking the Status of the Arbitration Process. 4.6.3 Surring and Stopping Mirroring Controller. 4.6.3 Surring Start Status of the Arbitration Server. 4.6.3 Surring Mirroring Controller. 4.7.3 Surring Mirroring Controller Startenses. 4.7.4 Surring Mirroring Controller Startenses. 4.8 Surring Mirroring Controller Startenses. 4.9 Changes in Operation. 4.9 Changes Required when an Earter Occurs in Database Multiplexing Mode. 4.1 Changes in Operation. 4.1 Surring Parameters. 4.1 Surring Parameters. 4.1 Surring Parameters. 4.1 Surring Parameters. 4.1		
2.11.4.1.4 Tuning Abnormality Monitoring for Operations that Perform Automatic Degenerate Unconditionally due to Hearbeat Abnormality Monitoring of Darabase Processes		
Ahnormality		
2.11.4.2 Tuning for Abnormality Monitoring of Darabase Processes. 2.11.4.3 Tuning for Abnormality Monitoring of Streaming Replication		
2.11.4.3 Tuning for Disk Abnormality Monitoring of Streaming Replication.  5.57 2.1.1.4.4 Tuning for Disk Abnormality Monitoring.  5.72 2.1.2 Setting Automatic Start and Stop of Mirroring Controller and Multiplexed Instances.  5.88 2.1.3 Setting Automatic Start and Stop of the Mirroring Controller Arbitration Process.  5.92 2.1.4 Backing up Database Multiplexing Mode Information.  6.2 2.1.4.1 Backing up Database Multiplexing Mode.  6.2 2.1.4.1 Backing up Database Multiplexing Mode.  6.3 3.1 Starting and Stopping the Mirroring Controller Arbitration Process.  6.4 3.1.1 Storping the Mirroring Controller Arbitration Process.  6.4 3.1.2 Stopping the Mirroring Controller Arbitration Process.  6.4 3.2.3 Starting and Stopping Mirroring Controller.  6.5 3.2.1 Starting Mirroring Controller.  6.5 3.2.1 Starting Mirroring Controller.  6.5 3.2.2 Stopping Mirroring Controller.  6.7 3.3 Checking the Database Multiplexing Mode Status.  6.8 3.3.1 Checking the Status of the Database Server  6.8 3.3.2 Checking the Status of the Database Server  6.8 3.3.4 Amaulty Disconnecting the Standby Server.  6.9 3.5 Manually Disconnecting the Standby Server.  6.9 3.8 Server Maintenance.  7.1 3.8 Fortion Required when a Heartbeat Abnormality is Detected.  7.0 3.9.1 Changes Required when the Standby Server.  7.1 3.8.2 Stopping for Maintenance.  7.2 3.3.3 Changing From Database Multiplexing Mode value in Single Server Mode.  8.3 3.9.2 Changing from Database Multiplexing Mode when the Arbitration Server is Stopped.  8.0 3.9.1 Changes Required when the Standby Server is Stopped.  8.0 3.9.2 Changing from Database Multiplexing Mode on Single Server Mode.  8.1 3.9.3 Changes in Doreation.  8.1 3.9.4 Changing parameters.  8.2 3.9.5 Changing parameters.  8.3 3.9.5 Changing parameters.  8.6 3.9.6 Changing parameters.  8.7 3.9.6 Changing parameters.  8.8 4.1 Action Required when the Standby Server is Stopped.  8.8 4.1 Action Required when the Standby Server is Stopped.  8.8 4.1 Action Required when the Standby Server is Stopped.  8.8 4.1 Action Reguired	·	
2.11.4.4 Tuning for Disk Abnormality Monitoring		
2.13 Setting Automatic Start and Stop of the Mirroring Controller and Multiplexed Instances.  2.14 Backing up Database Multiplexing Mode Information.  2.2 1.14 Backing up Database Multiplexing Mode Information.  2.2 2.14.1 Backing up Database Multiplexing Mode Information.  2.2 2.14.2 Database Rackup Operation  2.2 2.14.2 Database Multiplexing Mode.  2.14.2 Database Backup Operation  2.2 2.14.2 Database Multiplexing Mode.  3.1 Starting and Stopping the Mirroring Controller Arbitration Process.  3.1 Starting the Mirroring Controller Arbitration Process.  3.2 Starting and Stopping the Mirroring Controller Arbitration Process.  3.2 Starting and Stopping Mirroring Controller.  3.2 Stopping Mirroring Controller.  3.3 Checking the Database Multiplexing Mode Status.  3.3 Checking the Database Multiplexing Mode Status.  3.3 Checking the Status of the Database Server.  3.3 Checking the Status of the Database Server.  3.4 Manually Disconnecting the Standsy Server.  3.5 Manually Disconnecting the Standsy Server.  3.6 Action Required when a Heartbeat Abnormality is Detected.  3.7 Monitoring Mirroring Controller Messages.  3.8.1 Stopping For Maintenance.  3.8.2 Stopping for Maintenance.  3.8.3 As Arbitration Server Maintenance.  3.9 Changes in Operation  3.9.1 Changes Required when the Standby Server is Stopped.  80  3.9.2 Changing from Single Server Mode to Single Server Mode.  3.9.3 Changing from Database Multiplexing Mode.  81  4.1 Action Required when Rever Degradation Occurs.  82  4.1 Action Required when Server Degradation Occurs.  84  4.1 Action Required when Server Degradation Occurs.  85  4.1.1 Rebuild the Standby Server.  96  4.1.1.1 Rebuild the Standby Server.  97  4.1.2 Rebuild the Standby Server.  97  4.1.1.2 Rebuild the Standby Server is Stopped.  88  4.1.1 Operations when the Server has Started Degrading after a Switch has Occurred.  88  4.1.1.1 Rebuild the Standby Server.  97  4.1.2.1 Rebuild the Standby Server has Started Degrading after a Switch has Occurred.  4.1.2.1.2 Recovery of the Mirroring Controller		
2.13 Setting Automatic Start and Stop of the Mirroring Controller Arbitration Process. 2.14 Backing up Database Multiplexing Mode Information. 2.2.14.12 Database Backup Operation. 2.3.13 Chapter 3 Operations in Database Multiplexing Mode. 3.15 Starting and Stopping the Mirroring Controller Arbitration Process. 46.3.1.1 Starting the Mirroring Controller Arbitration Process. 46.3.1.2 Stopping the Mirroring Controller Arbitration Process. 46.3.1.2 Stopping the Mirroring Controller Arbitration Process. 46.3.2.2 Starting and Stopping Mirroring Controller Arbitration Process. 46.3.2.2 Starting and Stopping Mirroring Controller. 46.3.2.2 Starting and Stopping Mirroring Controller. 46.3.2.3 Starting and Stopping Mirroring Controller. 46.3.3.2 Stopping Mirroring Controller. 46.3.3.2 Stopping Mirroring Controller. 46.3.3.2 Stopping Mirroring Controller. 46.3.3.3 Checking the Database Multiplexing Mode Status. 46.3.3.1 Checking the Status of the Arbitration Server. 46.3.3.2 Checking the Status of the Arbitration Server. 46.3.3 Manually Disconnecting the Standby Server. 47.3 Manually Switching the Primary Server. 48.4 Manually Switching the Primary Server. 49.3 Manually Disconnecting the Standby Server. 40.4 Annually Switching the Primary Server. 40.3 Manually Disconnecting the Standby Server. 40.4 Notice Required when a Heartbeat Abnormality is Detected. 40.3 Miniterian Required when a Heartbeat Abnormality is Detected. 40.3 Miniterian Mirroring Controller Messages. 41.3 Server Maintenance. 47.3 Server Maintenance. 47.3 Server Maintenance. 47.3 Server Maintenance. 48.3 Server Maintenance. 48.3 Server Maintenance. 49.3 Changes in Operation. 40.3 Permaneters Multiplexing Mode to Single Server Mode. 40.3 Server Maintenance. 40.4 Server Maintenance. 41.4 Action Required when the Standby Server is Stopped. 41.4 Action Required when an Error Occurs in Database Multiplexing Mode. 41.4 Action Required when the Server has Started Degrading after a Switch has Occurred. 41.4.1.1 Parations when the Server has Started Degrading after		
2.14.1 Backing up Database Multiplexing Mode Information		
2.14.1 Backing up Database Multiplexing Mode Information. 62 2.14.2 Database Backup Operation. 62 2.14.2 Database Backup Operation. 62 2.14.2 Database Backup Operation. 62 3.1. Starting and Stopping the Mirroring Controller Arbitration Process. 64 3.1.1 Starting the Mirroring Controller Arbitration Process. 64 3.1.2 Stopping the Mirroring Controller Arbitration Process. 64 3.2 Starting and Stopping Mirroring Controller. 65 3.2.1 Starting Mirroring Controller. 65 3.2.1 Starting Mirroring Controller. 65 3.2.2 Stopping Mirroring Controller. 67 3.3 Checking the Database Multiplexing Mode Status. 68 3.3.1 Checking the Status of the Database Server. 68 3.3.1 Checking the Status of the Arbitration Server. 69 3.4 Manually Disconnecting the Standsy Server. 69 3.5 Manually Disconnecting the Standsy Server. 70 3.6 Action Required when a Heartbeat Abnormality is Detected. 70 3.7 Monitoring Mirroring Controller Messages. 71 3.8 Server Maintenance. 73 3.8 Server Maintenance. 73 3.9.1 Changes in Operation 80 3.9.2 Changing from Maintenance. 73 3.9.1 Changes Required when the Standby Server is Stopped. 80 3.9.2 Changing from Single Server Mode to Database Multiplexing Mode 81 3.9.3 Changing from Database Multiplexing Mode when the Arbitration Server is Used for Automatic Degradation 83 3.9.4 Changing from Database Multiplexing Mode when the Arbitration Server is Used for Automatic Degradation 83 3.9.4 Changing from Database Multiplexing Mode when the Arbitration Server is Used for Automatic Degradation 84 4.1 Action Required when a Error Occurs in Database Multiplexing Mode. 81 4.1 Action Required when Server Degradation Occurs. 84 4.1.1 Operations when the Server has Started Degrading after a Switch has Occurred. 88 4.1.1.1 Identify Cause of Error and Restore the Standby Server. 90 4.1.1.1.2 Recovery of the Mirroring Controller management directory. 91 4.1.1.2 Recovery of the Mirroring Controller management directory. 91 4.1.2.1 Identify Cause of Error and Restore the Standby Server. 94 4.1.2.1 Identify Cause of Error and Res		
Chapter 3 Operations in Database Multiplexing Mode	• •	
Chapter 3 Operations in Database Multiplexing Mode.  3.1 Starting and Stopping the Mirroring Controller Arbitration Process.  4.3.1.1 Starting the Mirroring Controller Arbitration Process.  4.3.1.2 Stopping the Mirroring Controller Arbitration Process.  4.3.1.2 Stopping the Mirroring Controller Arbitration Process.  4.3.2 Starting and Stopping Mirroring Controller.  5.3.2.1 Starting Mirroring Controller.  5.3.2.2 Stopping Mirroring Controller.  5.3.2.2 Stopping Mirroring Controller.  5.3.3.3 Checking the Database Multiplexing Mode Status.  3.3.1 Checking the Status of the Database Server.  5.3.3.2 Stocking the Status of the Arbitration Server.  5.3.3 Admually Switching the Primary Server.  5.3 Admually Switching the Primary Server.  5.3 Amually Disconnecting the Standby Server.  5.3 Anonitoring Mirroring Controller Messages.  7.3 Monitoring Mirroring Controller Messages.  7.3 As Server Maintenance.  7.3 S. Server Maintenance.  7.3 S. Stopping for Database Multiplexing Mode.  8.0 S. Stopping for Database Multiplexing Mode to Single Server Mode.  8.1 S. Stopping for Database Multiplexing Mode when the Arbitration Server is Used for Automatic Degradation.  8.2 Stopping parameters.  8.3 S. Othanging Form Single Server Mode to Database Multiplexing Mode.  8.4 L1.1 Identify Cause of Error and Restore the Standby Server.  9.4 L1.1.1 Secovery of the Mirroring Controller management directory.  9.4 L1.1.1 Rebuild the Standby Server.  9.5 L1.1.1 Identify Cause of Error and Restore the Standby Server.  9.6 L1.1.1 Rebuild the Standby Server.  9.7 L1.2.1 Rebuild the Standby Server.  9.8 L1.2.1 Identify Cause of Error and Restore the Standby Server.  9.9 L1.2.1 R		
3.1.1 Starting the Mirroring Controller Arbitration Process. 64 3.1.1 Starting the Mirroring Controller Arbitration Process. 64 3.1.2 Stopping the Mirroring Controller Arbitration Process. 64 3.2 Starting and Stopping Mirroring Controller. 65 3.2.1 Starting Mirroring Controller. 65 3.2.2 Stopping Mirroring Controller. 65 3.2.2 Stopping Mirroring Controller. 67 3.3.1 Checking the Database Multiplexing Mode Status. 68 3.3.1 Checking the Database Server. 68 3.3.2 Checking the Status of the Arbitration Server. 68 3.3.3 Checking the Status of the Arbitration Server. 69 3.4 Manually Switching the Primary Server. 69 3.5 Manually Disconnecting the Standby Server. 70 3.6 Action Required when a Heartheat Abnormality is Detected. 70 3.7 Monitoring Mirroring Controller Messages. 71 3.8 Server Maintenance. 73 3.8 Server Maintenance. 73 3.8.1 Rolling Updates. 73 3.8.2 Stopping for Maintenance. 73 3.9.1 Changes Required when the Standby Server is Stopped. 80 3.9.1 Changes Required when the Standby Server is Stopped. 80 3.9.2 Changing from Single Server Mode to Database Multiplexing Mode 81 3.9.3 Changing from Database Multiplexing Mode when the Arbitration Server Mode. 82 3.9.4 Changing Form Database Multiplexing Mode when the Arbitration Server is Used for Automatic Degradation. 83 3.9.5 Changing Parameters. 86 3.9.6 Uninstalling in Database Multiplexing Mode. 86 Chapter 4 Action Required when an Error Occurs in Database Multiplexing Mode. 86 Chapter 4 Action Required when Server has Started Degrading after a Switch has Occurred. 88 4.1.1.1 Identify Cause of Error and Restore the Standby Server. 99 4.1.1.1.2 Recovery of the Mirroring Controller management directory. 91 4.1.1.1 Reposition of the Mirroring Controller management directory. 91 4.1.1.2 Reposition of the Mirroring Controller management directory. 91 4.1.2.1 Identify Cause of Error and Restore the Standby Server. 97 4.1.2.1 Reposition of the Mirroring Controller management directory. 97 4.1.2.1 Reposition of the Mirroring Controller management directory. 97 4.1.2.1	2.14.2 Database Backup Operation	62
3.1.1 Starting the Mirroring Controller Arbitration Process. 64 3.1.1 Starting the Mirroring Controller Arbitration Process. 64 3.1.2 Stopping the Mirroring Controller Arbitration Process. 64 3.2 Starting and Stopping Mirroring Controller. 65 3.2.1 Starting Mirroring Controller. 65 3.2.2 Stopping Mirroring Controller. 65 3.2.2 Stopping Mirroring Controller. 67 3.3.1 Checking the Database Multiplexing Mode Status. 68 3.3.1 Checking the Database Server. 68 3.3.2 Checking the Status of the Arbitration Server. 68 3.3.3 Checking the Status of the Arbitration Server. 69 3.4 Manually Switching the Primary Server. 69 3.5 Manually Disconnecting the Standby Server. 70 3.6 Action Required when a Heartheat Abnormality is Detected. 70 3.7 Monitoring Mirroring Controller Messages. 71 3.8 Server Maintenance. 73 3.8 Server Maintenance. 73 3.8.1 Rolling Updates. 73 3.8.2 Stopping for Maintenance. 73 3.9.1 Changes Required when the Standby Server is Stopped. 80 3.9.1 Changes Required when the Standby Server is Stopped. 80 3.9.2 Changing from Single Server Mode to Database Multiplexing Mode 81 3.9.3 Changing from Database Multiplexing Mode when the Arbitration Server Mode. 82 3.9.4 Changing Form Database Multiplexing Mode when the Arbitration Server is Used for Automatic Degradation. 83 3.9.5 Changing Parameters. 86 3.9.6 Uninstalling in Database Multiplexing Mode. 86 Chapter 4 Action Required when an Error Occurs in Database Multiplexing Mode. 86 Chapter 4 Action Required when Server has Started Degrading after a Switch has Occurred. 88 4.1.1.1 Identify Cause of Error and Restore the Standby Server. 99 4.1.1.1.2 Recovery of the Mirroring Controller management directory. 91 4.1.1.1 Reposition of the Mirroring Controller management directory. 91 4.1.1.2 Reposition of the Mirroring Controller management directory. 91 4.1.2.1 Identify Cause of Error and Restore the Standby Server. 97 4.1.2.1 Reposition of the Mirroring Controller management directory. 97 4.1.2.1 Reposition of the Mirroring Controller management directory. 97 4.1.2.1	Chapter 3 Operations in Database Multiplexing Mode	64
3.1.1 Starting the Mirroring Controller Arbitration Process. 3.1.2 Stropping the Mirroring Controller Arbitration Process. 3.2 Starting and Stopping Mirroring Controller. 55 3.2.1 Starting Mirroring Controller. 55 3.2.2 Stropping Mirroring Controller. 57 3.2.2 Stropping Mirroring Controller. 57 3.3. Checking the Database Multiplexing Mode Status. 58 3.3.1 Checking the Status of the Database Server. 58 3.3.2 Checking the Status of the Database Server. 59 3.4 Manually Switching the Primary Server. 50 3.5 Manually Disconnecting the Standby Server. 50 3.5 Admually Disconnecting the Standby Server. 50 3.6 Action Required when a Hearrbeat Abnormality is Detected. 50 3.7 Monitoring Mirroring Controller Messages. 51 3.8 Server Maintenance. 51 3.8 Server Maintenance. 52 3.8 Server Maintenance. 53 3.8 Server Maintenance. 53 3.9 Changes in Operation. 53 3.9 Changes in Operation. 53 3.9 Changes Required when the Standby Server is Stopped. 53 3.9 Changing from Single Server Mode to Database Multiplexing Mode 53 3.9 Changing from Database Multiplexing Mode to Single Server Mode. 53 3.9 Changing Parameters. 53 3.9 Changing Parameters. 53 50 50 50 50 50 50 60 60 60 60 60 60 60 60 60 60 60 60 60		
3.1.2 Stopping the Mirroring Controller Arbitration Process.  3.2 Starting and Stopping Mirroring Controller.  53.2.1 Starting Mirroring Controller.  53.2.1 Starting Mirroring Controller.  53.2.2 Stopping Mirroring Controller.  53.3.2 Stopping Mirroring Controller.  53.3.3.1 Checking the Database Multiplexing Mode Status.  53.3.3.1 Checking the Status of the Arbitration Server.  58.3.3.2 Checking the Status of the Arbitration Server.  59.3.4 Manually Switching the Primary Server.  50.3.5 Manually Disconnecting the Standby Server.  50.3.6 Action Required when a Heartbeat Abnormality is Detected.  50.3.7 Monitoring Mirroring Controller Messages.  71.3.8 Server Maintenance.  73.3.8 Server Maintenance.  73.3.1 Rolling Updates.  73.3.3.2 Stopping for Maintenance.  73.3.3.3 Arbitration Server Maintenance.  73.3.3 Changes in Operation.  80.3.9.1 Changes Required when the Standby Server is Stopped.  80.3.9.2 Changing from Single Server Mode to Database Multiplexing Mode.  81.3.9.3 Changing from Database Multiplexing Mode when the Arbitration Server is Used for Automatic Degradation.  82.3.9.4 Changing to Database Multiplexing Mode when the Arbitration Server is Used for Automatic Degradation.  83.9.6 Changing from Database Multiplexing Mode when the Arbitration Server is Used for Automatic Degradation.  85.3.9.6 Uninstalling in Database Multiplexing Mode.  86.6 Chapter 4 Action Required when an Error Occurs in Database Multiplexing Mode.  87.4 Al.1.1 Identify Cause of Error and Restore the Standby Server.  88.4 Al.1.1.1 Stop Mirroring Controller.  89.4 Al.1.1.1 Stop Mirroring Controller management directory.  91.4 Al.1.2 Recovery of the Mirroring Controller management directory.  91.4 Al.1.2 Recovery of the Mirroring Controller management directory.  91.4 Al.2.1 Identify Cause of Error and Restore the Standby Server.  93.4 Al.2.1 Identify Cause of Error and Perform recovery.  94.4 Al.2.1 Identify Cause of Error and Perform recovery.  94.4 Al.2.1.3 Identify cause of Error and Perform recovery.		
3.2 Starting and Stopping Mirroring Controller		
3.2.1 Starting Mirroring Controller		
3.2 2 Stopping Mirroring Controller. 67 3.3 Checking the Database Multiplexing Mode Status. 68 3.3.1 Checking the Status of the Database Server. 68 3.3.2 Checking the Status of the Database Server. 68 3.3.2 Checking the Status of the Arbitration Server. 69 3.4 Manually Switching the Primary Server. 70 3.5 Manually Disconnecting the Standby Server. 70 3.6 Action Required when a Heartbeat Abnormality is Detected. 70 3.7 Monitoring Mirroring Controller Messages. 71 3.8 Server Maintenance. 73 3.8.1 Rolling Updates. 73 3.8.1 Rolling Updates. 73 3.8.2 Stopping for Maintenance. 73 3.8.3 Arbitration Server Maintenance. 73 3.9 Changes in Operation 79 3.9 Changes in Operation 79 3.9 Changes Required when the Standby Server is Stopped. 79 3.9.1 Changes Required when the Standby Server is Stopped. 80 3.9.2 Changing from Single Server Mode to Database Multiplexing Mode. 81 3.9.3 Changing to Database Multiplexing Mode when the Arbitration Server is Used for Automatic Degradation. 85 3.9.5 Changing Parameters. 85 3.9.6 Uninstalling in Database Multiplexing Mode. 86 Chapter 4 Action Required when an Error Occurs in Database Multiplexing Mode. 86 Chapter 4 Action Required when an Error Occurs in Database Multiplexing Mode. 86 4.1 Action Required when Server Degradation Occurs. 88 4.1.1 Operations when the Server has Started Degrading after a Switch has Occurred. 88 4.1.1.1 Identify Cause of Error and Restore the Standby Server. 90 4.1.1.1.2 Rebuild the Standby Server. 91 4.1.1.3 Reiback of the Primary Server. 91 4.1.1.4 Repositions when the Server has Started Degrading after a Disconnection has Occurred. 95 4.1.2.1 Bentify cause of Error and perform recovery. 91 4.1.2 Republid the Standby Server. 93 4.1.2 Repositions when the Server has Started Degrading after a Disconnection has Occurred. 95 4.1.2.1 Bentify cause of Error and perform recovery. 91 4.1.2 Republid the Standby Server. 95 4.1.2.1 Identify Cause of Error and perform recovery. 97 4.1.2.2 Rebuild the Standby Server. 97 4.1.2.1 Republic the Standby Server. 97 4.1.2.2 R		
3.3.1 Checking the Status of the Database Server. 68 3.3.2 Checking the Status of the Arbitration Server. 69 3.4 Manually Switching the Primary Server. 69 3.5 Manually Disconnecting the Standby Server. 70 3.6 Action Required when a Heartbeat Abnormality is Detected. 70 3.7 Monitoring Mirroring Controller Messages. 71 3.8 Server Maintenance. 73 3.8.1 Rolling Updates. 73 3.8.2 Stopping for Maintenance 78 3.8.3 Stopping for Maintenance 78 3.9 Changes in Operation 80 3.9.1 Changes Required when the Standby Server is Stopped 80 3.9.2 Changing from Single Server Mode to Database Multiplexing Mode. 81 3.9.3 Changing from Database Multiplexing Mode to Single Server Mode 83 3.9.4 Changing from Database Multiplexing Mode to Single Server is Used for Automatic Degradation 85 3.9.5 Changing Parameters. 86 3.9.6 Uninstalling in Database Multiplexing Mode. 82 3.9.4 Changing from Single Server Mode 86 Chapter 4 Action Required when an Error Occurs in Database Multiplexing Mode. 86 Chapter 4 Action Required when an Error Occurs in Database Multiplexing Mode. 88 4.1.1 Operations when the Server has Started Degrading after a Switch has Occurred. 88 4.1.1.1 Identify Cause of Error and Restore the Standby Server 90 4.1.1.1.1 Stop Mirroring Controller. 90 4.1.1.1 Recovery of the Mirroring Controller management directory 91 4.1.1.2 Rebuild the Standby Server. 93 4.1.2 Operations when the Server has Started Degrading after a Disconnection has Occurred. 95 4.1.2.1 Identify cause of Error and Restore the Standby Server. 94 4.1.2.1 Identify cause of Error and Restore the Standby Server. 94 4.1.2.1 Republication of the Primary Server. 94 4.1.2.2 Republication when the Server has Started Degrading after a Disconnection has Occurred. 95 4.1.2.1 Republication of the Primary Server. 96 4.1.2.2 Republication when the Server has Started Degrading after a Disconnection has Occurred. 97 4.1.2.2 Republication when the Server has Started Degrading after a Disconnection has Occurred. 97 4.1.2.2 Republication when the Server has Started Degrading af		
3.3.1 Checking the Status of the Database Server. 68 3.3.2 Checking the Status of the Arbitration Server. 69 3.4 Manually Switching the Primary Server. 69 3.5 Manually Disconnecting the Standby Server. 70 3.6 Action Required when a Heartbeat Abnormality is Detected. 70 3.7 Monitoring Mirroring Controller Messages. 71 3.8 Server Maintenance. 73 3.8.1 Rolling Updates. 73 3.8.2 Stopping for Maintenance 78 3.8.3 Stopping for Maintenance 78 3.9 Changes in Operation 80 3.9.1 Changes Required when the Standby Server is Stopped 80 3.9.2 Changing from Single Server Mode to Database Multiplexing Mode. 81 3.9.3 Changing from Database Multiplexing Mode to Single Server Mode 83 3.9.4 Changing from Database Multiplexing Mode to Single Server is Used for Automatic Degradation 85 3.9.5 Changing Parameters. 86 3.9.6 Uninstalling in Database Multiplexing Mode. 82 3.9.4 Changing from Single Server Mode 86 Chapter 4 Action Required when an Error Occurs in Database Multiplexing Mode. 86 Chapter 4 Action Required when an Error Occurs in Database Multiplexing Mode. 88 4.1.1 Operations when the Server has Started Degrading after a Switch has Occurred. 88 4.1.1.1 Identify Cause of Error and Restore the Standby Server 90 4.1.1.1.1 Stop Mirroring Controller. 90 4.1.1.1 Recovery of the Mirroring Controller management directory 91 4.1.1.2 Rebuild the Standby Server. 93 4.1.2 Operations when the Server has Started Degrading after a Disconnection has Occurred. 95 4.1.2.1 Identify cause of Error and Restore the Standby Server. 94 4.1.2.1 Identify cause of Error and Restore the Standby Server. 94 4.1.2.1 Republication of the Primary Server. 94 4.1.2.2 Republication when the Server has Started Degrading after a Disconnection has Occurred. 95 4.1.2.1 Republication of the Primary Server. 96 4.1.2.2 Republication when the Server has Started Degrading after a Disconnection has Occurred. 97 4.1.2.2 Republication when the Server has Started Degrading after a Disconnection has Occurred. 97 4.1.2.2 Republication when the Server has Started Degrading af		
3.3.2 Checking the Status of the Arbitration Server		
3.5 Manually Disconnecting the Standby Server		
3.5 Manually Disconnecting the Standby Server		
3.7 Monitoring Mirroring Controller Messages		
3.8 Server Maintenance	3.6 Action Required when a Heartbeat Abnormality is Detected	70
3.8.1 Rolling Updates	3.7 Monitoring Mirroring Controller Messages	71
3.8.2 Stopping for Maintenance	3.8 Server Maintenance	73
3.8.3 Arbitration Server Maintenance	3.8.1 Rolling Updates	73
3.9 Changes in Operation		
3.9.1 Changes Required when the Standby Server is Stopped		
3.9.2 Changing from Single Server Mode to Database Multiplexing Mode	3.9 Changes in Operation	80
3.9.3 Changing from Database Multiplexing Mode to Single Server Mode	3.9.1 Changes Required when the Standby Server is Stopped	80
3.9.4 Changing to Database Multiplexing Mode when the Arbitration Server is Used for Automatic Degradation		
3.9.5 Changing Parameters	3.9.3 Changing from Database Multiplexing Mode to Single Server Mode	82
3.9.6 Uninstalling in Database Multiplexing Mode	3.9.4 Changing to Database Multiplexing Mode when the Arbitration Server is Used for Automatic Degradation	85
Chapter 4 Action Required when an Error Occurs in Database Multiplexing Mode		
4.1 Action Required when Server Degradation Occurs	3.9.6 Uninstalling in Database Multiplexing Mode	86
4.1 Action Required when Server Degradation Occurs.  4.1.1 Operations when the Server has Started Degrading after a Switch has Occurred.  4.1.1.1 Identify Cause of Error and Restore the Standby Server.  4.1.1.1 Stop Mirroring Controller.  90 4.1.1.1.2 Recovery of the Mirroring Controller management directory.  91 4.1.1.3 Identify cause of error and perform recovery.  91 4.1.1.2 Rebuild the Standby Server.  93 4.1.1.3 Failback of the Primary Server.  94 4.1.2 Operations when the Server has Started Degrading after a Disconnection has Occurred.  95 4.1.2.1 Identify Cause of Error and Restore the Standby Server.  96 4.1.2.1.2 Recovery of the Mirroring Controller management directory.  97 4.1.2.1.3 Identify cause of error and perform recovery.  97 4.1.2.1.3 Identify cause of error and perform recovery.  97 4.1.2.1 Recovery of the Mirroring Controller management directory.  97 4.1.2.1 Identify cause of error and perform recovery.  97 4.1.2.2 Rebuild the Standby Server.  97	Chanter 4 Action Required when an Error Occurs in Database Multiplexing Mode	88
4.1.1 Operations when the Server has Started Degrading after a Switch has Occurred		
4.1.1.1 Identify Cause of Error and Restore the Standby Server		
4.1.1.1 Stop Mirroring Controller	·	
4.1.1.1.2 Recovery of the Mirroring Controller management directory		
4.1.1.3 Identify cause of error and perform recovery		
4.1.1.2 Rebuild the Standby Server		
4.1.1.3 Failback of the Primary Server944.1.2 Operations when the Server has Started Degrading after a Disconnection has Occurred954.1.2.1 Identify Cause of Error and Restore the Standby Server964.1.2.1.1 Stop Mirroring Controller964.1.2.1.2 Recovery of the Mirroring Controller management directory974.1.2.1.3 Identify cause of error and perform recovery974.1.2.2 Rebuild the Standby Server97		
4.1.2 Operations when the Server has Started Degrading after a Disconnection has Occurred.954.1.2.1 Identify Cause of Error and Restore the Standby Server.964.1.2.1.1 Stop Mirroring Controller.964.1.2.1.2 Recovery of the Mirroring Controller management directory.974.1.2.1.3 Identify cause of error and perform recovery.974.1.2.2 Rebuild the Standby Server.97		
4.1.2.1 Identify Cause of Error and Restore the Standby Server.964.1.2.1.1 Stop Mirroring Controller.964.1.2.1.2 Recovery of the Mirroring Controller management directory.974.1.2.1.3 Identify cause of error and perform recovery.974.1.2.2 Rebuild the Standby Server.97	·	
4.1.2.1.1 Stop Mirroring Controller		
4.1.2.1.2 Recovery of the Mirroring Controller management directory.       97         4.1.2.1.3 Identify cause of error and perform recovery.       97         4.1.2.2 Rebuild the Standby Server.       97		
4.1.2.1.3 Identify cause of error and perform recovery		
4.1.2.2 Rebuild the Standby Server		
$\cdot$		
2 2 2 2 2 2 3 3 3 3 3 3 3 3 3	4.1.3 Addressing Errors During Degrading Operation	

4.2 Action Required when Automatic Switch Fails	
4.3 Action Required when Automatic Disconnection Fails	
4.4 Action Required when All Database Servers or Instances Stopped	99
4.5 Recovering from an Incorrect User Operation.	102
Chapter 5 Managing Mirroring Controller Using WebAdmin	104
5.1 Mirroring Controller Setup	105
5.2 Edit Mirroring Controller Setup	107
5.3 Mirroring Controller Configuration	108
5.4 Stopping Mirroring Controller	109
5.5 Starting Mirroring Controller	110
5.6 Disabling Failover Mode	110
5.7 Enabling Failover Mode	
5.8 Deleting Mirroring Controller Setup	
5.9 Status Update after Failover	
5.10 Action Required when an Error Occurs in the Combined Admin Network and Log Transfer Network	
5.11 Performing Automatic Degradation Using the Arbitration Server	112
Appendix A Parameters	114
A.1 Parameters Set on the Primary Server.	114
A.2 Parameters Set on the Standby Server	116
A.3 Network Configuration File	118
A.4 Server Configuration File	121
A.4.1 Server Configuration File for the Database Servers	
A.4.2 Arbitration Configuration File	129
Appendix B Supplementary Information on Building the Primary Server and Standby Server on the Same Server	
B.1 Backup Data Storage Destination Directory	134
B.2 Registering Service Names and Event Source Names in the Windows Service	
B.3 How to Execute the mc_ctl Command	134
Appendix C User Commands	136
C.1 Fencing Command	
C.2 Arbitration Command.	137
C.3 State Transition Commands	
C.3.1 Post-switch Command	
C.3.2 Pre-detach Command	
C.3.3 Post-attach Command	139
Appendix D Notes on Performing Automatic Degradation Immediately after a Heartbeat Abnormality	141
Appendix E Supplementary Procedure on Configuring for Operation in Database Multiplexing Mode	144
E.1 Security Policy Settings	
E.2 Windows Firewall Settings.	144
Appendix F WebAdmin Disallow User Inputs Containing Hazardous Characters	146
Appendix G Collecting Failure Investigation Data	147
Index	148

## Chapter 1 Overview of Database Multiplexing Mode

This chapter provides an overview of database multiplexing mode.



In this and subsequent chapters, the word "Mirroring Controller" may be used in the process or management directory name or explanation.

## 1.1 What is Database Multiplexing Mode

Database multiplexing mode is an operation mode (log shipping mode) based on PostgreSQL streaming replication. Other software such as cluster software is not required.

This mode replicates the database on all servers that comprise the cluster system. It achieves this by transferring the updated transaction logs of the database from the server that receives the updates (primary server) to another server (standby server), and then reflecting them on the standby server. The client driver automatically distinguishes between the primary and standby servers, so applications can be connected transparently regardless of the physical server.

It consists of a feature that detects faults in the elements that are essential for the continuity of the database operation (such as the database process, disk, and network), as well as simplified switchover and standby server disconnection features. Furthermore, referencing can be performed on the standby server. The database will be copied in synchronous mode.



If using WebAdmin or Mirroring Controller, FUJITSU Enterprise Postgres supports cluster systems comprising one primary server and one standby server.

- Although it is possible to connect an asynchronous standby server to the cluster system as an additional server, the standby server is not targeted for monitoring by Mirroring Controller.
- A synchronous standby server cannot be connected to the cluster system as an additional server.



The streaming replication feature is not described in this manual.

Refer to "High Availability, Load Balancing, and Replication" in the PostgreSQL Documentation for information on the streaming replication feature.

Figure 1.1 Failover from the primary server to the standby server

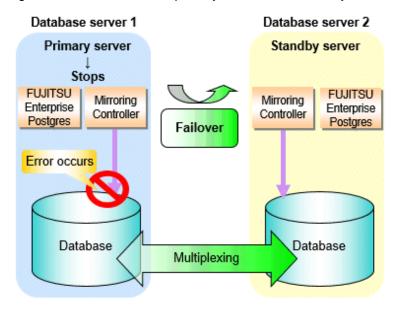
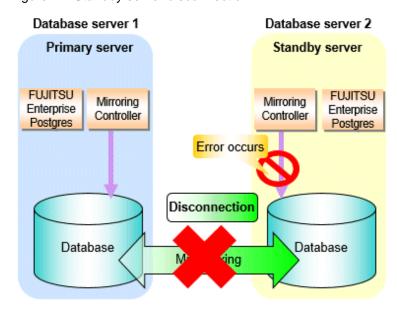


Figure 1.2 Standby server disconnection



#### Database degradation using the arbitration server

FUJITSU Enterprise Postgres provides the Server Assistant that objectively determines the status of database servers as a third party, and if necessary, isolates affected databases if the database servers are unable to accurately ascertain their mutual statuses in database multiplexing mode, such as due to a network error between database servers, or server instability. Database degradation can be performed by using the server (arbitration server) on which the Server Assistant is installed.

For database degradation using the arbitration server, if the database servers are unable to check their mutual statuses (due to a network error between database servers or server instability), then the database server queries the arbitration server for the status of the other database server. If it is determined based on the heartbeat result that the status is unstable, the applicable database server will be isolated from the cluster system (fencing). The arbitration server periodically heartbeats the database server so that it can respond immediately to queries from the database server. The fencing process can be customized according to the environment where Mirroring Controller is used.

Additionally, the database servers are always performing their heartbeats for the arbitration server so that it can perform check requests any time.

Database server 1 Database server 2 Primary server Standby server Abnormality Network issue between servers, detection or one server is unresponsive **FUJITSU FUJITSU** Mirroring Enterprise Enterprise Controller Postgres Postgres Failover Database Database Check status Arbitration server Fencing FUJITSU Enterprise Postgres Server Assistant

Figure 1.3 Database degradation using the arbitration server

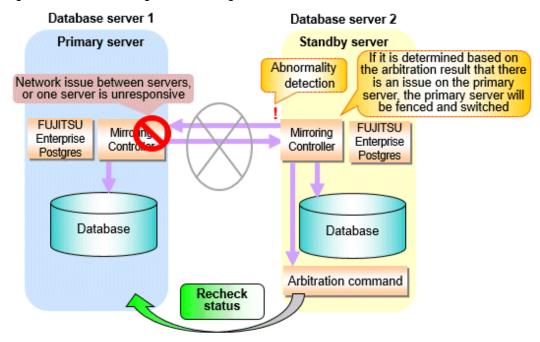


Install the arbitration server on a different physical server to that of the database server. Refer to "1.2 System Configuration for Database Multiplexing Mode" for information on the system configuration when using the arbitration server.

#### Database degradation using the arbitration command

The arbitration command is a user command that performs arbitration processing in lieu of the arbitration server. If an arbitration server cannot be deployed, arbitration of the database server can be performed using the arbitration command.

Figure 1.4 Database degradation using the arbitration command





Refer to "2.6 Creating a User Exit for a Database Server" or "Appendix C User Commands" for information on user commands.

## 1.1.1 Monitoring Using Database Multiplexing Mode

In database multiplexing mode, perform the monitoring below.

- Operating system or server failures, and no-response state

By generating a heartbeat between Mirroring Controller on each server, operating system or server errors are detected and acknowledged between the relevant servers.

The optimal operating method for environments where database multiplexing mode is performed can be selected from the following:

- Use the arbitration server to perform automatic degradation (switch/disconnect)

This is the default method.

The arbitration server objectively determines the status of database servers, then isolates and degrades from the cluster system the ones with an unstable status.

Refer to "Database degradation using the arbitration server" for details.

- Call the user exit (user command) that will perform the degradation decision, and perform automatic degradation

If the arbitration server cannot be installed, select if arbitration processing can be performed by the user instead.

Mirroring Controller queries the user exit on whether to degrade. The user exit determines the status of the database server, and notifies Mirroring Controller whether to perform degradation.

Refer to "Database degradation using the arbitration command" for details.

- Notification messages

Use this method if using a two-database server configuration.

Mirroring Controller outputs messages to the event log when an abnormality is detected. This ensures that a split brain will not occur due to a heartbeat abnormality - however, automatic switching will not be performed if the primary server operating system or server fails or becomes unresponsive.

- Perform automatic degradation unconditionally after a heartbeat abnormality

This method is handled as in FUJITSU Enterprise Postgres 9.6 or earlier versions.

This method is not recommended, because Mirroring Controller unconditionally will perform automatic degradation after heartbeat abnormalities.

- Database process failures, and no-response state

Mirroring Controller periodically accesses the database processes and checks the status. A process error is detected by monitoring whether an access timeout occurs.

- Disk failure

Mirroring Controller periodically creates files on the data storage destination disk below. A disk error is detected when an I/O error occurs.

- Data storage destination disk
- Transaction log storage destination disk
- Tablespace storage destination disk

Failures that can be detected are those that physically affect the entire system, such as disk header or device power failures.

- Streaming replication issue

Mirroring Controller detects streaming replication issues (log transfer network and WAL send/receive processes) by periodically accessing the PostgreSQL system views.

- Mirroring Controller process failure and no response

In order to continue the monitoring process on Mirroring Controller, Mirroring Controller process failures and no responses are also monitored.

The Mirroring Controller monitoring process detects Mirroring Controller process failures and no responses by periodically querying the Mirroring Controller process. If an issue is detected, Mirroring Controller is automatically restarted by the Mirroring Controller monitoring process.



- If output of messages is selected as the operation to be performed when a heartbeat abnormality is detected during heartbeat monitoring of the operating system or server, automatic degradation will not be performed.
  - However, if an issue in the WAL send process is detected on the primary server, then the standby server will be disconnected, and as a result an automatic disconnection may be performed even if the standby server operating system or server fails or becomes unresponsive.
- You can select in the parameters if the primary server will be switched if a database process is unresponsive or if tablespace storage destination disk failure is detected. Refer to "Appendix A Parameters" for details.
- If the standby server was disconnected, Mirroring Controller will automatically comment out the synchronous\_standby\_names parameter in the postgresql.conf file of the primary server. Accordingly, you can prevent the application processing for the primary server being stopped.



If the role of primary server was switched to another server and then starts degrading, the original primary server will not become the standby server automatically. Remove the cause of the error, and then change the role of the original primary server to the server currently acting as standby server. Refer to "4.1 Action Required when Server Degradation Occurs" for details.

## 1.1.2 Referencing on the Standby Server

#### 1.1.2.1 If Prioritizing the Main Job on the Primary Server

If a reference job is performed on the standby server and the primary server is switched, this may impact the main job from the point of view of load and conflict. This is because, on the new primary server (that is, the original standby server), both the main job that was being executed on the original primary server and the reference job that was being continued on the original standby server will be processed.

Therefore, to degrade the reference job (so that the impact on the main job is reduced), you can select the user exit below to disconnect the reference job that was performed on the original standby server.

- Post-switch command



If continuing with the referencing job after switching the primary server, give careful consideration to the server resource estimates, and the likely impact on performance.

#### 1.1.2.2 If Performing the Referencing Job on the Synchronous Standby Server

If an issue such as a log transfer network failure obstructs the continuation of a job on the primary server, the standby server may be automatically disconnected from the cluster system.

Therefore, if operating the reference job on the assumption that the connection destination is the synchronous standby server, you can select to temporarily stop the job by using the user exit or the feature below, so that unexpected referencing of past data does not occur as a result of the disconnection.

- Pre-detach command
- Forced stoppage of the standby server instance on disconnection (specify in the parameter of the server configuration file)

Additionally, if the standby server is incorporated into the cluster system, reference jobs can be started or resumed by using the user exit below.

- Post-attach command



#### See

- Refer to "2.6 Creating a User Exit for a Database Server" or "Appendix C User Commands" for information on each user exit.
- Refer to "A.4.1 Server Configuration File for the Database Servers" for information on the server configuration file of the database server.



Mirroring Controller will continue processing regardless of the processing result of the above user exits and features.

## 1.2 System Configuration for Database Multiplexing Mode

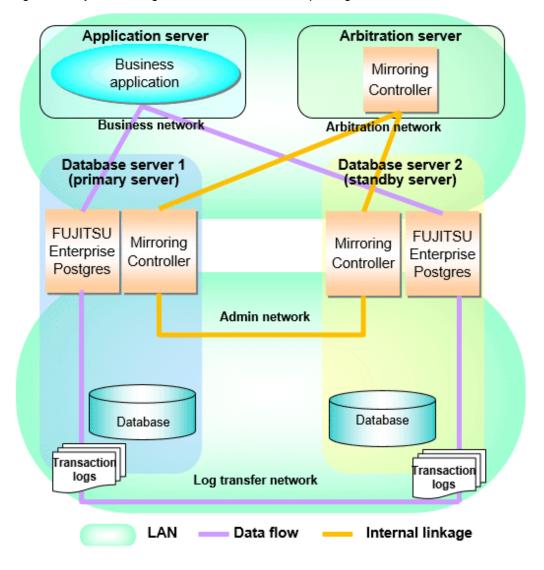
This section explains the products, features, and networks that are part of a database multiplexing system.

The following table shows the network types uses by database multiplexing systems.

Network type	Description
Job network	Network between the application that accesses the database, and the database
	server.

Network type	Description
Arbitration network	Network used by the arbitration server to check the status of the primary server and standby server, and communicate with Mirroring Controller of the database servers. Additionally, if the job network is disconnected from outside, it can also be used as the arbitration network. Refer to "1.4 Security in Database Multiplexing" for details on network security.
Admin network	Network used by the primary server and the standby server to monitor each other using Mirroring Controller, and to control Mirroring Controller of other servers.
Log transfer network	Network used to transfer the transaction logs of the database, which is part of database multiplexing.

Figure 1.5 System configuration for database multiplexing mode



The arbitration server is installed to check the database server status as a third party, and to perform fencing. Therefore, to obtain the intended benefits, consider the following.

- Install the arbitration server on a different server to that of the database server.
- For the arbitration network, use a network that will not be impacted by line faults or the load on the admin network or log transfer network. This is necessary to correctly determine issues on the admin network or log transfer network.



- The arbitration server can also be used as an application server. However, consider the server load.
- It is recommended to link the arbitration server with other cluster systems, in order to provide redundancy.
- Use the arbitration server in combination with the same version of FUJITSU Enterprise Postgres as that of the primary server and standby server.
- The arbitration server can be built on a different platform to that of the database server.



Because the ping command of the operating system is used for heartbeat monitoring of the database server, configure the network so that ICMP can be used on the admin network and the arbitration network.

#### 1.2.1 Mirroring Controller Resources

This section describes the database server and arbitration server resources of Mirroring Controller.

#### 1.2.1.1 Database Server Resources

The only Mirroring Controller resource is the Mirroring Controller management directory, which stores the files that define the Mirroring Controller behavior, and the temporary files that are created when Mirroring Controller is active.

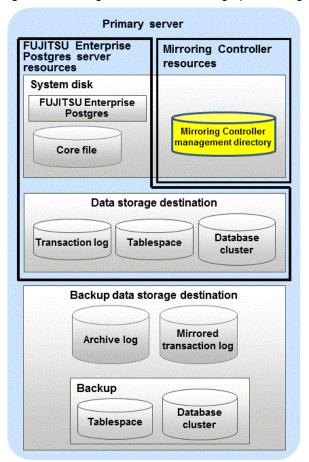


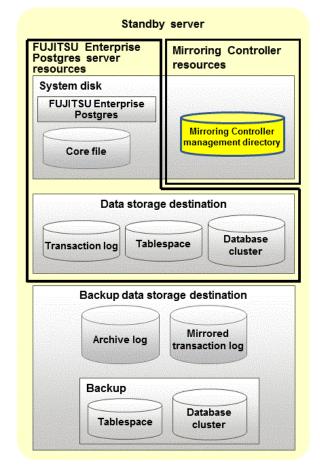
- Do not create the Mirroring Controller management directory in a directory managed by FUJITSU Enterprise Postgres, otherwise it may be deleted by mistake or may cause unexpected problems when FUJITSU Enterprise Postgres recovery is performed (such as old version of files being restored).

- Refer to "Preparing Directories for Resource Deployment" in the Installation and Setup Guide for Server for information on the directories managed by FUJITSU Enterprise Postgres.
- The backup methods described in "Backing Up the Database" in the Operation Guide cannot be used to back up the Mirroring Controller resources. Therefore, users must obtain their own backup of Mirroring Controller resources, in addition to FUJITSU Enterprise Postgres server resources. Retrieve backups after stopping Mirroring Controller.
- If the automatic switch/disconnection is enabled, do not edit synchronous\_standby\_names for the Mirroring Controller monitoring target instance. Otherwise, if Mirroring Controller is switched after editing, data may be lost or SQL access may stop.
- If you are building on a virtual machine or cloud, make sure the virtual machines are on different physical servers. Refer to your virtual machine software and cloud vendor documentation for instructions on how to deploy virtual machines.

The content on the primary server will be backed up. You cannot tell which server is the primary server to be backed up, because switching and failback may be performed between the servers. It is also impossible to tell which server is to be restored using the backed up data. Accordingly, ensure that you create a backup of each server when it is working as the primary server.

Figure 1.6 Configuration when backing up Mirroring Controller resources





#### 1.2.1.2 Arbitration Server Resources

The only arbitration server resource is the Mirroring Controller arbitration process management directory. This directory stores the files that define the Mirroring Controller arbitration process behavior and the temporary files created when Mirroring Controller is active.

## 1.2.2 Mirroring Controller Processes

This section describes the database server and arbitration server processes of Mirroring Controller.

#### 1.2.2.1 Database Server Processes

The database server processes comprise the Mirroring Controller process and Mirroring Controller monitoring process.

Process type	Description
Mirroring Controller process	Performs operating system/server and process heartbeat monitoring and disk abnormality monitoring between database servers. Additionally, it issues arbitration requests to the arbitration server.
Mirroring Controller monitoring process	Performs heartbeat monitoring of the Mirroring Controller process. If the Mirroring Controller process returns no response or is down, the monitoring process is restarted automatically.

#### 1.2.2.2 Arbitration Server Process

The only arbitration process is the Mirroring Controller arbitration process.

Process type	Description
Mirroring Controller arbitration process	Performs rechecks for issues detected on the primary server or the standby server.  Additionally, this process performs fencing if it determines that there is an issue on the primary server or the standby server.

#### 1.2.3 Redundancy of the Admin and Log Transfer Networks

The admin network is an important one, because it is used by Mirroring Controller to check the status of each server.

Additionally, the log transfer network is an important one, because it is necessary to ensure data freshness.

Accordingly, configure a failure-resistant network by implementing network redundancy via the NIC teaming feature provided by the operating system or network driver vendor.



NIC teaming feature is provided by the operating system by default.

#### 1.2.4 Notes on CPU Architecture and Products

Use the same CPU architecture (endian) for the primary server, standby server, and the arbitration server.

A server using only PostgreSQL streaming replication cannot be specified as the database multiplexing system log transfer destination.

## 1.3 Deciding on Operation when a Heartbeat Abnormality is Detected

The operation to be performed when a heartbeat abnormality is detected using operating system/server heartbeat monitoring is decided on according to the environment where database multiplexing mode is performed or the operating method.

It is possible to select from the four operations below, and specify this in the parameters of Mirroring Controller:

- Use the arbitration server to perform automatic degradation (switch/disconnect)
- Call the user exit (user command) that will perform the degradation decision, and perform automatic degradation
- Notification messages
- Perform automatic degradation unconditionally (switch/disconnect)

The table below shows if jobs can be continued on the primary server when an issue is detected during heartbeat monitoring of the operating system/server.

## Continuation of jobs on the primary server when an issue is detected during heartbeat monitoring of the operating system/server

Operation	Abnormal event						
	Server/operating system failures or no responses		Admin network	Log transfer	Issue on a network for both admin and		
	Primary server	Standby server	issue	Hetwork issue	log transfer		
Automatic degradation using	Y	Y	Y	Y	Y		
the arbitration server	(switch)	(disconnect)		(disconnect)	(disconnect)		
Call a user exit and perform automatic degradation	Y (switch)	Y (disconnect)	Y	Y (disconnect)	Y (disconnect)		

	Abnormal event						
Operation	Server/operating system failures or no responses		Admin network	Log transfer network issue	Issue on a network for both admin and		
	Primary server	Standby server	issue	network issue	log transfer		
Notification messages	N (message notification only)	N (message notification only)	Y	Y (disconnect)	Y (disconnect)		
Unconditional automatic degradation	Y (switch)	Y (disconnect)	N (split brain occurs)	Y (disconnect)	N (split brain occurs)		

Y: Job can be continued

N: Job cannot be continued

## 1.4 Security in Database Multiplexing

The database server replicates the database on all servers that comprise the cluster system. It achieves this by transferring and reflecting the updated transaction logs of the database from the primary server to the standby server.

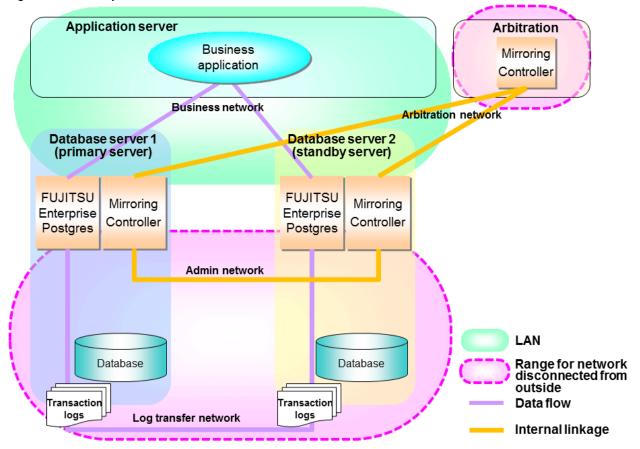
To safeguard the database against unauthorized access and preserve data confidentiality in transaction log transfers, carefully consider security and take note of the following when performing database multiplexing:

- Do not use trust authentication when using replication connection.
- Configure the admin network and the log transfer network so that they cannot be connected from the outside, as shown in Figure 1.7 Security.

Additionally, for the line on which Mirroring Controller connects from the database server to the arbitration server, take note of the following points and consider security carefully.

- Build a network with the arbitration server disconnected from outside, as shown in Figure 1.7 Security.

Figure 1.7 Security



However, it may not always be possible to adopt the configuration mentioned above. For example, you may want to place the servers in a nearby/neighboring office to minimize network delays.

In this case, combine the following features to enhance security:

- Authentication of the Standby Server
- Encryption of Transaction Logs Transferred to the Standby Server

When these features are combined, security will be achieved as shown below.



If the job network is disconnected from outside, it can be used as the arbitration network. However, if a network is to be used as both a job network and arbitration network, consider the load on the network.



If a port is blocked (access permission has not been granted) by a firewall, etc., enable use of the target port by granting access. Refer to the vendor document for information on how to open (grant access permission to) a port. Consider the security risks carefully when opening ports.

Job application **Primary server** Standby server Encrypted Database Database Encrypted Encrypted Obtain updated information Reflected Data transfer Encrypted SSL-encrypted communication Connection using standby server authentication Storage data encryption

Figure 1.8 Security achieved when standby server authentication is combined with transaction log encryption



Refer to "Performing Database Multiplexing" under "Configuring Secure Communication Using Secure Sockets Layer" in the Operation Guide for information on encrypting SSL communications.

### 1.4.1 Authentication of the Standby Server

You can prevent spoofing connections from an external server purporting to be the standby server by using authentication with a user name and password.

Configure the setting in the primary server  $pg_h$ ba.conf file so that authentication is performed for connections from the standby server in the same way as for connections from the client.



 $Refer \ to \ "Client \ Authentication" \ in \ the \ PostgreSQL \ Documentation \ for \ information \ on \ content \ that \ can \ be \ configured \ in \ pg\_hba.conf.$ 

## 1.4.2 Encryption of Transaction Logs Transferred to the Standby Server

In case the authentication of the standby server is breached so that a malicious user purporting to be the standby server can spoof data, the transaction log data can be encrypted to prevent it from being deciphered. The transparent data encryption feature is used to encrypt the data.



## Chapter 2 Setting Up Database Multiplexing Mode

This chapter describes how to set up database multiplexing mode, and how to check it.

#### Users who perform setup and operations on the database server

Setup and operations of the database server must be performed by the instance administrator user with administrator privileges (operating system user ID belonging to the Administrators group). The instance administrator user must be assigned the "Log on as a service" user right.

### Users who perform setup and operations on the arbitration server

The following users may perform setup and operations on the arbitration server when it is used for automatic degradation.

#### Linux

Any operating system user.

#### Windows

Any user with administrator privileges. This user must be assigned the "Log on as a service" user right.



- Mirroring Controller selects a database superuser as the user who will connect to the database instance. This enables instance administrator users and database superusers who operate the Mirroring Controller commands to run database multiplexing mode in different environments.

- The application name for connecting to the database instance is "mc\_agent".

### Matching the system times

Before starting the setup, ensure that the times in the primary server, standby server and arbitration server match, by using the operating system time synchronization feature, for example.

The tolerated difference is approximately one second.

If the system times are not synchronized (because the tolerated difference is exceeded, for example), problem investigation may be affected.

#### **Configuring ICMP**

Because the ping command of the operating system is used for heartbeat monitoring of the database server, configure the network so that ICMP can be used on the admin network and the arbitration network. Refer to the relevant operating system procedure for details.

### Setup

The setup procedure is shown in the table below. However, the procedure on the arbitration server should be performed only when the arbitration server is used for automatic degradation. A distinction is made between the procedures on the primary server and standby server according to whether the arbitration server is used.

Ston		Task	Refer to	
Step	Primary server	Standby server	Arbitration server	Kelei to
1		Installation		2.1 Installation
2	Preparing the database server		Preparing the arbitration server	2.2 Preparing for Setup
3			Configuring the arbitration server	2.3.1 Configuring the Arbitration Server
4			Creating a user exit	2.3.2 Creating a User Exit for the Arbitration Server

Cton		Defer to		
Step	Primary server	Standby server	Arbitration server	Refer to
5			Starting the arbitration process	2.3.3 Starting the Mirroring Controller Arbitration Process
6	Setting up database multiplexing mode			2.4.1 Setting Up Database Multiplexing Mode on the Primary Server
7	Creating, setting, and registering the instance			2.4.2 Creating, Setting, and Registering the Primary Server Instance
8	Creating a user exit			2.6 Creating a User Exit for a Database Server
9	Starting Mirroring Controller			2.4.3 Starting Mirroring Controller on the Primary Server
10		Setting up database multiplexing mode		2.5.1 Setting Up Database Multiplexing Mode on the Standby Server
11		Creating, setting, and registering the instance		2.5.2 Creating, Setting, and Registering the Standby Server Instance
12		Creating a user exit		2.6 Creating a User Exit for a Database Server
13		Starting Mirroring Controller		2.5.3 Starting Mirroring Controller on the Standby Server
14	Confirming the streaming replication status			2.7 Confirming the Streaming Replication Status
15	Checking the connection status			2.8.1 Checking the Connection Status on a Database Server
16		Checking the connection status		2.8.1 Checking the Connection Status on a Database Server
17			Checking the connection status	2.8.2 Checking the Connection Status on the Arbitration Server
18	Creating applications			2.9 Creating Applications
19	Checking the behavior			2.10 Checking the Behavior

Explanations for each step are provided below.



- The setup procedure is also the same when changing the mode on a single server to database multiplexing mode. In this case, omit the installation of FUJITSU Enterprise Postgres and the creation of the instance.

Refer to "3.9.2 Changing from Single Server Mode to Database Multiplexing Mode" for details.

- The primary and standby server can be pseudo-configured on the same server for system testing, for example. In this case, the setup can be performed using the same procedure, however there will be some supplementary steps.

Before performing the setup, refer to "Appendix B Supplementary Information on Building the Primary Server and Standby Server on the Same Server".

### 2.1 Installation

Refer to the manuals below, and then install the product.



- Refer to the Installation and Setup Guide for Server for details on how to install FUJITSU Enterprise Postgres.
- Refer to the Installation and Setup Guide for Server Assistant for information on installing the Server Assistant on the arbitration server.



Do not use the arbitration server also as a database server. The arbitration server is installed to check the database server status as a third party, and to perform fencing. Using the arbitration server also as a database server nullifies the effectiveness of the arbitration server.

### 2.2 Preparing for Setup

This section describes the preparation required before setting up Mirroring Controller.

### 2.2.1 Preparing the Database Server

### 2.2.1.1 Preparing the Backup Disk

In Mirroring Controller, by performing a backup, recovery is possible even if all server disks are corrupted.

The content on the primary server should be backed up. However, through switching and failback, the standby server may also become the primary server. Accordingly, prepare each of the backup disk devices for the primary and standby servers. Perform backup on the primary server used at the time of the backup.

### 2.2.1.2 Preparatory Tasks for the Output of Error Logs to the Event Log

This section explains the preparatory tasks required to output error logs to the event log.



If you do not register an event source name, the message content output to the event log may be incomplete.

### Setting each server

You should register this default event source name beforehand because the default event source name "MirroringControllerOpen" may be output to the event log when Mirroring Controller commands on the database server are used.

Example)

The following is an example in which the DLL of a 64-bit product is registered under the default event source name. Note that "<x>" indicates the product version.

> regsvr32 "c:\Program Files\Fujitsu\fsepv<x>server64\lib\mcevent.dll"

#### Setting each instance

You can output messages to any event source named by the user, so that messages output to the event log can be identified by each instance.

Example)

The following is an example in which the DLL of a 64-bit product is registered under the event source name "Mirroring Controller instl". Note that "<x>" indicates the product version.

> regsvr32 /n /i:"Mirroring Controller instl" "c:\Program Files\Fujitsu\fsepv<x>server64\lib
\mcevent.dll"

The parameter must be edited for each instance. Refer to "A.4.1 Server Configuration File for the Database Servers" to set the event\_source parameter.

#### If installing multiple versions

If FUJITSU Enterprise Postgres is already installed on the same machine, search for the key below in a registry editor, and make a note of the path of the registered DLL. Afterwards, register a new DLL using the default event source name.

Use the DLL path that you made a note of in the above step when re-registering the default event source name during uninstallation.

MirroringControllerOpen

### 2.2.1.3 Security Policy Settings

Security settings that allow logon as a service are required in Mirroring Controller for the operating system user account of the instance administrator user in order to start and stop Mirroring Controller and an instance using a Windows service.

If the security settings have not been configured, refer to "E.1 Security Policy Settings" for information on configuring the settings.

### 2.2.2 Preparing the Arbitration Server

### 2.2.2.1 Preparing to Output Error Logs to the Event Log (Windows)

This section explains the preparatory tasks for outputting error logs to the event log.



If no event source name is registered, messages output to the event log may be incomplete.

### Configuring each server

Event logs for the Mirroring Controller commands on the arbitration server may be output with the default event source name "MirroringControllerArbiter". Therefore, register this default event source name beforehand.

Example)

The following is an example in which the DLL of a 64-bit product is registered under the default event source name. Note that "<x>" indicates the product version.

> regsvr32 "c:\Program Files\Fujitsu\fsepv<x>assistant64\lib\mcarbevent.dll"

### Setting each instance

You can output messages to any event source named by the user, so that messages output to the event log can be identified by each instance.

Example)

The following is an example in which the DLL of a 64-bit product is registered under the event source name "Mirroring Controller arbiterl". Note that "<x>" indicates the product version.

> regsvr32 /n /i:"Mirroring Controller arbiter1" "c:\Program Files\Fujitsu\fsepv<x>assistant64\lib \mcarbevent.dll"

The parameter must be edited for each instance. Refer to "A.4.2 Arbitration Configuration File" and set the event\_source parameter.

### If installing multiple versions

If FUJITSU Enterprise Postgres is already installed on the same machine, search for the key below in Registry Editor, and make a note of the path of the registered DLL. Afterwards, register a new DLL using the default event source name.

Use the DLL path that you made a note of in the above step when re-registering the default event source name during uninstallation.

MirroringControllerArbiter

### 2.2.2.2 Security Policy Settings (Windows)

On the arbitration server, operating system user accounts that operate the Mirroring Controller arbitration process must be assigned the "Log on as a service" user right in order to use Windows Services to start and stop the Mirroring Controller arbitration process.

If the security settings to enable this have not been configured, refer to "E.1 Security Policy Settings" and configure the settings.

### 2.3 Setting Up the Arbitration Server

This section explains how to set up the arbitration server.

### 2.3.1 Configuring the Arbitration Server

This section explains how to set up database multiplexing mode on the arbitration server.

In database multiplexing mode, the files that are required for operations are managed in the Mirroring Controller arbitration process management directory.

There is one Mirroring Controller arbitration process management directory for each arbitration process.



The arbitration process for each database multiplexing system can be started on a single arbitration server.



#### See

- Refer to the Reference for information on the mc\_arb command.
- Refer to "Appendix A Parameters" for information on the parameters to be edited for the setup.

Perform the following procedure:

#### Linux

- 1. On the arbitration server, log in as any operating system user who starts and stops the arbitration process.
- 2. Configure the environment variables.

Set the following environment variables:

- PATH

Add the installation directory "/bin".

- MANPATH

Add the installation directory "/share/man".

- LD\_LIBRARY\_PATH

Add the installation directory "/lib".

### Example

The following example configures environment variables when the installation directory is "/opt/fsepv<x>assistant".

Note that "<*x*>" indicates the product version.

#### sh, bash

- \$ PATH=/opt/fsepv<x>assistant/bin:\$PATH ; export PATH
- \$ MANPATH=/opt/fsepv<x>assistant/share/man:\$MANPATH ; export MANPATH
- \$ LD\_LIBRARY\_PATH=/opt/fsepv<x>assistant/lib:\$LD\_LIBRARY\_PATH ; export LD\_LIBRARY\_PATH

#### csh, tcsh

```
$ setenv PATH /opt/fsepv<x>assistant/bin:$PATH
$ setenv MANPATH /opt/fsepv<x>assistant/share/man:$MANPATH
$ setenv LD_LIBRARY_PATH /opt/fsepv<x>assistant/lib:$LD_LIBRARY_PATH
```

- 3. Create the Mirroring Controller arbitration process management directory that will store the files required by the arbitration server.

  Use ASCII characters in the Mirroring Controller arbitration process management directory.
- 4. In the network configuration file (network.conf), define the Mirroring Controller network configuration that will be managed by the Mirroring Controller arbitration process.

Create network.conf in the Mirroring Controller arbitration process management directory, based on the sample file. For network.conf, set read and write permissions only for the operating system user who starts and stops the arbitration process in step 1

If users other than this are granted access permissions, the mc\_arb command will not work. Accordingly, users other than the operating system user who starts and stops the arbitration process in step 1 are prevented from operating the Mirroring Controller arbitration process.

Sample file

```
/installDir/share/mcarb_network.conf.sample
```

In network.conf, specify the IP address or host name and port number of the primary server and standby server, and define the Mirroring Controller network configuration that will be managed by the Mirroring Controller arbitration process.

Refer to "A.3 Network Configuration File" for details.

A definition example is shown below.

```
Example) The IDs of the servers are set to "server1" and "server2", and their port numbers are set to "27541".
```

```
server1 192.0.3.100 27541
server2 192.0.3.110 27541
```

5. In the arbitration configuration file (arbitration.conf), define the information related to control of the Mirroring Controller arbitration process.

Create arbitration.conf in the Mirroring Controller arbitration process management directory, based on the sample file. For arbitration.conf, set read and write permissions only for the operating system user who starts and stops the arbitration process in step 1. If users other than this are granted access permissions, the mc\_arb command will not work.

Sample file

```
/installDir/share/mcarb_arbitration.conf.sample
```

Set the parameters shown in the table below in arbitration.conf.

Table 2.1 Parameters

Parameter	Content specified	Remarks
port	Port number of the Mirroring Controller arbitration process	The port number must be 0 to 65535. Ensure that the port number does not conflict with other software. Do not specify an ephemeral port that may temporarily be assigned by another program.
my_address	'ipAddrOrHostNameThatAcceptsConnectionFro mMirroringControllerProcessOnDbServer' [Setting example] my_address = '192.0.3.120'	IPv4 and IPv6 addresses can be specified.  Specify the IP address, enclosed in single quotation marks (').

Parameter	Content specified	Remarks
syslog_ident	'programName'	Specify using single quotation marks (') to enclose the program name used to identify the Mirroring Controller arbitration process message in the system log. Use ASCII characters excluding spaces to specify this parameter.  The default is 'MirroringControllerArbiter'.
fencing_command	'fencingCmdFilePath'  [Setting example]  fencing_command = '/arbiter/fencing_dir/ execute_fencing.sh'	Specify the full path of the fencing command that fences a database server where it is determined that an error has occurred.  Enclose the path in single quotation marks (').  Specify the path using less than 1024 bytes.
fencing_command_ti meout	Timeout for fencing command (seconds)	If the command does not respond within the specified number of seconds, it is determined that fencing has failed and a signal (SIGTERM) is sent to the fencing command execution process.  Specify a value between 1 and 2147483647.  The default is 20 seconds.



Refer to "A.4.2 Arbitration Configuration File" for information on the parameters and for other parameters.

### Windows

- 1. On the arbitration server, log in as any operating system user who starts and stops the arbitration process.
- 2. Configure the environment variables.

Set the following environment variable:

- PATH

Add the installation folders "\bin" and "\lib".

### Example

The following example configures environment variables when the installation folder is "c:\Program Files\Fujitsu \fsepv<*x*>assistant64".

Note that "<x>" indicates the product version.

> SET PATH=c:\Program Files\Fujitsu\fsepv<x>assistant64\bin;c:\Program Files\Fujitsu
\fsepv<x>assistant64\lib;%PATH%

- 3. Create the Mirroring Controller arbitration process management directory that will store the files required by the arbitration server.

  Use ASCII characters in the Mirroring Controller arbitration process management directory.
- 4. In the network configuration file (network.conf), define the Mirroring Controller network configuration that will be managed by the Mirroring Controller arbitration process.

Create network.conf in the Mirroring Controller arbitration process management directory, based on the sample file.

Sample file

```
installDir\share\mcarb_network.conf.sample
```

In network.conf, specify the IP address or host name and port number of the primary server and standby server, and define the Mirroring Controller network configuration that will be managed by the Mirroring Controller arbitration process.

Refer to "A.3 Network Configuration File" for details.

A definition example is shown below.

#### Example)

The IDs of the servers are set to "server1" and "server2", and their port numbers are set to "27541".

```
server1 192.0.3.100 27541
server2 192.0.3.110 27541
```

5. Change the access permissions for network.conf.

For network.conf, set read and write permissions only for the operating system user who starts and stops the arbitration process in step 1. If users other than this are granted access permissions, the mc\_arb command will not work. Accordingly, users other than the operating system user who starts and stops the arbitration process in step 1 are prevented from operating the Mirroring Controller arbitration process.

#### Example)

The following is an execution example, in which the operating system user who starts and stops the arbitration process in step 1 is granted full access permissions as the owner when the user is "fsepuser". The following procedure applies when the user is logged in to the Windows server as "fsepuser".

- > takeown /f network.conf
- > icacls network.conf /reset
- > icacls network.conf /inheritance:r
- > icacls network.conf /grant fsepuser:F
- 6. In the arbitration configuration file (arbitration.conf), define the information related to control of the Mirroring Controller arbitration process.

Create arbitration.conf in the Mirroring Controller arbitration process management directory, based on the sample file.

Sample file

```
installDir\share\mcarb_arbitration.conf.sample
```

Set the parameters shown in the table below in arbitration.conf.

Table 2.2 Parameters

Parameter	Content specified	Remarks	
port	Port number of the Mirroring Controller arbitration process	The port number must be 0 to 65535. Ensure that the port number does not conflict with other software. Do not specify an ephemeral port that may temporarily be assigned by another program.	
my_address	'ipAddrOrHostNameThatAcceptsConnectionF romMirroringControllerProcessOnDbServer' [Setting example] my_address = '192.0.3.120'	IPv4 and IPv6 addresses can be specified.  Enclose the parameter value in single quotation marks (').	
service_name	'registeredServiceNameOfArbitrationProcessI nWindowsServices'	Use ASCII characters excluding forward slash (/) and backslash (\) to specify this parameter.  Enclose the parameter value in single quotation marks (').	

Parameter	Content specified	Remarks
		The maximum length of the service name is 124 bytes.
event_source	'eventSourceNameUsedToIdentifyArbitration ProcessMsgInEventLog'	Use ASCII characters excluding spaces to specify this parameter.
		Enclose the parameter value in single quotation marks (').
		The maximum length of the event source name is 255 bytes.
fencing_command	'fencingCmdFilePath'  [Setting example]  fencing_command = 'c:\\arbiter\\fencing_dir\\execute_fencing.bat'	Specify the full path of the fencing command that fences a database server where it is determined that an error has occurred.  Specify "\\" as the delimiter.  Enclose the path in single quotation marks (').  Specify the path using less than 260 bytes.  Any multibyte characters must use the same encoding as the operating system.
fencing_command_time out	Timeout for fencing command (seconds)	If the command does not respond within the specified number of seconds, it is determined that fencing has failed and a signal (SIGTERM) is sent to the fencing command execution process.  Specify a value between 1 and 2147483647.  The default is 20 seconds.



Refer to "A.4.2 Arbitration Configuration File" for information on the parameters and for other parameters.

7. Change the access permissions for arbitration.conf.

For arbitration.conf, set read and write permissions only for the operating system user who starts and stops the arbitration process in step 1. If users other than this are granted access permissions, the mc\_arb command will not work.

### Example)

The following is an execution example, in which the operating system user who starts and stops the arbitration process in step 1 is granted full access permissions as the owner when the user is "fsepuser". The following procedure applies when the user is logged in to the Windows server as "fsepuser".

- > takeown /f arbitration.conf
- > icacls arbitration.conf /reset
- > icacls arbitration.conf /inheritance:r
- > icacls arbitration.conf /grant fsepuser:F

### 8. Configure Windows Firewall.

If Windows Firewall is used, enable the port number of Mirroring Controller specified in the network configuration file in step 3. Refer to "E.2 Windows Firewall Settings" for details.

9. Register the Mirroring Controller arbitration process as a Windows service.

Execute the mc\_arb command in register mode.

For the -P option of the mc\_arb command, specify the password of the operating system user who executes the command.

Example)

> mc\_arb register -M D:\mcarb\_dir\arbiter1 -P \*\*\*\*\*\*\*



When specifying the password in the -P option of the mc\_arb command, for security reasons, you should be careful not to allow other users to access it.



Use the mc\_arb command with the -S option to specify automatic start and stop of the Mirroring Controller arbitration process. Refer to "2.13 Setting Automatic Start and Stop of the Mirroring Controller Arbitration Process" for details.

The Mirroring Controller arbitration process is registered as a Windows service using the service name specified in the service\_name parameter of arbitration.conf in step 6.

You can execute the sc qc command to check the registration status.

### 2.3.2 Creating a User Exit for the Arbitration Server

The only user exit for the arbitration server is the fencing command.

The fencing command is a user command that is called by the Mirroring Controller arbitration process if Mirroring Controller performs arbitration processing and determines that a database server is unstable.

In the fencing command, the user implements a process that isolates a database server from a cluster system by, for example, stopping the target operating system or server. The fencing command that was created is to be specified for the parameter in the arbitration configuration file. Refer to "A.4.2 Arbitration Configuration File" for information on the parameters.

- Fencing the primary server during the switch
  - Prevent the Mirroring Controller management process on the primary server from communicating with the Mirroring Controller management process on the other server.
  - Prevent applications from connecting to the primary server instance.
- Fencing the standby server during disconnection
  - Prevent the Mirroring Controller management process on the standby server from communicating with the Mirroring Controller management process on the other server.
  - Prevent applications from connecting to the standby server instance.
  - Prevent the standby server from continuing streaming replication.



Refer to "Appendix C User Commands" for information on user exits.

### 2.3.3 Starting the Mirroring Controller Arbitration Process

This section explains how to start the Mirroring Controller arbitration process.

An operating system user who has logged in to the arbitration server can start the Mirroring Controller arbitration process by executing the mc\_arb command in start mode.

#### Linux

Example)

\$ mc\_arb start -M /mcarb\_dir/arbiter1

#### Windows

Example)

> mc\_arb start -M D:\mcarb\_dir\arbiter1

### 2.4 Setting Up the Primary Server

This section explains how to set up the primary server.

### 2.4.1 Setting Up Database Multiplexing Mode on the Primary Server

This section explains how to set up database multiplexing mode on the primary server.

In database multiplexing, the files that are required for operations are managed in the Mirroring Controller management directory.

There is one Mirroring Controller management directory for each instance.



- Do not place the Mirroring Controller management directory in a directory managed by FUJITSU Enterprise Postgres, otherwise it may be deleted by mistake with the directories managed by FUJITSU Enterprise Postgres, and an old version of files may be restored.



- Refer to "Preparing Directories for Resource Deployment" in the Installation and Setup Guide for Server for details on the directories that are managed by FUJITSU Enterprise Postgres.

- Refer to "mc\_ctl" in Reference for information on the command.
- Refer to "Appendix A Parameters" for details on each parameter to be edited for the setup.

Perform the following procedure:

- 1. Log in to the primary server.
- 2. Create the Mirroring Controller management directory that will store the files required by database multiplexing.

Use ASCII characters in the Mirroring Controller management directory.

Additionally, grant "Write" permission to the instance administrator user for the Mirroring Controller management directory.

In the network configuration file (network.conf), define the network configuration that will link between the Mirroring Controller processes.

Create the network.conf file in the Mirroring Controller management directory, based on the sample file.

Sample file

installDir\share\mc\_network.conf.sample

In network.conf, specify the IP address or host name and port number of the primary server and standby server, and define the network configuration that will link between the Mirroring Controller processes, and between Mirroring Controller processes and the Mirroring Controller arbitration process.

Refer to "A.3 Network Configuration File" for details.

A definition example is shown below.

The content to be defined depends on the operation settings at the time a heartbeat abnormality is detected.

#### When automatic degradation by the arbitration server is selected

```
Example
```

The IDs of the primary server and standby server are set to "server1" and "server2", and their port numbers are set to "27540" and "27541". The ID of the server of the Mirroring Controller arbitration process is set to "arbiter", and its port number is set to "27541".

```
server1 192.0.2.100,192.0.3.100 27540,27541 server
server2 192.0.2.110,192.0.3.110 27540,27541 server
arbiter 192.0.3.120 27541 arbiter
```

Ensure that the port numbers set for the primary server, standby server, and arbitration server do not conflict with other software. In addition, when the arbitration server is used for automatic degradation, use a network in which the arbitration network is not affected by a line failure in the admin network.

When the server type is "server", two IP addresses or host names, and two port numbers need to be specified in the following order:

- IP address or host name of the database server used as the admin network
- IP address or host name of the database server used as the arbitration network
- Port number of the database server used as the admin network
- Port number of the database server used as the arbitration network

If the server type is "arbiter", specify the IP address or host name set for the my\_address parameter and the port number set for the port parameter in arbitration.conf of the arbitration server.

#### When operation other than automatic degradation by the arbitration server is selected

```
Example) The IDs of the servers are set to "server1" and "server2", and their port numbers are set to "27540".
```

```
server1 192.0.2.100 27540
server2 192.0.2.110 27540
```

Ensure that the port numbers for the primary and standby server do not conflict with other software.

Register the port number of the primary server in the services file, because there are programs, such as WebAdmin, that search an available port number using the services file.

Register any name as the service name.

4. Change the access permissions for the network.conf file.

For network.conf, set read and write permissions for the instance administrator user only.

If users other than the instance administrator user are granted access permissions, the mc\_ctl command will not work. Accordingly, users other than the instance administrator user are prevented from operating Mirroring Controller.

```
Example)
```

The following is an execution example, in which the instance administrator user is granted full access permissions as the owner when the operating system user name of the instance administrator user is "fsepuser". The following procedure applies when the user is logged in to the Windows server as "fsepuser":

```
> takeown /f network.conf
> icacls network.conf /reset
> icacls network.conf /inheritance:r
> icacls network.conf /grant fsepuser:F
```

5. Define the information related to Mirroring Controller monitoring and control in the serverIdentifier.conf file.

Create the serverIdentifier.conf file in the Mirroring Controller management directory, based on the sample file.

As the file name for the *serverIdentifier*.conf file, use the server identifier name that was specified in the network.conf file in step 3. Sample file

 $install Dir \verb|\share| mc\_server.conf.sample|$ 

Set the parameters shown in the table below in the serverIdentifier.conf file.

Table 2.3 Parameters

Parameter	Content specified	Remarks
db_instance	'dataStorageDestinationDir'	Use ASCII characters, and specify "\\" as the path delimiter.
		Enclose the parameter value in single quotation marks (').
db_instance_service_na me	'registeredServiceNameOfFujitsuEnterprisePost gresInstance'	Specify the registered service name of the FUJITSU Enterprise Postgres instance in the Windows service.
		Use ASCII characters, enclosed in single quotation marks (').
db_instance_password	'passwordOfInstanceAdminUser'	If password authentication is performed, specify this parameter in the settings used when Mirroring Controller connects to a database instance.
		Use ASCII characters, enclosed in single quotation marks (').
		If the specified value of this parameter includes ' or  write \' or \ respectively.
enable_hash_in_passwor d	on or off	Specify on to treat the # in the db_instance_password specification as a password character, or off to treat it as a comment.
		The default is "off".
mc_service_name	'registeredServiceNameOfMirroringController'	Specify the Mirroring Controller service name registered in the Windows service. Use ASCII characters excluding forward slash (/) and backslash (\) to specify this parameter.
		The service name is up to 124 bytes.
event_source	'eventSourceName'	Specify the event source name to be used to identify the Mirroring Controller message in the event log. Use ASCII characters to specify this parameter.
		The maximum length of the event source name is 255 bytes.
		By using a similar event source name as the postgresql.conf file parameter, the Mirroring Controller output content can be referenced transparently, so log reference is easy.
remote_call_timeout	Admin communication timeout	Specify the timeout value (milliseconds) of the Mirroring Controller agent process for communication between servers.

Parameter	Content specified	Remarks	
		Specify a value that is less than the operation system TCP connection timeout.	
heartbeat_error_action	Operation when a heartbeat abnormality is detected using operating system or server	arbitration: Perform automatic degradation using the arbitration server.	
	heartbeat monitoring	command: Call a user exit to determine degradation, and perform automatic degradation if required.	
		message: Notify messages.	
		fallback: Perform automatic degradation unconditionally.	
		Set the same value on the primary server and standby server.	
heartbeat_interval	Interval time for abnormality monitoring during heartbeat monitoring of the operating system or server (milliseconds)	Abnormality monitoring of the operating system or server is performed at the interval (milliseconds) specified in	
heartbeat_timeout	Timeout for abnormality monitoring during heartbeat monitoring of the operating system or server (seconds)	heartbeat_interval.  This parameter setting is used as the default for database process heartbeat monitoring,	
heartbeat_retry	Number of retries for abnormality monitoring during heartbeat monitoring of the operating system or server (number of times)	streaming replication abnormality monitoring, and disk abnormality monitoring.	
		When setting the monitoring time, there are some considerations to take into account to optimize degradation using abnormality monitoring. Refer to "2.11.4.1 Tuning for Abnormality Monitoring of the Operating System or Server" for details.	
fencing_command	'fencingCmdFilePath' [Setting example]	Specify the full path of the fencing command that fences a database server where an error is determined to have occurred.	
	fencing_command = 'c:\\mc\\fencing_dir\ \execute_fencing.bat'	Enclose the path in single quotation marks (').	
	(o.tecute_tonemg.sut	Any multibyte characters must use the same encoding as the operating system.	
		This parameter must be specified when "command" is set for heartbeat_error_action.	
		Specify the path using less than 260 bytes.	
fencing_command_time out	Fencing command timeout (seconds)	If the command does not respond within the specified number of seconds, fencing is determined to have failed and a signal (SIGTERM) is sent to the fencing command execution process.	
		Specify a value between 1 and 2147483647.	
		The default is 20 seconds.	
arbitration_timeout	Timeout for arbitration processing in the Mirroring Controller arbitration process (seconds)	The specified value must be at least equal to the heartbeat monitoring time of the operating system or server + fencing_command_timeout in the arbitration configuration file.	

Parameter	Content specified	Remarks
		If there is no response for at least the number of seconds specified, the primary server will not be switched and the standby server will not be disconnected. Therefore, perform degradation manually.  Specify a value between 1 and 2147483647.  This parameter does not need to be set for operation that does not use the arbitration server.
arbitration_command	'ArbitrationCmdFilePath'  [Setting example]  arbitration_command = 'c:\\mc\\arbitration_dir\\execute_arbitration_command.bat'	Specify the full path of the arbitration command to be executed when an abnormality is detected during heartbeat monitoring of the operating system or server.  Enclose the path in single quotation marks (').  Any multibyte characters must use the same encoding as the operating system.  This parameter must be specified when "command" is set for heartbeat_error_action.  Specify the path using less than 260 bytes.
arbitration_command_ti meout	Timeout for arbitration command (seconds)	If the command does not respond within the specified number of seconds, it is determined that execution of the arbitration command has failed and a signal (SIGTERM) is sent to the arbitration command execution process.  Specify a value between 1 and 2147483647.  This parameter can be specified only when "command" is set for heartbeat_error_action.



Refer to "A.4.1 Server Configuration File for the Database Servers" for information on the parameters and for other parameters.

#### 6. Change the access permissions for the *serverIdentifier*.conf file.

For *serverIdentifier*.conf, set read and write permissions for the instance administrator user only. If users other than the instance administrator user are granted access permissions, the mc\_ctl command will not work.

### Example)

The following is an execution example, in which the instance administrator user is granted full access permissions when the operating system user name of the instance administrator user is "fsepuser". The following procedure applies when the user is logged in to the Windows server as "fsepuser":

- > takeown /f serverIdentifier.conf
- > icacls serverIdentifier.conf /reset
- > icacls serverIdentifier.conf /inheritance:r
- > icacls serverIdentifier.conf /grant fsepuser:F

### 7. Configure the Windows firewall.

If the Windows firewall feature is to be enabled, you should enable the port number of Mirroring Controller that you specified in the network definition file in step 3. Refer to "E.2 Windows Firewall Settings" for details.

8. Register Mirroring Controller to the Windows service.

Execute the mc\_ctl command in the register mode.

For the -P option of the mc\_ctl command, specify the password of the operating system user who executes the command.

Example)

> mc\_ctl register -M D:\mcdir\inst1 -P \*\*\*\*\*\*\*



When specifying the password in the -P option of the mc\_ctl command, for security reasons, you should be careful not to allow other users to access it.



You can use the mc\_ctl command with the -S option to specify automatic start and stop of Mirroring Controller. Refer to "2.12 Setting Automatic Start and Stop of Mirroring Controller and Multiplexed Instances" for details.

Using the service name specified in the mc\_service\_name parameter of *serverIdentifier*.conf in step 5, Mirroring Controller is registered to the Windows service.

You can execute the sc qc command to check the registration status.

### 2.4.2 Creating, Setting, and Registering the Primary Server Instance

This section explains how to create, set, and register the primary server instance.



- Mirroring Controller supports instances that are registered in the Windows service.



- Refer to "Client Authentication" in the PostgreSQL Documentation for information on the pg\_hba.conf file.
- Refer to "A.1 Parameters Set on the Primary Server" for information on the postgresql.conf file.
- Refer to "mc\_ctl" in Reference for information on the command.

Perform the following procedure:

- 1. Refer to "Setup" in the Installation and Setup Guide for Server, and then perform the FUJITSU Enterprise Postgres setup and create the FUJITSU Enterprise Postgres instance.
  - Use ASCII characters in the data storage destination directory.
  - When registering an instance to the Windows service, perform the settings required to enable Mirroring Controller to start and stop the instance. Execute the pg\_ctl command with the following specified for the register mode:
    - For the service name of the -N option, specify the name set for the db\_instance\_service\_name parameter in the server definition file
    - Specify "demand" for the -S option, so that the service does not start automatically on startup of the system



- If degradation starts occurring due to an error during operations in database multiplexing mode, recovery is required for the standby server. There are some conditions to execute the pg\_rewind command to recover the standby server. One of the conditions can be satisfied by enabling checksums when executing the initdb command. This is not mandatory. Refer to "4.1.1.1.3 Identify cause of error and perform recovery" for details.
- Do not configure the Windows service of a multiplexed instance to perform automatic start, as it is started by Mirroring Controller.
- 2. When using transparent data encryption, configure the encryption settings for the storage data.

Create the keystore file.

Refer to "Database Multiplexing Mode" in the Operation Guide for details, and then configure the settings.

3. Add the following entry to the pg\_hba.conf file to authenticate connections from the standby server.

Copy the file to the standby server later.

# TYPE	DATABASE	USER	ADDRESS	METHOD	
host	replication	fsep	standbyServerAddress	$authentication  exttt{Method}$	
host	replication	fsep	primaryServerAddress	${\it authentication} {\it Method}$	

For the primary and standby server addresses, specify the IP address that will connect to the log transfer network.

Additionally, all servers can be used as the primary server or the standby server, so add entries for the addresses of all servers that comprise the database multiplexing system.



#### Setting an authentication method other than trust authentication

If the primary server becomes the standby server, to perform automatic authentication of connections to the primary server, create a password file (%APPDATA%\postgresql\pgpass.conf), and then specify a password for the replication database. Accordingly, the instance administrator operating system user and the user registered in the database will be the same, so you can verify that the connection was not made by an unspecified user. Additionally, the password that was set beforehand will be used in the authentication, so that the connection will be automatic.



If trust authentication is set, all OS users who can log in to the primary server will be able to connect, and if one of these is a malicious user, then that user can corrupt the standby server data, or cause the job system to fail, by sending an erroneous transaction log. Therefore, decide on the authentication method according to the security requirements of the system using database multiplexing mode.

Refer to "Authentication Methods" in the PostgreSQL Documentation for details on the authentication methods that can be set.

4. Configure this setting to enable the instance administrator user of the primary server to connect as a database application.

This setting enables the connection to the instance using the user name of the instance administrator user, so that Mirroring Controller can monitor instance errors. Configure this setting to enable the connection to the postgres database.

- If password authentication is used

In the db\_instance\_password parameter of the *serverIdentifier*.conf file, specify the password for the instance administrator user. This password is used to connect to the database instance. If a password is not specified in the db\_instance\_password parameter, the connection to the database instance from Mirroring Controller will fail, and it will not be possible to perform the process monitoring of the instance.

- If password authentication is not used

There is no need to specify the password in the db\_instance\_password parameter.

Even if the password for the instance administrator user is specified in the db\_instance\_password parameter, it will be ignored.

- If certificate authentication using SSL is used

Specify connection parameters for SSL in the db\_instance\_ext\_pq\_conninfo parameter and db\_instance\_ext\_jdbc\_conninfo parameter in the *serverIdentifier*.conf file. If the pasameters are not specified, the connection to the database instance from Mirroring Controller will fail, and it will not be possible to perform the process monitoring of the instance. If certificate authentication using SSL is not performed, the parameters specification is not required.

An example of setting the authentication method is shown below.

# TYPE	DATABASE	USER	ADDRESS	METHOD
host	postgres	fsep	127.0.0.1/32	authentication Method



Mirroring Controller uses the PostgreSQL JDBC 4.2 driver to connect to the database instance. Therefore, for the authentication method, specify a method supported by the JDBC driver. If an authentication method not supported by the JDBC driver is specified, Mirroring Controller will fail to start. Refer to the PostgreSQL JDBC Driver Documentation for information on authentication methods supported by the JDBC driver.

5. To use database multiplexing mode, specify the parameters shown in the table below in the postgresql.conf file.

The postgresql.conf file is copied when the standby server instance is created. Accordingly, set the required parameters in the standby server.

To use database multiplexing mode, specify the parameters shown in the table below in the postgresql.conf file. After editing the postgresql.conf file, restart the instance.

Table 2.4 Parameters

Parameter	Content specified	Remarks
wal_level	replica or logical	Specify "logical" when logical decoding is also to be used.
max_wal_senders	2 or more	Specify "2" when building a Mirroring Controller cluster system.
		When additionally connecting asynchronous standby servers to the cluster system, add the number of simultaneous connections from these standby servers.
synchronous_standby_names	'standbyServerName'	Specify the name that will identify the standby server.
		Enclose the name in single quotation marks (').
		Do not change this parameter while Mirroring Controller is running.
		Do not specify multiple names to this parameter as the Mirroring Controller can manage only one standby server.
hot_standby	on	Specify this to execute reference jobs on the standby server.
wal_keep_segments	Number of file segments	If a delay exceeding the value set in this parameter occurs, the WAL segment required later by the primary server may be deleted.
		Additionally, if you stop a standby server (for maintenance, for example), consider the stop time and set a value that will not cause the WAL segment to be deleted.
		Refer to "Estimating Transaction Log Space Requirements" in the Installation and Setup Guide for Server for information on estimating the WAL segment.

Parameter	Content specified	Remarks
wal_log_hints	on	When using the pg_rewind command to recover a standby server, specify this parameter or enable checksums when executing the initdb command.
wal_sender_timeout	Timeout (milliseconds)	Specify the time period after which it is determined that an error has occurred in the transaction log transfer on the primary server.
		By aligning this value with the value for the database process heartbeat monitoring time, you can unify the time after which it is determined that an error has occurred.
archive_mode	on	Specify the archive log mode.
archive_command	'cmd /c ""installDir\\bin\ \pgx_walcopy.cmd" "%p" "backupDataStorageDestin ationDir\\archived_wal\\ %f""	Specify the command and storage destination to save the transaction log.
backup_destination	Backup data storage destination directory	Specify the name of directory where to store the backup data.
		Set the permissions so that only the instance administrator user can access the specified directory.
		Specify the same full path on all servers, so that the backup data of other servers can be used to perform recovery.
max_connections	Number of simultaneous client connections to the instance + superuser_reserved_connections	The value specified is also used to restrict the number of connections from client applications and the number of connections for the management of instances.
		Refer to "When an Instance was Created with the initdb Command" in the Installation and Setup Guide for Server, and "Connections and Authentication" in the PostgreSQL Documentation, for details.
superuser_reserved_connections	Add the number of simultaneous executions of	Specify the number of connections reserved for connections from database superusers.
	mc_ctl status (*1) + 2	Add the number of connections from Mirroring Controller processes. Also reflect the added value in the max_connections parameter.
wal_receiver_timeout	Timeout (milliseconds)	Specify the time period after which it is determined that an error has occurred when the transaction log was received on the standby server.
		By aligning this value with the value for the heartbeat monitoring time of the database process, you can unify the time after which it is determined that an error has occurred.
restart_after_crash	off	If "on" is specified, or the default value is used for this parameter, behavior equivalent to restarting FUJITSU Enterprise Postgres, including crash recovery, will be performed when some server processes end abnormally.
		However, when database multiplexing monitoring is used, a failover will occur after an error is detected when some server processes end abnormally, and the restart of those server processes is forcibly stopped. Specify "off" to prevent behavior such as this from occurring for no apparent reason.

Parameter	Content specified	Remarks
synchronous_commit	on or remote_apply	Specify up to what position WAL send is to be performed before transaction commit processing returns a normal termination response to a client.  The recommended value is on or remote_apply to prevent data loss caused by operating system or server down immediately after a switch or switch.
recovery_target_timeline	latest	Specify "latest" so that the new standby server (original primary server) will follow the new primary server when a switch occurs.  This parameter is required when the original primary server is incorporated as a new standby server after the primary server is switched.

<sup>\*1:</sup> Number of simultaneous executions of the mc\_ctl command in the status mode.

### 2.4.3 Starting Mirroring Controller on the Primary Server

This section explains how to start Mirroring Controller on the primary server.

When the arbitration server is used for automatic degradation, start the Mirroring Controller arbitration process on the arbitration server in advance.

1. Start the Mirroring Controller process.

Enabling automatic switch/disconnection

As the instance administrator user, execute the mc\_ctl command in start mode.

Example)

```
> mc_ctl start -M D:\mcdir\inst1
```

### Disabling automatic switch/disconnection

As the instance administrator user, execute the mc\_ctl command in start mode with the -F option specified.

Example)

> mc\_ctl start -M D:\mcdir\inst1 -F



- When the arbitration server is used for automatic degradation, the database server must connect to the arbitration server, and as a result, Mirroring Controller startup may take longer than when the arbitration server is not used.
- If the parameter for heartbeat monitoring of operating systems or servers set by the abitration server is greater than parameter for heartbeat monitoring of operating systems and servers of the Mirroring Controller, the Mirroring Controller may fail to start. In this case, check the contents of the message notification and review the parameters for heartbeat monitoring of operating systems or servers for the arbitration server or Mirroring Controller.
- If the heartbeat\_error\_action parameter in *serverIdentifier*.conf is set to "message", even if automatic switch/disconnection is enabled and Mirroring Controller is started, only message output is performed when a heartbeat abnormality is detected during heartbeat monitoring of operating systems and servers switch/disconnection is not performed.
- Mirroring Controller startup usually fails if the standby server is mistakenly started as the primary server or if the old primary server is not recovered after the switch and is then mistakenly started as the primary server. However, if the admin network is disconnected, then startup does not fail, and both servers may become primary servers. Therefore ensure that the admin network is connected before starting Mirroring Controller.



- The mc\_ctl command fails if the Mirroring Controller arbitration process has not been started on the arbitration server when the arbitration server is used for automatic degradation. However, if the Mirroring Controller arbitration process cannot be started in advance, it can be started by specifying the --async-connect-arbiter option in the mc\_ctl command.
- After Mirroring Controller is started, automatic switch/disconnection can be enabled or disabled using the enable-failover or disable-failover mode of the mc\_ctl command.
- 2. Obtain the backup.

Use the pgx\_dmpall command to collect the backup.

### 2.5 Setting Up the Standby Server

This section explains how to set up the standby server.

### 2.5.1 Setting Up Database Multiplexing Mode on the Standby Server

This section explains how to set up database multiplexing mode on the standby server.

In database multiplexing, the files that are required for operations are managed in the Mirroring Controller management directory.

There is one Mirroring Controller management directory for each instance.



- Do not place the Mirroring Controller management directory in a directory managed by FUJITSU Enterprise Postgres, otherwise it may be deleted by mistake with the directories managed by FUJITSU Enterprise Postgres, and an old version of files may be restored.

- When creating a standby server for a large database, stop job system operations, specify a large value for the wal\_keep\_segments parameter, or use replication slots.

This is because WALs generated after the standby server is built using the pg\_basebackup command, but before it is started, need to be retained. However, the number of WAL segments that can be retained is constrained by the wal\_keep\_segments parameter. Additionally, setting the wal\_keep\_segments parameter requires consideration regarding stabilization of the database multiplexing mode (refer to "2.11.1 Tuning to Stabilize the Database Multiplexing Mode" for details).



- Refer to "Preparing Directories for Resource Deployment" in the Installation and Setup Guide for Server for details on the directories that are managed by FUJITSU Enterprise Postgres.

......

- $\ Refer to \ "pg\_base backup" in \ "Reference" in the \ PostgreSQL \ Documentation for information on the \ pg\_base backup \ command.$
- Refer to "mc\_ctl" in Reference for information on the command.
- Refer to "Appendix A Parameters" for details on each parameter to be edited for the setup.
- Refer to "Replication Slots" in the PostgreSQL Documentation for information on replication slots.

Perform the following procedure:

- 1. Log in to the standby server.
- 2. Create the Mirroring Controller management directory that will store the files required by database multiplexing.

Use ASCII characters in the Mirroring Controller management directory.

Additionally, grant "Write" permission to the instance administrator user for the Mirroring Controller management directory.

3. Copy, and then deploy, the network.conf file of the primary server.

Copy the network.conf file that was defined in the primary server setup, and deploy it to the Mirroring Controller management directory of the standby server.

Register the port number of the standby server that was specified in the network.conf file in the services file, because there are programs, such as WebAdmin, that search an available port number using the services file.

Register any name as the service name.

4. Change the access permissions for the network.conf file.

Set read and write permissions for the instance administrator user only. If users other than the instance administrator user are granted access permissions, the mc\_ctl command will not work. Accordingly, users other than the instance administrator user are prevented from operating Mirroring Controller.

#### Example)

The following is an execution example, in which the instance administrator user is granted full access permissions as the owner when the operating system user name of the instance administrator user is "fsepuser". The following procedure applies when the user is logged in to the Windows server as "fsepuser".

```
> takeown /f network.conf
> icacls network.conf /reset
> icacls network.conf /inheritance:r
> icacls network.conf /grant fsepuser:F
```

5. Copy, and then deploy, the *serverIdentifier*.conf file of the primary server.

Copy the *serverIdentifier*.conf file that was defined in the primary server setup, and deploy it to the Mirroring Controller management directory of the standby server.



If the primary server and standby server are to be built within the same server, change the following parameters in the serverIdentifier.conf file, ensuring that the names are not duplicated with those on the primary server:

- db\_instance\_service\_name (registered service name of the FUJITSU Enterprise Postgres instance)
- mc\_service\_name (registered service name of Mirroring Controller)

- event\_source (event source name)

#### 6. Change the access permissions for the *serverIdentifier*.conf file.

Set read and write permissions for the instance administrator user only. If users other than the instance administrator user are granted access permissions, the mc\_ctl command will not work.

#### Example)

The following is an execution example, in which the instance administrator user is granted full access permissions as the owner when the operating system user name of the instance administrator user is "fsepuser". The following procedure applies when the user is logged in to the Windows server as "fsepuser":

```
> takeown /f serverIdentifier.conf
> icacls serverIdentifier.conf /reset
> icacls serverIdentifier.conf /inheritance:r
> icacls serverIdentifier.conf /grant fsepuser:F
```

### 7. Configure the Windows firewall.

If the Windows firewall feature is to be enabled, you should enable the port number of Mirroring Controller that you specified in the network definition file in step 3. Refer to "E.2 Windows Firewall Settings" for details.

8. Register Mirroring Controller to the Windows service.

Execute the mc\_ctl command in the register mode.

For the -P option of the mc\_ctl command, specify the password of the operating system user who executes the command.

Example)

> mc\_ctl register -M D:\mcdir\inst1 -P \*\*\*\*\*\*\*



When specifying the password in the -P option of the mc\_ctl command, for security reasons, you should be careful not to allow other users to access it.

......



You can use the mc\_ctl command with the -S option to specify automatic start and stop of Mirroring Controller. Refer to "2.12 Setting Automatic Start and Stop of Mirroring Controller and Multiplexed Instances" for details.

Using the service name specified in the mc\_service\_name parameter of *serverIdentifier*.conf in step 5, Mirroring Controller is registered to the Windows service.

......

You can execute the sc qc command to check the registration status.

### 2.5.2 Creating, Setting, and Registering the Standby Server Instance

This section explains how to create, set, and register the standby server instance.



Mirroring Controller supports instances that are registered in the Windows service.



- Refer to "Appendix A Parameters" for details on each parameter.
- Refer to "mc\_ctl" in Reference for information on the command.

Perform the following procedure:

1. Prepare for setup.

Refer to "Preparations for Setup" in the Installation and Setup Guide for Server for information on the preparatory tasks to be performed before creating an instance on the standby server.



If the primary server and standby server are to be built within the same server, perform preparation to ensure that the event source names of FUJITSU Enterprise Postgres are not duplicated with that of the primary server.

2. When using transparent data encryption, configure the encryption settings for the storage data.

Deploy a copy of the keystore file of the primary server on the standby server.

Refer to "Database Multiplexing Mode" in the Operation Guide for details.

3. Execute the pg\_basebackup command to create a copy of the primary server instance on the standby server.

Example)

> pg\_basebackup -D D:\database\inst1 -X fetch --waldir=E:\transaction\inst1 --progress --verbose -R --dbname="application\_name=standbyServerName" -h primaryServerIpAddress -p primaryServerPortNumber



- Use the pg\_basebackup command with the -R option to create a standby.signal file. If you do not create the standby.signal file, the Mirroring Controller cannot be started as a standby server.
- If using a method that requires password authentication for connections to the primary server, you will need to ensure that authentication is performed automatically. If the -R option is specified for the pg\_basebackup command and the password parameter is specified for the --dbname option, the pg\_basebackup command will set the password in the primary\_conninfo parameter in postgresql.auto.conf file, enabling connections to be performed automatically.

If a password is not set in the primary\_conninfo parameter in postgresql.auto.conf file, it will be necessary to create a password file (%APPDATA%postgresql\pgpass.conf), and then specify a password for the replication database.

- The primary\_conninfo parameter should not be set in the postgresql.conf file, but only in the postgresql.auto.conf file using the pg\_basebackup command.
- When executing the pg\_basebackup command, consider the following for collection of transaction logs.
  - When "fetch" is specified for the -X option of the command
    - Transaction logs are collected at the end of the backup, so it is necessary to ensure that transaction logs that occur during backup are not deleted from the primary server. Therefore, allow for a sufficient value for the wal\_keep\_segments parameter in postgresql.conf.
  - When the -X option is omitted or "stream" is specified for the -X option of the command

Transaction logs are streamed, so when Mirroring Controller is running on the primary server, the connection is changed to a synchronous standby server on detection of a streaming replication connection using this command. Therefore, if a job has started on the primary server, the primary server will be impacted, therefore execute this command after stopping only the Mirroring Controller process on the primary server.



#### See

Refer to "Hot Standby" in the PostgreSQL Documentation for information on the standby signal file.

4. Set the parameters shown in the table below in the postgresql.conf file.

Table 2.5 Parameters

Parameter	Content specified	Remarks
synchronous_standby_names	'primaryServerName'	Required after switching the primary server and then changing the original primary server to the new standby server.  Enclose the name in single quotation marks (').  Do not change this parameter while Mirroring Controller is running.  Do not specify multiple names to this parameter as the Mirroring Controller can manage only one standby server.

5. Register an instance to the Windows service.

Refer to "Creating an Instance" in the Installation and Setup Guide for Server for information on how to register an instance to the Windows service. Note that you should execute the pg\_ctl command with the following specified for the register mode to enable Mirroring Controller to start and stop an instance:

- For the service name of the -N option, specify the name set for the db\_instance\_service\_name parameter in the server definition file

- Specify "demand" for the -S option, so that the service does not start automatically on startup of the system



- Do not configure the Windows service of a multiplexed instance to perform automatic start, as it is started by Mirroring Controller.
- If the primary server and standby server are to be built within the same server, ensure that the registered service name of the FUJITSU Enterprise Postgres instance is not duplicated with that of the primary server.

### 2.5.3 Starting Mirroring Controller on the Standby Server

This section explains how to start Mirroring Controller on the standby server.

When the arbitration server is used for automatic degradation, start the Mirroring Controller arbitration process on the arbitration server in advance.

 $1. \ \ After ensuring that the Mirroring Controller process of the primary server has started, start Mirroring Controller on the standby server.$ 

#### Enabling automatic switch/disconnection

As the instance administrator user, execute the mc\_ctl command in start mode with the -f option specified. This action enables automatic switch/disconnection.

If you start Mirroring Controller and the instance without specifying the -f option, automatic switch/disconnection will not be enabled. To enable both, start Mirroring Controller and then execute the mc\_ctl command in enable-failover mode or restart Mirroring Controller with the -f option specified.

Example)

> mc\_ctl start -M D:\mcdir\inst1

### Disabling automatic switch/disconnection

As the instance administrator user, execute the mc\_ctl command in start mode with the -F option specified.

Example)

> mc\_ctl start -M D:\mcdir\instl -F

2. Check the status of the Mirroring Controller process.

As the instance administrator user, execute the mc\_ctl command in status mode. Ensure that "mirroring status" is switchable.

Example)

> mc\_ctl status -M D:\mcdir\inst1



- When the arbitration server is used for automatic degradation, the time required for the database server to connect to the arbitration server is added on. Therefore, Mirroring Controller startup may take longer than when the arbitration server is not used.
- If the parameter for heartbeat monitoring of operating systems or servers set by the abitration server is greater than parameter for heartbeat monitoring of operating systems and servers of the Mirroring Controller, the Mirroring Controller may fail to start. In this case, check the contents of the message notification and review the parameters for heartbeat monitoring of operating systems or servers for the arbitration server or Mirroring Controller.
- If the heartbeat\_error\_action parameter in *serverIdentifier*.conf is set to "message", even if automatic switch/disconnection is enabled and Mirroring Controller is started, only message output is performed when a heartbeat abnormality is detected during heartbeat monitoring of operating systems and servers switch/disconnection is not performed.

- Mirroring Controller startup usually fails if the standby server is mistakenly started as the primary server or if the old primary server is not recovered after the switch and is then mistakenly started as the primary server. However, if the admin network is disconnected, then startup does not fail, and both servers may become primary servers. Therefore, ensure that the admin network is connected before starting Mirroring Controller.



- The mc\_ctl command fails if the Mirroring Controller arbitration process has not been started on the arbitration server when the arbitration server is used for automatic degradation. However, if the Mirroring Controller arbitration process cannot be started in advance, it can be started by specifying the --async-connect-arbitration in the mc\_ctl command.

•••••

- After Mirroring Controller is started, automatic switch/disconnection can be enabled or disabled using the enable-failover or disable-failover mode of the mc\_ctl command.

### 2.6 Creating a User Exit for a Database Server

This section explains how to create a user exit for a database server.

The user command types explained below can be used as user exits. These commands are called by Mirroring Controller management processes.

The user can create user exits as required.

Specify the user commands that were created for the parameters in the server configuration file of the database server. Refer to "A.4.1 Server Configuration File for the Database Servers" for information on these parameters.

### **User command types**

- Fencing command

This user command performs fencing if Mirroring Controller performs arbitration processing and determines that a database server is unstable.

- Arbitration command

This user command performs arbitration processing in lieu of the arbitration server when there is no arbitration server.

- State transition commands

These user commands are called when Mirroring Controller performs state transition of a database server. It includes the following types:

- Post-switch command

This user command is called after a promotion from standby server to primary server.

- Pre-detach command

This user command is called before the standby server is disconnected from a cluster system.

If the pre-detach command is specified on both the primary server and standby server, it is called first on the standby server and then on the primary server.

If the settings are configured to forcibly stop the instance on the standby server when the standby server is disconnected, the predetach command is called on the standby server and then the instance on the standby server is stopped.

- Post-attach command

This user command is called after the standby server has been attached to a cluster system.

If the post-attach command is specified on both the primary server and standby server, it is called first on the primary server and then on the standby server.



When the arbitration server is used for automatic degradation and the requirements can be satisfied using the fencing command on the arbitration server only, the fencing command on the database server is not required. In addition, if the requirements can be satisfied using the fencing command on the database server only, create a fencing command on the arbitration server for termination processing only (without implementation).

Table 2.6 Availability of user commands, and database server calling the command

User command	Operation when a heartbeat abnormality is detected using operating system or server heartbeat monitoring				Database server calling the command	
	Message output	Automatic degradation by arbitration server	Automatic degradation by arbitration command	Unconditio nal automatic degradatio n	Primary server	Standby server
Fencing command	Y (*1)	Y (*2)	R	N	Y	Y
Arbitration command	N	N	R	N	Y	Y
Post-switch command	Y	Y	Y	Y	Y	N
Pre-detach command	Y	Y	Y	Y	Y	Y (*3)
Post-attach command	Y	Y	Y	Y	Y	Y (*3)

- R: Required
- Y: Can be used
- N: Cannot be used
- \*1: Called only when the mc\_ctl command is used to execute forced switching or forced disconnection.
- \*2: Creation of a fencing command on a database server is optional, but it must be created on the arbitration server.
- \*3: If message output or unconditional automatic degradation is selected, this command is called only from the primary server.



Refer to "Appendix C User Commands" for information on the interface for each user command.

### 2.7 Confirming the Streaming Replication Status

Before performing the setup of the database multiplexing mode, ensure that the prerequisite streaming replication feature has been set up correctly.

Perform the following procedure:

1. On the primary server, ensure that single-row searches can be performed using the pg\_stat\_replication statistics view.

An example output of the psql command is shown below.

Example)

```
client_hostname |
client_port
               55098
backend_start
             2020-03-23 11:17:49.628793+09
backend_xmin
               streaming
sent_lsn
               0/3000060
write_lsn
               0/3000060
               0/3000060
flush_lsn
               0/3000060
replay_lsn
write_lag
flush_lag
replay_lag
sync_priority
                | 1
sync_state
                 sync
reply_time
                2020-03-23 11:23:27.703366+09
```

#### 2. Confirm the search results of step 1.

Ensure that the connection established with the intended standby server is in synchronous mode.

Table 2.7 Items to be checked

. date 21. Herris to be directed.		
Item	Required value	
application_name	Value specified for synchronous_standby_names parameter in the postgresql.conf file of the primary server.	
client_addr	IP address of the standby server.	
state	"streaming".	
sync_state	"sync".	



### See

- Refer to "The Statistics Collector" in "Server Administration" in the PostgreSQL Documentation for information on the pg\_stat\_replication statistics view.
- Note that the pg\_stat\_replication statistics view may change in the future.

### 2.8 Checking the Connection Status

This section explains how to check the connection status from a database server or the arbitration server.

### 2.8.1 Checking the Connection Status on a Database Server

This section explains how to use a database server to check the connection status of the Mirroring Controller arbitration process and the Mirroring Controller process on the primary server and standby server.

Perform the following procedure:

1. On the primary server and standby server, execute the mc\_ctl command in status mode with the --arbiter option specified.

```
Example)
```

The mc\_ctl command is executed with the --arbiter option specified, and the status is output.

```
> mc_ctl status --arbiter -M D:\mcdir\inst1

arbiter_id host status
------
arbiter 192.0.3.120 online
```

2. On the primary server and standby server, check the result displayed by executing the mc\_ctl command in status mode in step 1.

Items to be checked

Check that the output status" is "online".



Refer to the Reference for information on the mc\_ctl command.

### 2.8.2 Checking the Connection Status on the Arbitration Server

This section explains how to use the arbitration server to check the connection status of the Mirroring Controller arbitration process and the Mirroring Controller process on the primary server and standby server.

Perform the following procedure:

1. Execute the mc arb command in status mode on the arbitration server.

The example below executes the mc\_arb command, and shows the status.

#### Linux

Example)

\$ mc_arb st	atus -M /mcarb_c	dir/arbiter1
server_id	host	status
server1 server2	192.0.3.100 192.0.3.110	online online

### Windows

Example)

> mc_arb st	atus -M D:\mcarl	o_dir\arbiter1
server_id	host	status
server1 server2	192.0.3.100 192.0.3.110	online online

2. Check the result displayed by executing the  $mc_{arb}$  command in step 1.

Items to be checked

Check that the output status is "online" on both lines.



Refer to the Reference for information on the mc\_arb command.

### 2.9 Creating Applications

This section explains how to create applications using database multiplexing, and points that should be noted when you create the applications.

### 2.9.1 Application Connection Server Settings

If database multiplexing is used and a failover occurs, it will be necessary to switch the application connection server. Accordingly, use the application connection switch feature to create applications.



Refer to "Application Connection Switch Feature" in the Application Development Guide for details.

### 2.10 Checking the Behavior

To check if the environment setup was performed correctly, start the application and then check the behavior of the switch and rebuild.

### 2.11 Tuning

This section explains how to tune database multiplexing mode.

### 2.11.1 Tuning to Stabilize the Database Multiplexing Mode

When large amounts of data are updated, the write-to load for the database will become great, and the multiplexing state may become unstable.

Accordingly, by editing the parameters below in the postgresql.conf file, a stable multiplexing state can be maintained. Refer to "Estimating Transaction Log Space Requirements" in the Installation and Setup Guide for Server for information on transaction log space requirements.

#### Table 2.8 Parameters

Parameter	Content
wal_keep_segments	Refer to "2.4.2 Creating, Setting, and Registering the Primary Server Instance" for details.
max_wal_size	The transaction log is written out according to the checkpoint trigger.
	If a transaction log with the capacity of the value specified in this parameter is generated, the checkpoint will be executed.
	If a large value is specified in this parameter, the time required for crash recovery will increase.
	If a small value is specified in this parameter, many checkpoints will be generated, which will affect the performance of the applications that connect to the primary server.

### 2.11.2 Tuning to Stabilize Queries on the Standby Server

Queries made using reference jobs on the standby server may be canceled by jobs executed on the primary server.

To reduce the possibility of a job being canceled, specify as large a value as possible for the max\_standby\_archive\_delay parameter in the postgresql.conf file.



#### See

- Refer to "Handling Query Conflicts" in the PostgreSQL Documentation for details.
- Refer to "Standby Servers" in the PostgreSQL Documentation for details on the max\_standby\_archive\_delay parameter.

# 2.11.3 Tuning to Stabilize Queries on the Standby Server (when Performing Frequent Updates on the Primary Server)

If jobs are updated on the primary server regularly and frequently, it will be easy for the query made by the reference job on the standby server to be canceled. In this case, edit one of the postgresql.conf file parameters shown in the table below.

#### Table 2.9 Parameters

Parameter	Description
hot_standby_feedback	When "on" is set, the deletion (vacuum) of the data area that was deleted or updated on the primary server is suppressed.  Accordingly, the query on the standby server will not be canceled. (*1)
vacuum_defer_cleanup_age	The deletion (vacuum) of the data area that was deleted or updated on the primary server is delayed until the specified number of transactions is processed. Accordingly, the probability that the query on the standby server will be canceled decreases.

<sup>\*1:</sup> Because the vacuum is delayed, the data storage destination disk space of the primary server comes under pressure.

Additionally, if there is conflict between accesses and queries executed on the standby server, transaction logs indicating this conflict will be transferred.

Accordingly, specify as large a value as possible for the max\_standby\_archive\_delay parameter so that access conflicts do not occur.



#### See

- Refer to "Standby Servers" in the PostgreSQL Documentation for details on the hot\_standby\_feedback parameter.
- Refer to "Master Server" in the PostgreSQL Documentation for details on the vacuum\_defer\_cleanup\_age parameter.

### 2.11.4 Tuning for Optimization of Degradation Using Abnormality Monitoring

Mirroring Controller uses a monitoring method that outputs an error if the timeout or number of retries is exceeded when accessing resources targeted for monitoring. Setting inappropriate values in these settings may lead to misdetection or a delay in automatic degradation, so you must design these values appropriately.

For example, the following type of issue occurs if the tuning related to abnormality monitoring is not performed appropriately.

- If the timeout is too short
  - Results in redundant degradation and availability falls.
- If the timeout is too long

It takes longer for automatic degradation to be performed even when an error affecting operational continuity occurs, potentially causing downtime.

You can optimize degrading operation by editing the values for the parameters in the server configuration file described below in accordance with the system. Refer to "A.4 Server Configuration File" for information on how to edit these parameters.

### 2.11.4.1 Tuning for Abnormality Monitoring of the Operating System or Server

Tuning for abnormal monitoring of the operating system or server depends on the operation when heartbeat abnormality is detected by the heartbeat monitoring of operating systems or servers.



#### See

Refer to "1.1.1 Monitoring Using Database Multiplexing Mode" for the operation when heartbeat abnormality is detected in the the heartbeat monitoring of operating systems or servers.

# 2.11.4.1.1 Tuning Abnormality Monitoring for Operations that Use an Arbitration Server for Automatic Degeneration

In an operation that use an arbitration server for automatic degeneration, the database server is periodically monitored for abnormalities so that the Mirroring Controller arbitration process can immediately respond to an arbitration request from the Mirroring Controller. The

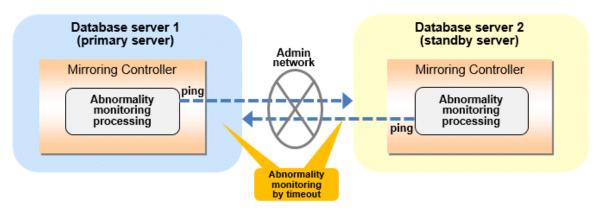
automatic degradation using the arbitration server can optimize the time from error detection to automatic degradation of the operating systems or servers by editing the following parameters.

- Parameters for the abnormality monitoring of the operating system or server in the server configuration file of the database server
- Parameters for the abnormality monitoring of the operating system or server in the arbitration configuration file
- Parameters for the arbitration processing and fencing

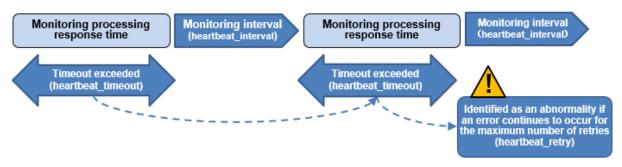
# Parameters for the abnormality monitoring of the operating system or server in the server configuration file of the database server

Table 2.10 Parameters for the abnormality monitoring of the operating system or server in the server configuration file of the database server

Parameter	Description
Abnormality monitoring interval (heartbeat_interval)	Mirroring Controller is configured so that abnormality monitoring does not place a load on the system. This parameter does not normally need to be set. (The default is 800 milliseconds.)
Abnormality monitoring timeout (heartbeat_timeout)	Take into account the time during which a load is placed continuously on the server or admin network performance. For example, it is envisaged that this parameter will be used in situations such as when performing high-load batch jobs or when a large number of online jobs occur continuously and concurrently. (The default is 1 second.)
Abnormality monitoring retries (heartbeat_retry)	This parameter can be set when needing a safety value for situations in which the value specified for heartbeat_timeout is exceeded, for example, when using systems with fluctuating loads, however, this parameter does not normally need to be set. (The default is 2 times.)



### Flow of abnormality monitoring by timeout



The expression for calculating the time required to detect an abnormality by Mirroring Controller is shown below.

```
Abnormality detection time of Mirroring Controller = ( heartbeat_timeout(seconds) + heartbeat_interval(milliseconds) / 1000 ) x ( heartbeat_retry(number of times) + 1)
```

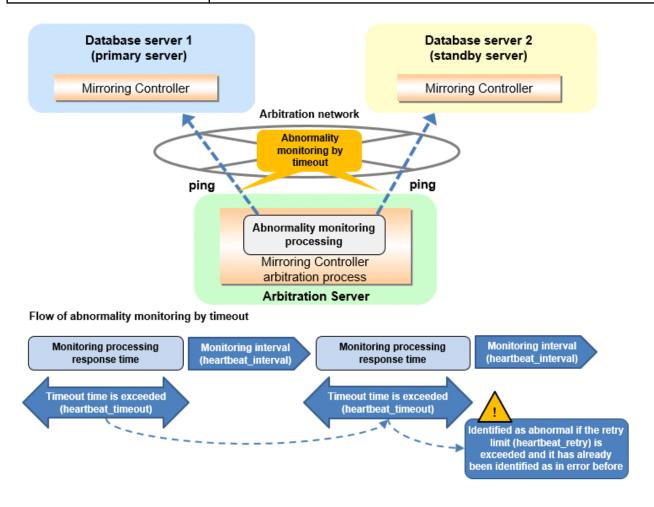
The abnormality detection time when the default value is used is shown below.

```
Abnormality detection time of Mirroring Controller = ( 1 + 800 / 1000 ) x ( 2 + 1 ) = 5.4(seconds)
```

# Parameters for the abnormality monitoring of the operating system or server in the arbitration configuration file

Table 2.11 Parameters for the abnormality monitoring of the operating system or server in the arbitration configuration file

Parameter	Description
Abnormality monitoring interval (heartbeat_interval)	Mirroring Controller arbitration process is configured so that abnormality monitoring does not place a load on the system. This parameter does not normally need to be set. (The default is the value set in heartbeat_interval in the server configuration file of the database server.) (milliseconds).
Abnormality monitoring timeout (heartbeat_timeout)	Take into account the time during which a load is placed continuously on the server and arbitration network capabilities. (The default is the value set in heartbeat_timeout in the server configuration file of the database server.) (seconds).
Abnormality monitoring retries (heartbeat_retry)	This parameter can be set when needing a safety value for situations in which the value specified for heartbeat_timeout is exceeded, for example, when using systems with fluctuating loads, however, this parameter does not normally need to be set. (The default is the value set in heartbeat_retry in the server configuration file of the database server.) (number of times)



The expression for calculating the time required to detect an abnormality by Mirroring Controller arbitration process is shown below.

Abnormality detection time of Mirroring Controller arbitration process = ( heartbeat\_timeout(seconds) + heartbeat\_interval(milliseconds) / 1000 ) x ( heartbeat\_retry(number of times) + 1)

The abnormality detection time when the default value is used is shown below.

Abnormality detection time of Mirroring Controller arbitration process = ( 1 + 800 / 1000 ) x ( 2 + 1 ) = 5.4(seconds)



The abnormality detection time of the operation for automatic degradation using the arbitration server can be calculated as follows.

Abnormality detection time = Max(Abnormality detection time by Mirroring Controller, Abnormality detection time by Mirroring Controller arbitration process)



If the heartbeat\_interval is set in the arbitration configuration file, the relationship between the parameter for operating system or server abnormality monitoring specified in the server configuration file of the database server file and the heartbeat\_interval of the arbitration configuration file must satisfy the following relational expression.

.....

Heartbeat\_interval in the arbitration configuration file (milliseconds) / 1000 ) <
( heartbeat\_timeout(seconds) + heartbeat\_interval(milliseconds) / 1000 ) \* heartbeat\_retry(number of times) + heartbeat\_timeout(seconds)</pre>

#### Parameters for the arbitration processing and fencing

Table 2.12 Parameters for the arbitration processing and fencing

Parameter	Description
Arbitration processing timeout (arbitration_timeout in the server configuration file of the database server)	Take into account the time to perform arbitration processing on the Mirroring Controller arbitration process. The value must be greater than or equal to abnormality detection time of Mirroring Controller arbitration process + fencing_command_timeout in the arbitration configuration file (seconds).
Fencing timeout (fencing_command_timeout in the arbitration configuration file)	Take into account the time to execute the fencing command (seconds).

### Flow from the abnormality detection to the automatic degeneracy

When performing automatic degradation using the arbitration server, the flow from the abnormality detection in the operating system or server to the occurrence of automatic degeneracy and the parameters is shown below.

Flow from the abnormality detection to the automatic degeneracy	Description	Parameter	
(1) Abnormality detection	Mirroring Controller detect the database server operating system or server errors.	Parameters for the abnormality monitoring of the operating system or server in the server configuration file of the database server	
(2) Arbitration request	Mirroring Controller that detect the operating system or server error asks the Arbitration Server	-	arbitration_timeout in the server configuration

Flow from the abnormality detection to the automatic degeneracy	Description	Parameter	
	to check the status of the other server's operating system or server.		file of the database server
(3) Arbitration processing	The Mirroring Controller arbitration process checks the status of the other server's operating system or server. However, if the result of the operating system or server abnormality monitoring by the arbitration server has been determined before the arbitration request from the Mirroring Controller of the database server, this process is not performed.	Parameters for the abnormality monitoring of the operating system or server in the arbitration configuration file	
(4) Fencing	If the Mirroring Controller arbitration process determines that the other server is an abnomaly of the operating system or server, it fences the other server and isolates it from the cluster system. If the Mirroring Controller arbitration process determines that the operating system or server status is normal, this process and the (6) are not performed.	fencing_command_ti meout in the arbitration configuration file	
(5) Return of the arbitration results	Returns the results of the arbitration to the Mirroring Controller of the database server that requested the arbitration.	-	
(6) Automatic degradation	The automatic degradation is performed. If fencing fails in (4), this procedure is not performed.	-	

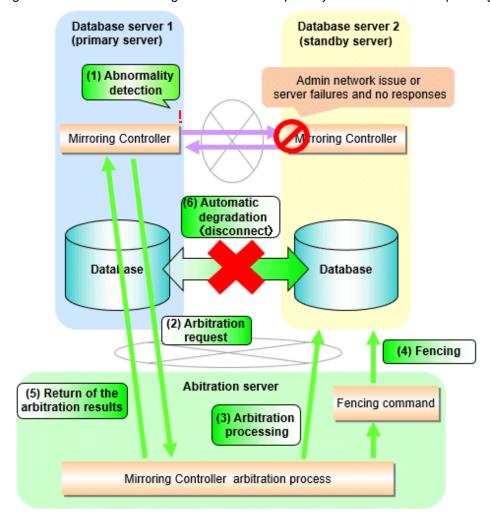
<sup>-:</sup> No associated parameters



If the fencing\_command parameter is specified in the server configuration file of the database server, the fencing command is invoked on the database server if fencing is successful on the arbitration server. In that case, add the value of the fencing\_command\_timeout parameter in the server configuration file of the database server to the estimate.

.....

Figure 2.1 When the Mirroring Controller on the primary server detects an operating system or server error



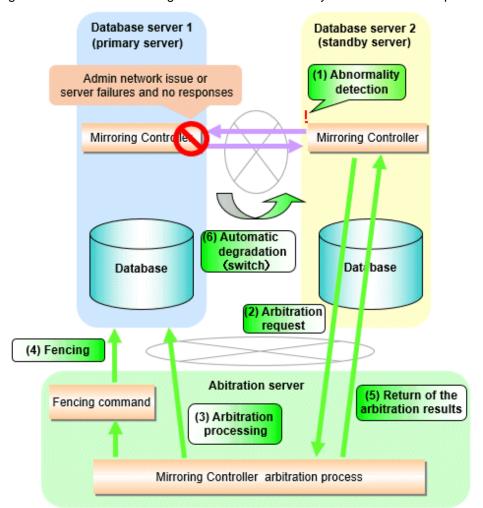


Figure 2.2 When the Mirroring Controller on the standby server detects an operating system or server error

# 2.11.4.1.2 Tuning Abnormality Monitoring for Operations that Perform Automatic Degeneration by Calling a User Exit that Determines Degeneration

In an operation that perform automatic degeneration by calling a user exit that determines degeneration, you can optimize the time from operating system or server abnormality detection to automatic degradation by editing the operating system or server abnormality monitoring parameters and parameters related to arbitration processing and fencing in the server configuration file of the database server. Refer to "Parameters for the abnormality monitoring of the operating system or server in the server configuration file of the database server" for information on the operating system or server abnormality monitoring parameters in the server configuration file of the database server.

Table 2.13 Parameters for the arbitration processing and fencing

Parameter	Description
Arbitration processing timeout (arbitration_command_timeout)	Take into account the time to execute the arbitration command(seconds).
Fencing timeout (fencing_command_timeout)	Take into account the time to execute the fencing command (seconds).

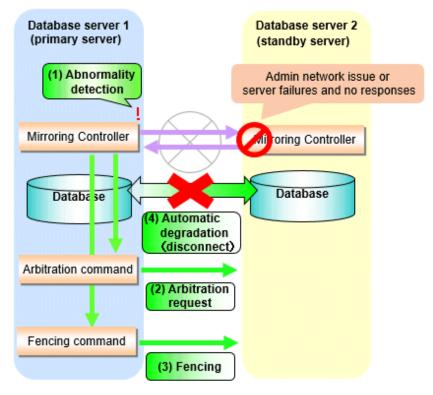
### Flow from the abnormality detection to the automatic degeneracy

When performing automatic degradation by calling a user exit that determines degeneration, the flow from the abnormality detection in the operating system or server to the occurrence of automatic degeneracy and the parameters is shown below.

Flow from the abnormality detection to the automatic degeneracy	Description	Parameter
(1) Abnormality detection	Mirroring Controller detect the database server operating system or server errors.	Parameters for the abnormality monitoring of the operating system or server in the server configuration file of the database server
(2) Arbitration processing	An arbitration command is executed to check the status of the other server's operating system or server.	arbitration_command_timeout in the server configuration file of the database server
(3) Fencing	If the operating system or server status of the other server is abnomal in (2), it fences the other server and isolates it from the cluster system. If the operating system or server status of the other server is nomal in (2), this process and (4) are not executed.	fencing_command_timeout in the server configuration file of the database server
(4) Automatic degradation	The automatic degradation is performed. If fencing fails in (3), this procedure is not performed.	-

<sup>-:</sup> No associated parameters

Figure 2.3 When the Mirroring Controller on the primary server detects an operating system or server error



Database server 1 Database server 2 (standby server) (primary server) (1) Abnormality Admin network issue or detection server failures and no responses Mirroring Contro Mirroring Controller Database Da tabase (4) Automatic degradation (switch) Arbitration command (2) Arbitration request Fencing command (3) Fencing

Figure 2.4 When the Mirroring Controller on the standby server detects an operating system or server error

### 2.11.4.1.3 Tuning Abnormality Monitoring for Operations that Notify Messages

In an operation that notify messages, you can optimize the abnormality detection time by editing the operating system or server abnormality monitoring parameters in the server configuration file of the database server. Refer to "Parameters for the abnormality monitoring of the operating system or server in the server configuration file of the database server" for information on the operating system or server abnormality monitoring parameters in the server configuration file of the database server. In addition, when the Mirroring Controller detects an error, it does not perform the arbitration processing, fencing, or automatic degradation, but only notification messages is performed.

## 2.11.4.1.4 Tuning Abnormality Monitoring for Operations that Perform Automatic Degenerate Unconditionally due to Heartbeat Abnormality

In an operation that perform automatic degenerate unconditionally due to heartbeat abnormality, you can optimize the time from operating system or server abnormality detection to automatic degradation by editing the operating system or server abnormality monitoring parameters in the server configuration file of the database server. Refer to "Parameters for the abnormality monitoring of the operating system or server in the server configuration file of the database server" for information on the operating system or server abnormality monitoring parameters in the server configuration file of the database server. In addition, when the Mirroring Controller detects an error, it does not perform the arbitration processing, fencing, or automatic degradation, but only automatic degenerate unconditionally is performed.



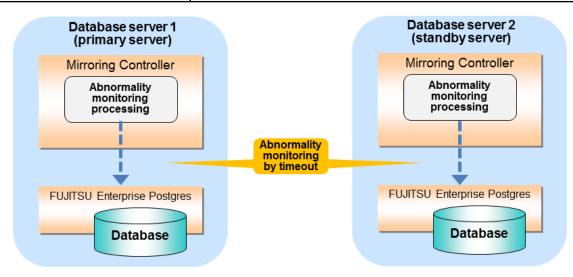
Refer to "Appendix D Notes on Performing Automatic Degradation Immediately after a Heartbeat Abnormality" for notes on the operation that perform automatic degenerate unconditionally due to heartbeat abnormality.

### 2.11.4.2 Tuning for Abnormality Monitoring of Darabase Processes

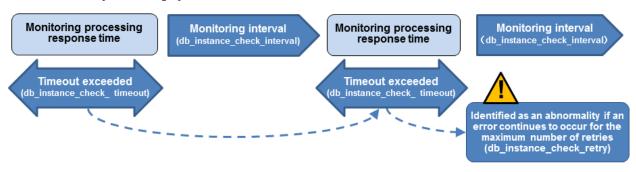
In an abnormality monitoring of database pocesses, you can optimize by editing the following parameters in the server configuration file of the database server.

Table 2.14 Parameters for abnormality monitoring of database processes

Parameter	Description
Abnormality monitoring interval (db_instance_check_interval)	Abnormality monitoring by Mirroring Controller is set so as not to place load on the system, but normally it does not need to be set. (The default is the value set in heartbeat_interval.) (milliseconds)
Timeout for abnormality monitoring of database processes (db_instance_check_timeout)	Take into account the time during which a load is placed continuously on the database. For example, it is envisaged that this parameter will be used in situations such as when performing high-load batch jobs or when a large number of online jobs occur continuously and concurrently. (The default is the value set in heartbeat_timeout.) (seconds)
Abnormality monitoring retries (db_instance_check_retry)	This parameter can be set when needing a safety value for situations in which the value specified for db_instance_check_timeout is exceeded, for example, when using systems with fluctuating loads, however, this parameter does not normally need to be set. (The default is the value set in heartbeat_retry.) (number of times)



### Flow of abnormality monitoring by timeout



The expression for calculating the time required to detect an abnormality is shown below.

```
Abnormality detection time = ( db_instance_check_timeout(seconds) + db_instance_check_interval(milliseconds) / 1000 ) x ( db_instance_check_retry(number of times) + 1 )
```

The abnormality detection time when the default value is used is shown below.

```
Abnormality detection time = ( 1 + 800 / 1000 ) x ( 2 + 1 ) = 5.4(seconds)
```



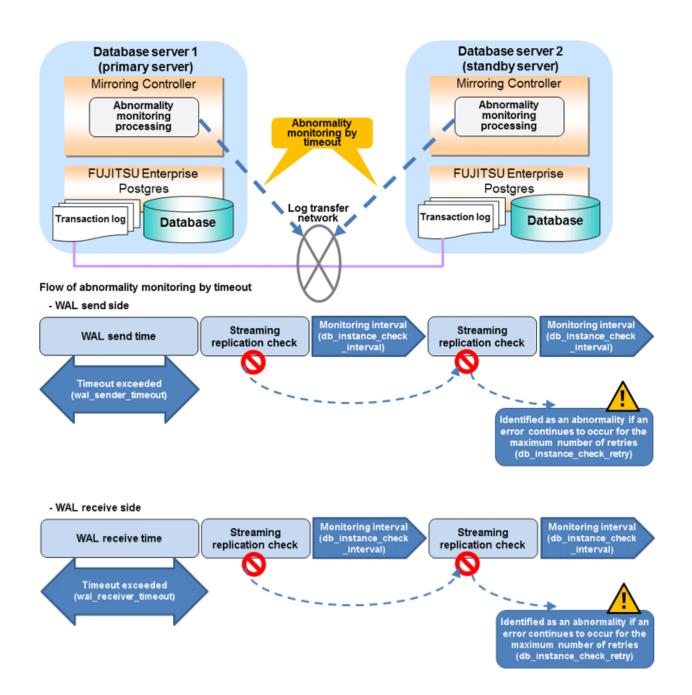
- If the db\_instance\_timeout\_action parameter in *serverIdentifier*.conf is set to "message", and the db\_instance\_check\_timeout parameter is set to a short value, a crash of the database process will be detected as "no response", and it may take time for automatic degradation to occur. Therefore, specify an appropriate timeout for db\_instance\_check\_timeout.
- If a high load on the database and an event that prevents connection to an instance occur at the same time, it is judged as abnormal without retrying monitoring.

## 2.11.4.3 Tuning for Abnormality Monitoring of Streaming Replication

In an abnormality monitoring of streaming replication, you can optimize by editing the following parameters in the server configuration file of the database server.

Table 2.15 Parameters for abnormality monitoring of streaming replication

Parameter	Description	
Abnormality monitoring interval (db_instance_check_interval)	Abnormality monitoring by Mirroring Controller is set so as not to place load on the system, but normally it does not need to be set. (The default is the value set in heartbeat_interval.) (milliseconds)	
Abnormality monitoring retries (db_instance_check_retry)	This parameter can be set when needing a safety value, such as when it is anticipated that a temporary log transfer LAN error may occur, but it does not normally need to be set. (The default is the value set in heartbeat_retry.) (number of times)	
Timeout for abnormality monitoring of streaming replication (wal_sender_timeout and wal_receiver_timeout in postgresql.conf)	Take into account the capacity and load of the log transfer network and the time during which a load is placed continuously on the database. For example, if there is a succession of data update jobs that generate a high WAL volume, you must configure the settings to avoid misdetection. (The default is 60 seconds.)	



The expression for calculating the time required to detect an abnormality is shown below.

The abnormality detection time when the default value is used is shown below.

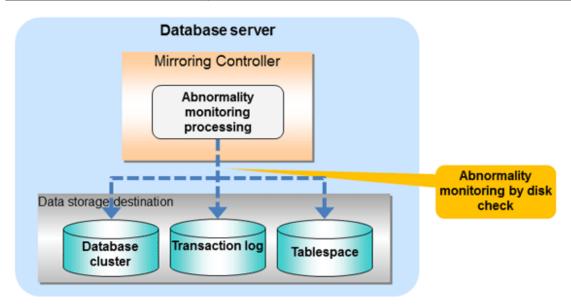
```
Abnormality detection time = 60 + (800 / 1000 \times (2 + 1))
= 62.4(seconds)
```

### 2.11.4.4 Tuning for Disk Abnormality Monitoring

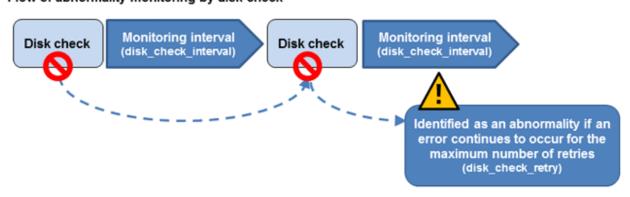
In an abnormality monitoring of the disk, you can optimize by editing the following parameters in the server configuration file of the database server.

Table 2.16 Parameters for disk abnormality monitoring

Parameter	Description
Abnormality monitoring interval (disk_check_interval)	Abnormality monitoring by Mirroring Controller is set so as not to place load on the system, but normally it does not need to be set. (The default is the value set in heartbeat_interval.) (milliseconds)
Abnormality monitoring retries (disk_check_retry)	This parameter can be set when needing a safety value, such as when it is anticipated that a temporary disk input/output error may occur, but normally it does not need to be set. (The default is the value set in heartbeat_retry.) (number of times)



### Flow of abnormality monitoring by disk check



The expression for calculating the time required to detect an abnormality is shown below.

```
Abnormality detection time = disk_check_interval (milliseconds) / 1000 x ( disk_check_retry(number of times) + 1 )
```

The abnormality detection time when the default value is used is shown below.

```
Abnormality detection time = 800 / 1000 \times (2 + 1)
= 2.4(seconds)
```



- The tuning described above impacts on the time taken from detection of a timeout until switching the primary server. Therefore, modify the values while taking into account the switch/disconnection time, using a design for which misdetection does not occur.
- Immediately selecting automatic degradation when a heartbeat abnormality occurs in operating system or server heartbeat monitoring risks causing split brain. Refer to "Appendix D Notes on Performing Automatic Degradation Immediately after a Heartbeat Abnormality" for details.



Mirroring Controller uses connections to database instances and SQL access to monitor abnormality in some resources targeted for monitoring. The connection destination database names and connection user names used for abnormality monitoring conform to the parameters in the server configuration file. The application name is "mc\_agent".

# 2.12 Setting Automatic Start and Stop of Mirroring Controller and Multiplexed Instances

Multiplexed instances and Mirroring Controller can be started and stopped automatically in line with the starting and stopping of the operating system of the database server.

However, to stop the operating system on the primary server, it is necessary to stop the operating system on the standby server so that no unnecessary automatic switching or errors occur.



To guarantee the startup sequence of Mirroring Controller on the primary and standby servers, first confirm that the primary server has started, and then start the standby servers in sequence.

The startup sequence of the Mirroring Controller process on the database server and the Mirroring Controller arbitration process on the arbitration server is not guaranteed. If the arbitration server cannot be started first, execute the mc\_ctl command in start mode with the --async-connect-arbiter option specified to start the Mirroring Controller process.

You can configure the Windows service to perform automatic start and stop of Mirroring Controller.

Setting automatic start and stop of a multiplexed instance

No settings are required for Mirroring Controller to start and stop an instance.



Do not configure the Windows service of a multiplexed instance to perform automatic start.

Configuring automatic start and stop for database multiplexing mode

Configuring during setup

When registering Mirroring Controller to the Windows service in "2.4.1 Setting Up Database Multiplexing Mode on the Primary Server" and "2.5.1 Setting Up Database Multiplexing Mode on the Standby Server", specify "auto" for the -S option of the register mode used with the mc\_ctl command.

 ${\tt Example)}$ 

> mc\_ctl register -M D:\mcdir\inst1 -P \*\*\*\*\*\*\* -S auto

### Changing the configuration after setup

Use the sc config command to change the configuration of the Windows service of Mirroring Controller.

Example)

The following is an example using the registered service name "Mirroring\_Controller\_instl".

......

> sc config "Mirroring\_Controller\_inst1" start= auto



Refer to documentation such as Windows Help and Support for the sc command for information on how to configure the service.



You can check the registration status in the Windows service window or by using the sc qc command.

# 2.13 Setting Automatic Start and Stop of the Mirroring Controller Arbitration Process

You can automatically start or stop the Mirroring Controller arbitration process when the operating system on the arbitration server is started or stopped.

### Linux

If using an rc script

Perform the following procedure:

1. Set the operating system user.

Configure the environment variables required for starting the Mirroring Controller arbitration process.

2. Edit the shell script.

Copy the sample shell script, and edit the environment variables in the shell script to suit the environment.

Sample shell script

```
/installDir/share/rc_mcarboi.sample
```

3. Copy the revised shell script and set access permissions.

Copy the revised shell script to /etc/rc.d/init.d, and then set and register the access rights.

Example)

In the following example, the shell script was created as "rc\_mcarb\_arbiter1".

```
# cp rc_mcarb_arbiter1 /etc/rc.d/init.d/
# chmod 755 /etc/rc.d/init.d/rc_mcarb_arbiter1
```

4. Register and enable automatic start and stop.

As the operating system superuser, execute the chkconfig command to register and enable the script. To register and enable the script, specify the options shown in the table below.

Option	Description	
add	Registers the shell script as the startup script.	
level	Sets and enables the startup script for each run level.	

#### Example)

In the following example, the shell script is registered, and then enabled for run levels 3 and 5.

# chkconfig --add rc\_mcarb\_arbiter1

# chkconfig --level 35 rc\_mcarb\_arbiter1 on

### If using systemd



If you use systemd to start the Mirroring Controller arbitration process, wait for time correction, network setup, and so on.

### Perform the following procedure:

### 1. Create a unit file.

Copy the unit file sample stored in the directory below, and revise it to match the target instance.

Sample file

/installDir/share/mcarboi.service.sample

#### Example)

In the following example, the installation directory is "/opt/fsepv<x>assistant", and the identifier of the arbitration process is "arbiter1". Note that "<x>" indicates the product version.

# cp /opt/fsepv<x>assistant/share/mcarboi.service.sample /usr/lib/systemd/system/
mcarboi\_arbiter1.service

Revise the underlined portions of the options below in the unit file.

Section	Option	Specified value	Description
Unit	Description	FUJITSU Enterprise Postgres Mirroring Controller Arbiter <arbitrationprocessid></arbitrationprocessid>	Specifies the feature overview.
			Specifies the identifier of the targeted arbitration process. (*1)
Service	ExecStart	/bin/bash -c 'installDir/bin/mc_arb_std start installDir mirroringControllerArbitrationProcessMgmtDir mc_arbOption'	Command to be executed when the service is started.
			Specify the option you want to add when the mc_arb command is executed without the -M option in the mc_arb option.
			Note that the content specified in this mc_arb option is carried over from the mc_arb_std command in "Specified
			value" to the mc_arb command.

Section	Option	Specified value	Description
	ExecStop	/bin/bash -c ' <u>installDir</u> 'bin/ mc_arb_std stop <u>installDir</u> <u>mirroringControllerArbitrationProcessMgmtDir</u> <u>mc_arbOption'</u>	Command to be executed when the service is stopped.
			Specify the option you want to add when the mc_arb command is executed without the -M option in the mc_arb option.  Note that the content specified in this me, or here.
			specified in this mc_arb option is carried over from the mc_arb_std command in "Specified value" to the mc_arb command.
	User	<u>User</u>	Specify the account of the operating system user.
	Group	Group	Specify the group to which the user belongs.

<sup>\*1:</sup> The arbitration process identifier used here is a name for identifying the Mirroring Controller arbitration process.

The naming conventions for identifying the Mirroring Controller arbitration process are as follows:

- Up to 16 bytes
- The first character must be an ASCII alphabetic character
- The other characters must be ASCII alphanumeric characters
- 2. Enable automatic start and stop.

As the operating system superuser, use the systemctl command to enable automatic start and stop.

Example)

# systemctl enable mcarboi\_arbiter1.service

### **Windows**

You can configure the Windows service to perform automatic start and stop.

### Configuring during setup

When registering the Mirroring Controller arbitration process as a Windows service in "2.3.1 Configuring the Arbitration Server", specify "auto" for the -S option of the register mode used with the mc\_arb command.

Example)

```
> mc_arb register -M D:\mcdir\inst1 -P ******* -S auto
```

### Changing the configuration after setup

Use the sc config command to change the configuration of the Windows service of the Mirroring Controller arbitration process.

Example)

The configuration of the registered service name "Mirroring\_Controller\_Arbiter1" is changed.

> sc config "Mirroring\_Controller\_Arbiter1" start= auto



Refer to sc command help for information on how to configure the service.



You can check the registration status in the Windows service window or by using the sc qc command.

# 2.14 Backup Operation

This section explains the backup operation for database multiplexing mode.

### 2.14.1 Backing up Database Multiplexing Mode Information

When changing the Mirroring Controller settings, in addition to backing up the database, back up the configuration file in the Mirroring Controller management directory so that the Mirroring Controller settings are not lost.

When the arbitration server is used for automatic degradation, also back up the configuration file in the Mirroring Controller arbitration process management directory.

### 2.14.2 Database Backup Operation

Using database multiplexing mode is the same as obtaining the backup data on the standby server as a safeguard against a disk failure. Note that all server disks may be corrupted due to some cause.

As a safeguard against this type of case, execute the pgx\_dmpall command on the primary server to create the backup data.

However, it is not definite as to which server runs as the primary server, so ensure that the pgx\_dmpall command is executed periodically on all servers, so that the backup data will be obtained. For example, create a script to obtain the backup data, and set it in the operation management software.



When the pgx\_dmpall command is executed on the standby server, it will not match the statuses, however the error message shown below will be output and return the value "1".

If a script that ignores only this type of error is executed on all servers, the backup data of the primary server can be obtained.

Error message

ERROR: recovery is in progress (10095)



- Consider the possibility that the server that runs as the primary server may be destroyed alongside the backup data, so it is recommended to promote another server to become the primary server, and then back up the data on the new primary server without waiting for the next scheduled backup.
- Specify the same backup directory name for the primary and standby servers. If different backup directory names are specified, and recovery is performed using the backup data of the other server, the recovery cannot be performed correctly.

# See See

- Period backups allow shorter recovery time and reduction in disk usage. Refer to "Backing Up the Database" in the Operation Guide for details on the backup operation.
- Refer to "Chapter 4 Action Required when an Error Occurs in Database Multiplexing Mode" for details on recovery based on the backup data that was obtained using the pgx\_dmpall command.

# Chapter 3 Operations in Database Multiplexing Mode

This chapter describes the periodic operations that are performed when running database multiplexing mode.

The periodic operations are the same as the operations on a single server.



See

Refer to "Periodic Operations" in the Operation Guide for information on the periodic operations.

# 3.1 Starting and Stopping the Mirroring Controller Arbitration Process

This section describes how to start and stop the Mirroring Controller arbitration process.

### 3.1.1 Starting the Mirroring Controller Arbitration Process

#### Linux

While the Mirroring Controller arbitration process is in a stopped state, execute the mc\_arb command in start mode to start the Mirroring Controller arbitration process.

......

Example)

\$ mc\_arb start -M /mcarb\_dir/arbiter1

#### Windows

The Mirroring Controller arbitration process can be started using one of the following options:

- Using the mc\_arb command
- Starting the service on system startup

### Using the mc\_arb command

While the Mirroring Controller arbitration process is in a stopped state, execute the mc\_arb command from the command prompt to start the Mirroring Controller arbitration process.

Example)

> mc\_arb start -M D:\mcarb\_dir\arbiter1

### Starting the service on system startup

Specify automatic start when registering the Mirroring Controller arbitration process to the Windows service during setup of database multiplexing mode. Accordingly, the Mirroring Controller arbitration process service will start on startup of the operating system.



Refer to the Reference for information on how to specify the mc\_arb command.

# 3.1.2 Stopping the Mirroring Controller Arbitration Process

#### Linux

While the Mirroring Controller arbitration process is running, execute the mc\_arb command in stop mode to stop the Mirroring Controller arbitration process.

Example)

\$ mc\_arb stop -M /mcarb\_dir/arbiter1

#### Windows

The Mirroring Controller arbitration process can be stopped using one of the following options:

- Using the mc\_arb command
- Stopping the service

### Using the mc\_arb command

While the Mirroring Controller arbitration process is running, execute the mc\_arb command in stop mode from the command prompt to stop the Mirroring Controller arbitration process.

Example)

> mc\_arb stop -M D:\mcarb\_dir\arbiter1

### Stopping the service

Select [Administrative Tools], then [Services] to open the [Services] window, and then select the Mirroring Controller service and click the [Stop] menu.



Refer to the Reference for information on how to specify the mc\_arb command.



- The arbitration server will be forcibly stopped when the service is stopped.
- Before shutting down the operating system on the arbitration server, either stop the Mirroring Controller on the primary server or standby server or shut down the operating system on the primary server or standby server.

# 3.2 Starting and Stopping Mirroring Controller

This section describes how to start and stop the Mirroring Controller.

# 3.2.1 Starting Mirroring Controller

Mirroring Controller can be started using one of the following:

- Using the mc\_ctl command
- Starting a service on system startup



Mirroring Controller must be started by a user with administrator privileges (user ID belonging to the Administrators group).

### Using the mc\_ctl command

While Mirroring Controller is in a stopped state, execute the mc\_ctl command from the command prompt to start the instance and Mirroring Controller.

### Enabling automatic switch/disconnection

Execute the mc\_ctl command in start mode.

Example)

> mc ctl start -M D:\mcdir\inst1

#### Disabling automatic switch/disconnection

Execute the mc\_ctl command in start mode with the -F option specified.

Example)

> mc\_ctl start -M D:\mcdir\inst1 -F



- To start the Mirroring Controller process only, execute the mc\_ctl command in start mode with the --mc-only option specified.
- After Mirroring Controller is started, automatic switch/disconnection can be enabled or disabled using the enable-failover or disable-failover mode of the mc\_ctl command.
- When the arbitration server is used for automatic degradation, the Mirroring Controller process startup fails on the database server if the Mirroring Controller arbitration process has not been started on the arbitration server in advance. However, even if the Mirroring Controller arbitration process cannot be started in advance, the Mirroring Controller process can be started by specifying the --async-connect-arbiter option in the mc\_ctl command.



See

Refer to the Command Reference for information on how to specify the mc\_ctl command.

### Starting a service on system startup

Specify automatic start when registering Mirroring Controller to the Windows service during setup of database multiplexing mode. Accordingly, the Mirroring Controller service will start on startup of the operating system.



See

Refer to "2.12 Setting Automatic Start and Stop of Mirroring Controller and Multiplexed Instances" for details.



- Mirroring Controller startup usually fails if the standby server is mistakenly started as the primary server or if the old primary server is not recovered after the switch and is then mistakenly started as the primary server. However, if the admin network is disconnected, then startup does not fail, and both servers may become primary servers. Therefore, ensure that the admin network is connected before starting Mirroring Controller.
- When only the instance is started without starting Mirroring Controller, the impact will be as follows:
  - Enabling automatic switch/disconnection

Features such as automatic switch and automatic disconnection will not work until Mirroring Controller is started.

- Disabling automatic switch/disconnection

Errors indicated in "1.1 What is Database Multiplexing Mode" will not be detected until Mirroring Controller is started.

### 3.2.2 Stopping Mirroring Controller

Mirroring Controller can be stopped using one of the following ways:

- Using the mc\_ctl command
- Stopping the service



Mirroring Controller must be stopped by a user with administrator privileges (user ID belonging to the Administrators group).

#### Using the mc\_ctl command

While Mirroring Controller is running, execute the mc\_ctl command from the command prompt in stop mode to stop Mirroring Controller.

Example)

> mc\_ctl stop -M D:\mcdir\inst1



To stop the Mirroring Controller process only, execute the mc\_ctl command in stop mode with the --mc-only option specified.



Refer to the Command Reference for information on how to specify the mc\_ctl command.

### Stopping the service

Select [Administrative Tools], then [Services] to open the [Services] window, and then select the Mirroring Controller service and click the [Stop] menu.



- To prevent an unintended automatic switch, before shutting down the operating system on the primary server, you must stop the Mirroring Controller, or shut down the operating system on the standby server.
- If you stop the Mirroring Controller by stopping the service, you must firstly exit all applications or programs that are using the instance that is to be stopped.
- When only the instance is stopped without stopping Mirroring Controller, the impact will be as follows:
  - Enabling automatic switch/disconnection

Mirroring Controller determines that an error has occurred in the instance, and performs an unnecessary automatic switch. Automatic switch may also stop working correctly in some cases.

- Disabling automatic switch/disconnection

Mirroring Controller determines that an error has occurred in the instance, and outputs an error to the event log.

# 3.3 Checking the Database Multiplexing Mode Status

### 3.3.1 Checking the Status of the Database Server

This section describes how to check the status of the database server.

Check the multiplexed database status by executing the mc\_ctl command in status mode.

Additionally, errors can be detected by monitoring the Mirroring Controller messages. If the status or messages are monitored periodically, you can react quickly following an automatic switch failure.

### Checking the status of the multiplexing database

When the mc\_ctl command is executed, the details of the multiplexing configuration, information about whether switch is possible following the error, and location and details of the error that caused the switch or disconnection are displayed.

After starting database multiplexing mode, execute the mc\_ctl command in status mode to check the multiplexing status.

An example of the status displayed when the mc\_ctl command is executed is shown below.

Example)

```
> mc_ctl status -M D:\mcdir\inst1
mirroring status
_____
switchable
server id host role
                                   host status db_proc_status
                                                                 disk status
server1 primary
                     192.0.2.100
                                               normal
                                                                   normal
                                  normal
server2 standby
                     192.0.2.110 normal
                                             normal
                                                                   normal
```

### Checking the status of connection to the Mirroring Controller arbitration process

When the arbitration server is used for automatic degradation, the status of the connection to the Mirroring Controller arbitration process can be checked by specifying the --arbiter option. If the output status is "online", it indicates that an arbitration request can be made from the database server to the arbitration server. When the arbitration server is used for automatic degradation, regularly execute the command in status mode with the --arbiter option specified and check that the output status is "online".

```
Example)
```

The mc\_ctl command is executed with the --arbiter option specified, and the status is output.

```
> mc_ctl status --arbiter -M D:\mcdir\inst1

arbiter_id host status
------
arbiter 192.0.3.120 online
```

### Checking the status of data synchronization

Additionally, by referencing the pg\_stat\_replication statistics view on the primary server, the data synchronization status can be confirmed. However, when creating the monitoring program, note that the content of pg\_stat\_replication may be changed in the future.

The following example shows that the locations of the transaction log after it is sent and received (sent\_lsn, replay\_lsn) match, and that they are fully synchronized.

Example)

```
application_name | standby
client_addr | 192.0.2.210
client_hostname
client_port | 55098
backend_start | 2020-03-23 11:17:49.628793+09
backend_xmin
state
              streaming
             0/3000060
sent_lsn
write_lsn
             0/3000060
             0/3000060
flush_lsn
             0/3000060
replay_lsn
write_lag
flush_lag
replay_lag
sync_priority
sync_state
               sync
reply_time
               2020-03-23 11:23:27.703366+09
```



- Refer to "mc\_ctl" in Reference for information on the command.
- Refer to "Notes on Application Compatibility" in the Application Development Guide for information on retaining application compatibility.
- Refer to "The Statistics Collector" in "Server Administration" in the PostgreSQL Documentation for details on pg\_stat\_replication.

### 3.3.2 Checking the Status of the Arbitration Server

This section describes how to check the status of the arbitration server.

The status of the connection between the Mirroring Controller arbitration process and primary server/standby server can be checked by executing the mc\_arb command in status mode.

The example below executes the mc\_arb command, and shows the status.

#### Linux

Example)

### Windows

Example)

```
> mc_arb status -M D:\mcarb_dir\arbiter1

server_id host status
------server1 192.0.3.100 online
server2 192.0.3.110 online
```

# 3.4 Manually Switching the Primary Server

The primary server cannot be switched automatically in the following case:

- If automatic switch/disconnection is disabled
- If output of messages is selected for heartbeat abnormalities during heartbeat monitoring of the operating system or server and the operating system/server crashes or becomes unresponsive

In this case, to manually switch the primary server, execute the mc\_ctl command in switch mode on either the primary server or the standby server.

Example)

> mc\_ctl switch -M D:\mcdir\inst1



If automatic switch/disconnection is enabled, it is possible to perform switch of primary server at any time.

# 3.5 Manually Disconnecting the Standby Server

The procedure to perform disconnection of the standby server differs depending on whether the automatic switch/disconnection is enabled or disabled.

If automatic switch/disconnection is enabled

Execute the mc\_ctl command in stop mode on the standby server.

Example)

```
> mc_ctl stop -M D:\mcdir\inst1
```

If automatic switch/disconnection is disabled

1. Execute the mc\_ctl command in stop mode on the standby server.

Example)

```
> mc_ctl stop -M D:\mcdir\inst1
```

- 2. Comment out the synchronous\_standby\_names parameter in the postgresql.conf file on the primary server.
- 3. Execute the pg\_ctl command in reload mode on the primary server.

Example)

```
> pg_ctl reload -D D:\database\inst1
```

# 3.6 Action Required when a Heartbeat Abnormality is Detected

The message below is output when a heartbeat abnormality is detected during heartbeat monitoring of operating systems or servers:

```
detected an error on the monitored object "server(server identifier name)": no response:ping timeout (MCA00019)
```

If the heartbeat\_error\_action parameter in *serverIdentifier*.conf is set to "message", even if automatic switch/disconnection is enabled and Mirroring Controller is started, automatic switch/disconnection is not performed when a heartbeat abnormality is detected. Therefore, user action will be necessary.

This section explains the action required when the heartbeat\_error\_action parameter is set to "message" and a heartbeat abnormality is detected.

- 1. Identify the cause of the heartbeat abnormality. The possible causes are below:
  - The remote operating system or server crashed or is unresponsive

- An admin network issue occurred
- 2. Address the cause identified in step 1.
  - The remote operating system or server crashed or is unresponsive
     Manually perform switch or disconnection using the mc\_ctl command.
  - An admin network issue occurred

Refer to "Chapter 4 Action Required when an Error Occurs in Database Multiplexing Mode", and recover the database multiplexing system.

# 3.7 Monitoring Mirroring Controller Messages

The messages that are output by Mirroring Controller are output to both the database server and the arbitration server. If the automatic switch fails, for example, an important message related to the continuation of the operation may be output, so ensure that the system log messages are monitored.

If the arbitration server is used for automatic degradation, monitor messages on both the database server and the arbitration server.

Message output destination on the database server

Messages are output to the event log.

Message output destination on the arbitration server

Linux

Messages are output to the system log.

Windows

Messages are output to the event log.



- To monitor message types considered to be important, an operating system setting must be configured beforehand. Refer to the operating system manuals, check if the message is of a message type that is monitored to be output to the system log, and configure the setting if required.
- If the heartbeat\_error\_action parameter in *serverIdentifier*.conf is set to "message", only message output is performed when a heartbeat abnormality is detected during heartbeat monitoring of operating systems and servers automatic switch/disconnection is not performed. Therefore users need to monitor the messages. Refer to "3.6 Action Required when a Heartbeat Abnormality is Detected" for details.

### Display format on the database server

eventSourceName[processId]: messageType: messageContent (messageNumber)

 $Specify \ the \ event \ source \ name \ in \ the \ event\_source \ parameter \ of \ the \ serverIdentifier. conf \ file \ of \ the \ database \ server.$ 

The message types output by Mirroring Controller, their severity, and their corresponding value in the event log are shown in the table below.

Table 3.1 Message type, severity, and corresponding value in the event log

Message type	Severity	Meaning	Event log
INFO	Information	Provides information that does not fall under LOG or NOTICE.	INFORMATI
LOG		Provides information recognized as a particularly important event in tracing the operation history.  (Example: Automatic switch is complete)	ON

Message type	Severity	Meaning	Event log
NOTICE	Notice	Outputs information that takes into account the user instructions within the program in response to an executed or automatically executed process.	
WARNING	Warning	Provides a warning, for example it will soon be impossible to maintain multiplexing capabilities.	WARNING
ERROR	Error	Reports that an error other than FATAL or PANIC has occurred.	ERROR
FATAL		Reports that an abnormality was detected in the multiplexed database systems requiring recovery of the system, and also the content and cause of the abnormality.	
PANIC		Reports that an abnormality was detected in all multiplexed database systems requiring immediate recovery of the system, and also the content and cause of the abnormality.	

The message severity has the following meanings:

- Information

Informational status. A message that was reported by the system is displayed. No action is required.

- Notice

Informational status, but a message that should be noted is displayed. If necessary, take the actions described in the "Action" section of the message.

- Warning

No error has occurred, but the user is requested to check, and take action. Take the actions described in the "Action" section of the message.

- Error

An error has occurred. Take the actions described in the "Action" section of the message.

### Display format on the arbitration server

### Linux

```
programName[processId]: messageType: messageText (messageNumber)
```

Specify the program name in the syslog\_ident parameter of the arbitration.conf file of the arbitration server.

### Windows

```
eventSourceName[processId]: messageType: messageText (messageNumber)
```

Specify the event source name in the event\_source parameter of the arbitration.conf file of the arbitration server.

The message types output by Mirroring Controller, their severity, and their corresponding value in the output destination log are shown in the table below.

Table 3.2 Message type, severity, and corresponding value in the output destination log

Message type	Severity	Meaning	System log (Linux)	Event log (Windows)
INFO	Information	Provides information not categorized as LOG or NOTICE.	INFO	INFORMATION
LOG		Provides information recognized as a particularly important event in tracing the operation history.		
		(Example: Automatic switch is complete)		

Message type	Severity	Meaning	System log (Linux)	Event log (Windows)
NOTICE	Notice	Outputs information that takes into account the user instructions within the program in response to an executed or automatically executed process.	NOTICE	
WARNING	Warning	Provides a warning, for example it will soon be impossible to perform the arbitration process.	WARNING	WARNING
ERROR	Error	Reports that an error other than FATAL or PANIC has occurred.	ERROR	ERROR
FATAL		Reports that an abnormality was detected in the arbitration server requiring recovery of the system, and also the content and cause of the abnormality.	CRIT	
PANIC		Reports that an abnormality was detected in the arbitration server requiring immediate recovery of the system, and also the content and cause of the abnormality.	ALERT	

The message severity has the following meanings:

- Information

Informational status. A message that was reported by the system is displayed. No action is required.

- Notice

Informational status, but a message that should be noted is displayed. If necessary, take the actions described in the "Action" section of the message.

- Warning

No error has occurred, but the user is requested to check, and take action. Take the actions described in the "Action" section of the message.

- Error

An error has occurred. Take the actions described in the "Action" section of the message.

### 3.8 Server Maintenance

To perform maintenance tasks such as periodic server inspections and the application of updates for software products including the operating system, you must perform a planned stop of the server, and then perform the maintenance.

# 3.8.1 Rolling Updates

In database multiplexing mode, rolling updates, that perform the maintenance for the servers that comprise the cluster system, can be performed while jobs continue.

First, perform the maintenance for the standby server, and then switch the standby server to the primary server. Then, perform the maintenance for the original primary server that was switched to the standby server. This enables maintenance to be performed while jobs continue.

Note that arbitration server maintenance can be performed without affecting database server operation, so it is not necessary to consider rolling update.



If the downtime due to the maintenance of the standby server is expected to be long, refer to "Standby server downtime" in "3.9.1 Changes Required when the Standby Server is Stopped".

The flow of a rolling update is shown below.

System administrator Application Maintenance tasks on the standby serv Standby server Primary server Database Database Change the server role to 'primary server' Application Standby server Primary server Database Database System administrator Application Standby server Primary server Data Database Maintenance tasks on the new standby server Application System administrator Standby server Primary server Database Database Failback to the Application Operation from the primary server (only application if required) System status transition

Figure 3.1 Performing a Rolling Update

Perform the following procedure as shown in the above figure:

Primary server

Database

### Standby server maintenance tasks

1. To perform the maintenance on the standby server, stop Mirroring Controller.

Example)

> mc\_ctl stop -M D:\mcdir\inst1

Standby server

Database

Operation to each

server

2. Ensure that Mirroring Controller has completely stopped.

If the multiplexed instances and Mirroring Controller have been configured on the standby server to start and stop automatically when the operating system of the database server is started or stopped, cancel the setting to start and stop automatically.



### See

Refer to "2.12 Setting Automatic Start and Stop of Mirroring Controller and Multiplexed Instances" for information on how to configure the multiplexed instances and Mirroring Controller to start and stop automatically when the operating system of the database server start and stops.

This task should be performed by the instance administrator user with administrator privileges.

Use the sc config command to disable automatic start of multiplexed instances and Mirroring Controller from the Windows service.

Example)

The following is an example using the registered service name "Mirroring\_Controller\_instl".

> sc config "Mirroring\_Controller\_inst1" start= demand



### Information

You can use the sc qc command to check the registration status.

Refer to documentation such as Windows Help and Support for the sc command for information on registry content.

- 3. Perform maintenance tasks.
- 4. Create a copy of the primary server instance on the standby server.

Execute the pg\_basebackup command to create data in the standby server by synchronizing with the primary server.

Example)

> pg\_basebackup -D D:\database\inst1 -X fetch --waldir=E:\transaction\inst1 --progress --verbose
-R --dbname="application\_name=standbyServerName" -h primaryServerHostName -p
primaryServerPortNumber



### See

The procedure for copying the primary server instance to the standby server is the same as the procedure for setting up the standby server

......

Refer to "2.5.2 Creating, Setting, and Registering the Standby Server Instance", and then perform the recovery.

5. Check the settings for automatic start and stop of the multiplexed instances and Mirroring Controller.

If the multiplexed instances and Mirroring Controller were configured in step 2 to not start and stop automatically when the operating system of the database server starts and stops, then change the settings back. This step can be skipped if automatic start and stop are not required.

This task should be performed by an instance administrator user with administrator privileges.

Use the sc config command to enable automatic start of multiplexed instances and Mirroring Controller from the Windows service.

Example)

The following is an example using the registered service name "Mirroring\_Controller\_instl".

> sc config "Mirroring\_Controller\_inst1" start= auto



You can use the sc qc command to check the registration status.

Refer to documentation such as Windows Help and Support for the sc command for information on registry content.

6. Start (rebuild) Mirroring Controller on the standby server.

This operation is required when determining the maintenance tasks on the standby server.

Enabling automatic switch/disconnection

As the instance administrator user, execute the mc\_ctl command in start mode.

Example)

```
> mc_ctl start -M D:\mcdir\inst1
```

### Disabling automatic switch/disconnection

As the instance administrator user, execute the mc\_ctl command in start mode with the -F option specified.

Example)

```
> mc_ctl start -M D:\mcdir\inst1 -F
```



After Mirroring Controller is started, automatic switch/disconnection can be enabled or disabled using the enable-failover or disablefailover mode of the mc\_ctl command.

### Switching to the primary server

To perform the maintenance on the primary server, execute the mc\_ctl command in the switch mode on the primary server or the standby server.

Example)

```
> mc_ctl switch -M D:\mcdir\inst1
```

When the switch is complete, the synchronous\_standby\_names parameter in the postgresql.conf file of the new primary server will be commented as follows:

Example)

```
#synchronous_standby_names = 'primary'
```

### New standby server maintenance tasks

1. Stop the Mirroring Controller.

On the new standby server (the primary server before the switch), execute the mc\_ctl command in stop mode.

Example)

```
> mc_ctl stop -M D:\mcdir\inst1
```

2. Ensure that Mirroring Controller has completely stopped.

If the multiplexed instances and Mirroring Controller have been configured on the new standby server to start and stop automatically when the operating system of the database server is started or stopped, cancel the setting to start and stop automatically now.



Refer to "2.12 Setting Automatic Start and Stop of Mirroring Controller and Multiplexed Instances" for information on how to configure the multiplexed instances and Mirroring Controller to start and stop automatically when the operating system of the database server starts and stops.

This task should be performed by an instance administrator user with administrator privileges.

Use the sc config command to disable automatic start of multiplexed instances and Mirroring Controller from the Windows service.

Example)

The following is an example using the registered service name "Mirroring\_Controller\_instl".

sc config "Mirroring\_Controller\_inst1" start= demand



### Information

You can use the sc qc command to check the registration status.

Refer to documentation such as Windows Help and Support for the sc command information on registry content.

- 3. Perform the maintenance on the new standby server that was stopped.
- 4. Create a copy of the new primary server instance on the new standby server.

Execute the pg\_basebackup command to create data in the new standby server by synchronizing with the new primary server.

Example)

> pg\_basebackup -D D:\database\inst1 -X fetch --waldir=\transaction\inst1 --progress --verbose R --dbname="application\_name=standbyServerName" -h primaryServerHostName -p
primaryServerPortNumber



### See

The procedure for copying the primary server instance to the standby server is the same as the procedure for setting up the standby server.

Refer to "2.5.2 Creating, Setting, and Registering the Standby Server Instance", and then perform the recovery.

5. Check the settings for automatic start and stop of the multiplexed instances and Mirroring Controller.

If the multiplexed instances and Mirroring Controller were configured in step 2 to not start and stop automatically when the operating system of the database server starts and stops, then change the settings back. This step can be skipped if automatic start and stop are not required.

This task should be performed by an instance administrator user with administrator privileges.

Use the sc config command to enable automatic start of multiplexed instances and Mirroring Controller from the Windows service.

Example)

The following is an example using the registered service name "Mirroring\_Controller\_instl".

> sc config "Mirroring\_Controller\_instl" start= auto



### Information

You can use the sc qc command to check the registration status.

Refer to the document such as Windows Help and Support for the sc command for information on registry content.

6. After the maintenance is complete, edit the following parameters in the postgresql.conf file of the standby server as required.

Copying an instance results in the value of the synchronous\_standby\_names parameter becoming the specified value on the primary server. Therefore, correct it to the specified value on the standby server. If the parameter was commented out, then you must uncomment it.

7. On the standby server, start (rebuild) Mirroring Controller.

### Enabling automatic switch/disconnection

As the instance administrator user, execute the mc\_ctl command in start mode.

Example)

> mc\_ctl start -M D:\mcdir\inst1

### Disabling automatic switch/disconnection

As the instance administrator user, execute the mc\_ctl command in start mode with the -F option specified.

Example)

> mc\_ctl start -M D:\mcdir\inst1 -F



After Mirroring Controller is started, automatic switch/disconnection can be enabled or disabled using the enable-failover or disable-failover mode of the mc\_ctl command.

......

### Failback of the Primary Server

Revert the primary server and standby server to the original server configuration. Do this to execute the main job on the previous primary server. Refer to "4.1.1.3 Failback of the Primary Server" for details.



Obtain a backup as soon as this task is complete.

# 3.8.2 Stopping for Maintenance

Perform this procedure to stop all servers for periodic inspections, for example. On the server on which Mirroring Controller is running, execute the mc\_ctl command in stop mode to stop the instance and Mirroring Controller.

After that, on the server where the Mirroring Controller arbitration process is running, execute the mc\_arb command in stop mode to stop the Mirroring Controller arbitration process.

### **Stopping Mirroring Controller**

Example)

> mc\_ctl stop -M D:\mcdir\inst1 -a

### Stopping the Mirroring Controller arbitration process

Linux

Example)

\$ mc\_arb stop -M /mcarb\_dir/arbiter1

### Windows

Example)

```
> mc_arb stop -M D:\mcarb_dir\arbiter1
```

# 3.8.3 Arbitration Server Maintenance

Arbitration server maintenance can be performed without affecting database server operation.

Follow the procedure below to perform arbitration server maintenance.

1. Execute the mc\_arb command in stop mode to forcibly stop the Mirroring Controller arbitration process.

### Linux

Example)

```
$ mc_arb stop -M /mcarb_dir/arbiter1 -e
```

#### Windows

Example)

```
> mc_arb stop -M D:\mcarb_dir\arbiter1 -e
```

- 2. Perform maintenance tasks.
- 3. Execute the mc\_arb command in start mode to restart the Mirroring Controller arbitration process.

#### Linux

Example)

```
$ mc_arb start -M /mcarb_dir/arbiter1
```

### Windows

Example)

```
> mc_arb start -M D:\mcarb_dir\arbiter1
```

4. Execute the mc\_arb command in status mode to check that the arbitration server is connected to the database server.

The example below executes the mc\_arb command, and shows the status.

### Linux

Example)

```
$ mc_arb status -M /mcarb_dir/arbiter1

server_id host status
------server1 192.0.3.100 online
server2 192.0.3.110 online
```

### Windows

Example)

```
> mc_arb status -M D:\mcarb_dir\arbiter1
server_id host status
```

### 5. Check the command output.

Items to be checked

Check that the output status is "online" on both lines.

# 3.9 Changes in Operation

The following changes in operation may be required:

- Changes required when the standby server is stopped
- Changing from single server mode to database multiplexing mode
- Changing from database multiplexing mode to single server mode
- Changing to database multiplexing mode when the arbitration server is used for automatic degradation
- Changing parameters
- Uninstalling in the database multiplexing mode

### 3.9.1 Changes Required when the Standby Server is Stopped

### Operation when the standby server is stopped

Before performing maintenance for the primary server instance when the standby server has been stopped, stop Mirroring Controller on the primary server, comment out the synchronous\_standby\_names parameter in the postgresql.conf file of the primary server, and then execute the pg\_ctl command in reload mode.

If this operation is not performed, operations performed on the primary server for the instance will remain in a wait state.



Refer to "pg\_ctl" in Reference for information on the command.

### Standby server downtime

If you specified the synchronous\_standby\_names parameter of the postgresql.conf file and then the standby server instance is stopped, consider the points below.

- The wal\_sender\_timeout parameter in the postgresql.conf file
  - If the standby server is stopped after the timeout time set in this parameter was exceeded, an error stating that the transaction log could not be received may be output to the primary server event log, and all transaction logs that should be transferred to the standby server are accumulated.
- The wal\_keep\_segments parameter in the postgresql.conf file

If a transaction log that exceeds the value set in this parameter was generated while the standby server was stopped, the transaction log may be deleted.

Additionally, setting this parameter requires consideration regarding stabilization of the database multiplexing mode. Refer to "2.11.1 Tuning to Stabilize the Database Multiplexing Mode" for details.



The standby server must be rebuilt if the pending transaction log to be transferred to the standby server is lost when the standby server is started after the maintenance task is complete.

Take the action advised in the recovery operation that starts from "4.1.1.1.3 Identify cause of error and perform recovery" through to "4.1.1.2 Rebuild the Standby Server".

### 3.9.2 Changing from Single Server Mode to Database Multiplexing Mode

The procedure for switching single server mode to database multiplexing mode for the purposes of high reliability and load distribution of the system is explained below.

This procedure is equivalent to the setup procedure explained in "Chapter 2 Setting Up Database Multiplexing Mode".



#### If the data storage destination directory name is not comprised of ASCII characters

Stop the application job and then perform the following procedure to migrate to a directory with a name that uses only ASCII characters:

- 1. Stop the database instance on the primary server.
- 2. Change the name of the data storage destination directory to one that uses only ASCII characters.

  For example, do not perform operations that will cause the symbolic link contained in the data storage destination directory to become lost, such as moving data to another drive.



When encrypting the storage data, refer to "Database Multiplexing Mode" in the Operation Guide, and then perform the setup for encryption on the primary and standby servers.

......

1. Install on the arbitration server

Perform this step only if the arbitration server is used for automatic degradation.

Install the Server Assistant on the server where the Mirroring Controller arbitration process is started.

Refer to "Installation" in the Installation and Setup Guide for Server Assistant for information on how to install the Server Assistant.

2. Install on the standby server

Install FUJITSU Enterprise Postgres on the server to be started as the standby server.

Refer to "Installation" in the Installation and Setup Guide for Server for information on how to install FUJITSU Enterprise Postgres.

Use ASCII characters in the data storage destination directory.

3. Stop the application jobs

Stop the application jobs to be connected to the primary server.

4. Change the primary server settings

To allow connections from the server to be started as the standby server, configure the settings in step 2 and thereafter of "2.4.2 Creating, Setting, and Registering the Primary Server Instance" on the primary server.

5. Set up the arbitration server

Refer to "2.3 Setting Up the Arbitration Server" for details.

Perform this step only if the arbitration server is used for automatic degradation.

6. Set up database multiplexing mode on the primary server

Refer to "2.4.1 Setting Up Database Multiplexing Mode on the Primary Server" for details.

7. Set up database multiplexing mode on the standby server

Refer to "2.5.1 Setting Up Database Multiplexing Mode on the Standby Server" for details.

8. Create the standby server instance and start it

Refer to "2.5.2 Creating, Setting, and Registering the Standby Server Instance" for details.

After the above steps are completed, refer to the remaining explanations in "Chapter 2 Setting Up Database Multiplexing Mode" and ensure that the required settings and operations are completed.

## 3.9.3 Changing from Database Multiplexing Mode to Single Server Mode

The procedure for stopping database multiplexing mode and changing to single server mode is explained below.

Some tasks must be performed on the database server, and others must be performed on the arbitration server.

The tasks on the arbitration server are required only if the arbitration server is used for automatic degradation.

#### Tasks on the database server

1. Determine the server for which the instance is to be stopped, and switch this server

Determine the server that is to be excluded as the database multiplexing mode target, and for which the instance is to be stopped.

If the server for which the instance is to be stopped is the primary server, execute the mc\_ctl command in the switch mode to switch the standby server to the primary server.

The standby server after the switch is complete will be the server for which the instance is to be stopped.

If the server for which the instance is to be stopped is the standby server, there is no need to perform the switch operation.

Example)

```
> mc_ctl switch -M D:\mcdir\inst1
```

2. Stop Mirroring Controller and the instance.

On the server that was determined in step 1, execute the mc\_ctl command in the stop mode to stop Mirroring Controller and the instance.

Example)

```
> mc_ctl stop -M D:\mcdir\inst1
```

3. Unregister Mirroring Controller from the Windows service.

Execute the mc\_ctl command in unregister mode to unregister Mirroring Controller from the Windows service.

Example)

```
> mc_ctl unregister -M D:\mcdir\inst1
```

4. Delete registries related to the event log

If error logs are output to the event log in "2.2.1.2 Preparatory Tasks for the Output of Error Logs to the Event Log", delete the registered event source name for each instance.

Example)

```
> regsvr32 /u /i:"Mirroring Controller instl" "c:\Program Files\Fujitsu\fsepv<x>server64\lib \mcevent.dll"
```

Note that "<x>" indicates the product version.

5. Delete the file resources

Delete the following file resources:

- Data storage destination directory

- Mirroring Controller management directory

Example)

```
> rmdir /S /Q D:\database\instl
> rmdir /S /Q D:\mcdir\instl
```



Refer "Security-Related Notes" in the Operation Guide for information on deleting the data securely.

6. Stop the application jobs

Stop the application jobs to be connected to the primary server.

7. Stop Mirroring Controller and the instance on the primary server

Execute the mc\_ctl command in stop mode on the primary server.

Example)

```
> mc_ctl stop -M D:\mcdir\inst1
```

8. Unregister Mirroring Controller from the Windows service on the primary server

Execute the mc\_ctl command in unregister mode to unregister Mirroring Controller from the Windows service.

Example)

```
> mc_ctl unregister -M D:\mcdir\inst1
```

9. Delete registries related to the event log on the primary server

If error logs are output to the event log in "2.2.1.2 Preparatory Tasks for the Output of Error Logs to the Event Log", delete the registered event source name for each instance.

Example)

```
> regsvr32 /u /i:"Mirroring Controller instl" "c:\Program Files\Fujitsu\fsepv<x>server64\lib \mcevent.dll"
```

Note that "<x>" indicates the product version.

10. Delete the database multiplexing mode settings that were configured for the primary server instance.

Reset the postgresql.conf file parameters to their values before the database multiplexing operation was set.

Delete the file resources from the Mirroring Controller management directory.

If the backup operation was performed, delete the following resources:

- Mirroring Controller management directory backup data obtained in database multiplexing mode
- Instance backup data obtained in database multiplexing mode

Additionally, if the primary\_conninfo parameter is set in the postgresql.auto.conf file, execute the ALTER SYSTEM RESET statement to delete the setting.

Example)

An example execution of the psql command is shown below.

```
postgres=# ALTER SYSTEM RESET primary_conninfo;
```

After these actions are performed, ensure that the backup data is collected when starting the single operation.



- Refer to "Security-Related Notes" in the Operation Guide for details on deleting the data securely.
- Refer to "2.14 Backup Operation" for details on the backup operation.
- Refer to "Appendix A Parameters" for details on the postgresql.conf file parameters.



In the above procedure, if the postgresql.conf file of the single primary server can be changed by reloading the file, the operation mode can be changed without stopping the application job.

In that case, execute the mc\_ctl command in stop mode with the --mc-only option specified to stop only Mirroring Controller in relation to stopping the primary server.

### Tasks on the arbitration server

#### Linux

1. Execute the mc\_arb command in stop mode to stop the Mirroring Controller arbitration process.

Example)

```
$ mc_arb stop -M /mcarb_dir/arbiter1
```

2. Delete the Mirroring Controller arbitration process management directory.

Example)

```
$ rm -rf /mcarb_dir/arbiter1
```

### Windows

1. Execute the mc\_arb command in stop mode to stop the Mirroring Controller arbitration process.

Example)

```
> mc_arb stop -M D:\mcarb_dir\arbiter1
```

2. Unregister the Mirroring Controller arbitration process from the Windows service.

Execute the mc\_arb command in unregister mode to unregister the Mirroring Controller arbitration process from the Windows service.

Example)

```
> mc_arb unregister -M D:\mcarb_dir\arbiter1
```

3. Delete registrations related to the event log

If error logs are output to the event log in "2.2.2.1 Preparing to Output Error Logs to the Event Log (Windows)", delete the registered event source name for each instance.

Example)

```
> regsvr32 /u /i:"Mirroring Controller arbtier1" "c:\Program Files\Fujitsu
\fsepv<x>assistant64\lib\mcarbevent.dll"
```

Note that "<x>" indicates the product version.

4. Delete the Mirroring Controller arbitration process management directory.

Example)

```
> rmdir /S /Q D:\mcarb_dir\arbiter1
```

# 3.9.4 Changing to Database Multiplexing Mode when the Arbitration Server is Used for Automatic Degradation

This section provides the procedure to change to database multiplexing mode using the Mirroring Controller only on the database server when the arbitration server is used for automatic degradation.

Some tasks must be performed on the database server, and others must be performed on the arbitration server.

#### Tasks on the arbitration server

1. Set up the arbitration server.

Refer to "2.3 Setting Up the Arbitration Server" for information on how to set up the arbitration server.

#### Tasks on the database server

1. On the server where Mirroring Controller is running, execute the mc\_ctl command in stop mode to stop Mirroring Controller on the primary server and standby server.

Example)

```
> mc_ctl stop -M D:\mcdir\inst1 -a --mc-only
```

2. Edit the network.conf file of the primary server and standby server to add the information of the arbitration server.

Refer to "A.3 Network Configuration File" for details.

The definition example of the network.conf file of the primary server is shown below:

```
Example)
```

The IDs of the primary server and standby server are set to "server1" and "server2", and their port numbers are set to "27540" and "27541". The ID of the server of the Mirroring Controller arbitration process is set to "arbiter", and its port number is set to "27541".

```
serverl 192.0.2.100,192.0.3.100 27540,27541 server
server2 192.0.2.110,192.0.3.110 27540,27541 server
arbiter 192.0.3.120 27541 arbiter
```



- Ensure that the port numbers set for the primary server, standby server, and arbitration server do not conflict with other software. Also do not configure the same segment for the admin network and arbitration network.
- If the server type is "server", two IP addresses or host names, and two port numbers need to be specified in the following order:
  - IP address or host name of the database server used as the admin network
  - IP address or host name of the database server used as the arbitration network
  - Port number of the database server used as the admin network
  - Port number of the database server used as the arbitration network
- If the server type is "arbiter", specify the IP address or host name set for the my\_address parameter and the port number set for the port parameter in arbitration.conf.

3. Edit the *serverIdentifier*.conf file of the primary server and standby server to add parameters required for the operation where the arbitration server is used for automatic degradation.

Refer to "A.4.1 Server Configuration File for the Database Servers" for information on the parameters required when the arbitration server is used for automatic degradation.

4. On the primary server and standby server, execute the mc\_ctl command in start mode to start the Mirroring Controller process.

Example)

```
> mc_ctl start -M D:\mcdir\inst1 --mc-only
```

#### Common tasks

1. Check the connection status from the database server or arbitration server.

Refer to "2.8 Checking the Connection Status" for details.

### 3.9.5 Changing Parameters

Stop Mirroring Controller before editing the Mirroring Controller server configuration file and network configuration file.

If the Mirroring Controller process crashes or becomes unresponsive, restart is performed automatically by the Mirroring Controller monitoring process, and the configuration file is reloaded. Therefore, if the configuration file was being edited, unintended behavior will occur.

### 3.9.6 Uninstalling in Database Multiplexing Mode

This section explains how to uninstall FUJITSU Enterprise Postgres on a server using database multiplexing mode.

Some tasks must be performed on the database server, and others must be performed on the arbitration server.

The tasks on the arbitration server are required only if the arbitration server is used for automatic degradation.

#### Tasks on the database server

1. Stop the multiplexed instances and Mirroring Controller

Refer to "3.2 Starting and Stopping Mirroring Controller" for information on how to stop the instance.

2. Unregister Mirroring Controller from the Windows service

Execute the mc\_ctl command in unregister mode to unregister Mirroring Controller from the Windows service.

Example)

```
> mc_ctl unregister -M D:\mcdir\inst1
```

3. Delete registries related to the event log

If messages are output to the event log, DLLs are registered in accordance with "2.2.1.2 Preparatory Tasks for the Output of Error Logs to the Event Log". Delete these registries so that no unnecessary issues occur.

```
Example)
```

The following is an example in which the DLL of a 64-bit product that is registered under the default event source name is deleted. Note that "<x>" indicates the product version.

```
> regsvr32 /u "c:\Program Files\Fujitsu\fsepv<x>server64\lib\mcevent.dll"
```

- Delete the registered DLL for each instance

DLL registration is performed so that messages output to the event log are identified by each instance, and are output to any event source named by the user.

Accordingly, it is necessary to delete the DLL registry for each instance. Delete the DLL registry for each event source name.

```
Example)
```

The following is an example in which the DLL of a 64-bit product that is registered under the

event source name "Mirroring Controller instl" is deleted. Note that "<x>" indicates the product version.

> regsvr32 /u /i:"Mirroring Controller instl" "c:\Program Files\Fujitsu\fsepv<x>server64\lib
\mcevent.dll"

#### - If installing multiple versions

If the database multiplexing system you set up using the FUJITSU Enterprise Postgres package has been set to output error logs to the event log, use the DLL path name that you took note of previously as explained in "2.2.1.2 Preparatory Tasks for the Output of Error Logs to the Event Log" to reregister the default event source name.



Ensure that you delete the DLLs before the uninstallation. If you perform the uninstallation without doing so, you may not be able to delete the DLLs at a later time.

#### 4. Uninstall FUJITSU Enterprise Postgres

Refer to "Uninstallation" in the Installation and Setup Guide for Server for information on how to uninstall FUJITSU Enterprise Postgres.

#### Tasks on the arbitration server

Refer to "Uninstallation" in the Installation and Setup Guide for Server Assistant, and uninstall the Server Assistant.

# Chapter 4 Action Required when an Error Occurs in Database Multiplexing Mode

This chapter describes the action required if an error occurs in database multiplexing mode.

In database multiplexing mode, when an error is detected, the switch or disconnection of the standby server is performed automatically, so that only the primary server starts degrading. In this case, the recovery tasks will be required for the standby server on which the switch or disconnection was performed.

Other possible cases are as follows:

- When automatic switch fails
- When automatic disconnection fails
- When all servers or instances were stopped

### 4.1 Action Required when Server Degradation Occurs

If the server has started degrading, the recovery tasks will vary depending on whether the cause was the switch (failover or switchover), or the disconnection.

Execute the mc\_ctl command in status mode, or refer to the event log, and check if the cause of the server to start degrading was the switch or the disconnection.

In the example below, the mc\_ctl command is executed in status mode.

If a switch has occurred, "switched" (the switch is complete and the server is in a degrading state) is displayed for "mirroring status".

Example)

```
> mc_ctl status -M D:\mcdir\inst1
mirroring status
-----switched
:
```

If a disconnection has occurred, "not-switchable" (disconnection was performed so the server cannot be switched) is displayed for "mirroring status".

Example)

```
> mc_ctl status -M D:\mcdir\inst1
mirroring status
-----
not-switchable
:
```



If Mirroring Controller detects any errors on the server on which operations are continuing during recovery to database multiplexing mode from a degrading operation state, perform the procedure in "4.1.3 Addressing Errors During Degrading Operation", and then recover to database multiplexing mode.

# 4.1.1 Operations when the Server has Started Degrading after a Switch has Occurred

This section explains the operations when the server has started degrading after a switch has occurred.



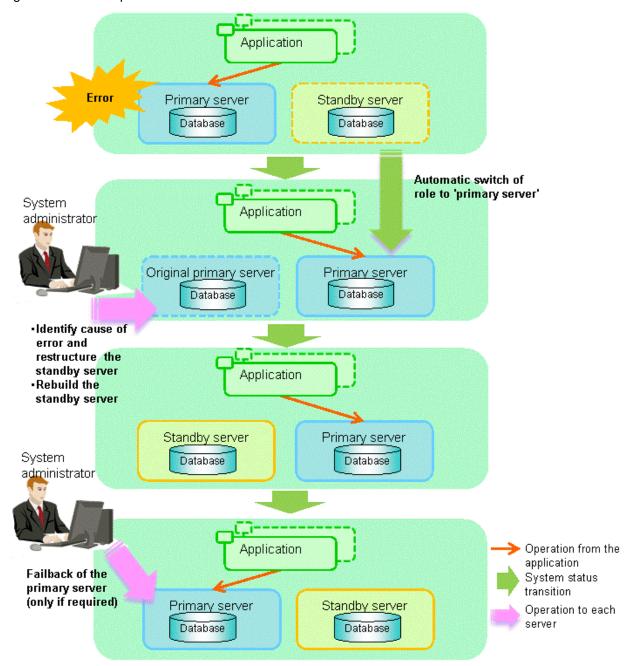
- After a switch has occurred as a result of an abnormality on the primary server, the database will not have a multiplexed configuration until the standby server is rebuilt. Remove the cause of the error as quickly as possible, and then rebuild the standby server.
- If the reference job was executed on the standby server, and the servers are switched because an error occurred on the primary server, the load is concentrated on the new primary server. Accordingly, pause the reference job on the original standby server, rebuild the original primary server as the new standby server, and then resume the reference job for the new standby server.
- If the instance on the new primary server is stopped before the original primary server where the error occurred is rebuilt as the new standby server, a split brain occurs at startup from the instance on the original primary server. Therefore, start the instance on the new primary server before rebuilding the standby server.

If the switch occurred and the server has started degrading, perform the following operations to recover the standby server and revert it to its original state:

- Identify Cause of Error and Restore the Standby Server
- Rebuild the Standby Server
- Failback of the Primary Server (only if required)

The flow of these operations is shown in the figure below.

Figure 4.1 Flow of operations



### 4.1.1.1 Identify Cause of Error and Restore the Standby Server

Perform the recovery according to the following procedure:

- 1. Stop Mirroring Controller
- 2. Recovery of the Mirroring Controller management directory
- 3. Identify cause of error and perform recovery

### 4.1.1.1.1 Stop Mirroring Controller

Execute the mc\_ctl command in stop mode for the original primary server on which the error occurred.

Example)

> mc\_ctl stop -M D:\mcdir\inst1

This also stops the instance that is required to perform the recovery.



If the instance does not stop, refer to "Actions in Response to Failure to Stop an Instance" in the Operation Guide, and then stop the instance.

Then, specify the -e option in the above command to forcibly stop Mirroring Controller.

### 4.1.1.1.2 Recovery of the Mirroring Controller management directory

Copy the files in the Mirroring Controller management directory from the backup data, and then perform the recovery.

### 4.1.1.3 Identify cause of error and perform recovery

Refer to the event log of the primary server and the standby server to identify the cause of the error, and then perform recovery.

The following commands can be used to recover a standby server. Select depending on the recovery and the situation.

- pg\_basebackup

Creates a copy of all resources of the primary server instance.

- pg\_rewind

Creates a copy of only the updated files on the new primary server. For this reason, if this command is used to incorporate a new standby server, recovery time can be shortened. To use this command to build the original primary server as a new standby server, at least one of the following must be met:

- a. Checksums were enabled when an instance was created and the original primary server was stopped normally, or
- b. The original primary server was stopped normally and the wal\_log\_hints parameter of postgresql.conf was enabled when an instance was started.

Additionally, full\_page\_writes must be enabled, which is its default value.

If it is required to start the original primary server again to then stop it normally, both servers will temporarily operate as primary servers. If data is updated on the original primary server, data corruption will occur, so take measures such as changing settings in pg\_hba.conf to disallow client connections.

When executing the pg\_rewind command, the timeline ID of the source server and the target server must be different. However, immediately after promotion of the new primary server, it may not be possible to execute the pg\_rewind command due to the application processing of unapplied update transaction logs and subsequent update processing of timeline IDs using checkpoint processing not being completed. Therefore, execute the pg\_rewind command after unapplied update transaction logs are no longer on the new primary server and updating of timeline IDs is complete.



 $- \ Refer to \ "pg\_basebackup" \ in \ "Reference" \ in the \ PostgreSQL \ Documentation \ for \ information \ on the \ pg\_basebackup \ command.$ 

•••••

- Refer to "pg\_rewind" in "Reference" in the PostgreSQL Documentation for information on the pg\_rewind command.

The example below executes the pg\_rewind command to perform recovery by synchronizing data on the original primary server with the new primary server.

1. Use the pg\_ctl command to start the original primary server instance.

Example)

> pg\_ctl start -D D:\database\inst1

2. Use the pg\_ctl command to stop the original primary server instance.

Example)

```
> pg_ctl stop -D D:\database\instl
```

3. Wait for the application of unapplied update transaction logs on the new primary server.

Execute the SQL below on the new primary server, and wait until the result is false.

```
# select pg_is_in_recovery();
```

Example)

```
> psql -h hostNameOfNewPrimaryServer -p portNumOfNewPrimaryServer -d dbName -c "select pg_is_in_recovery();"
```

Any database can be connected to.



If the pg\_rewind command is executed immediately after promotion of the new primary server, the processing in steps 3 and 4 is required. If update-type SQL can be executed on the new primary server and checkpoint processing is executed after promotion, the processing in steps 3 and 4 will not be necessary.

4. Update the timeline ID.

Execute checkpoint processing, and update the timeline ID.

```
> psql -h hostNameOfNewPrimaryServer -p portNumOfNewPrimaryServer -d dbName -c "checkpoint;"
```

Any database can be connected to.

5. Create a copy of the new primary server instance in the original primary server (new standby server).

Execute the pg\_rewind command to synchronize the new standby server data with the new primary server.

Example)

```
\verb|-pg_rewind_-D_D:\database\\instl_--source-server="user=userName_host=newPrimaryServerHostName_port=newPrimaryServerPortNumber_dbname=dbName"|
```

6. Specify parameters in the postgresql.conf file of the original primary server (new standby server).

Set the parameters required for the standby server in postgresql.conf.

Refer to "Table 2.5 Parameters" for information on the parameters to set in postgresql.conf.

7. Create the standby.signal file on the original primary server (new standby server).

The file will be used to connect to the new primary server for streaming replication.

8. Use the pg\_ctl command to start the original primary(new standby server) server instance.

Example)

```
> pg_ctl start -D D:\database\inst1
```

9. Use the ALTER SYSTEM SET statement to change the connection string specified for the primary\_conninfo parameter in postgresql.auto.conf file using as shown below.

Table 4.1 Changes to the connection string

Keyword	Content specified	Remarks
host	hostNameOfNewPrimaryServer	Specify informations for connecting to the
port	portNumOfNewPrimaryServer	primary server in the form of a PostgreSQL connection string.
application_name	newstandbyServerName	TostgresQL connection string.

Example)

An example execution of the psql command is shown below.

postgres=# ALTER SYSTEM SET primary\_conninfo='host=hostNameOfNewPrimaryServer
port=portNumOfNewPrimaryServer application\_name=newstandbyServerName';



- If you need to set a connection string other than host, port and application\_name, include it in the setting of the primary\_conninfo parameter.

......

- The primary\_conninfo parameter should not be set in the postgresql.conf file, but only in the postgresql.auto.conf file using the ALTER SYSTEM SET statement.
- 10. Use the pg\_ctl command to stop the original primary server (new standby server) instance.

Example)

> pg\_ctl stop -D D:\database\inst1



#### See

- Refer to "Hot Standby" in the PostgreSQL Documentation for details on the standby.signal file.
- Refer to "Setting Up a Standby Server" in the PostgreSQL Documentation for details on the primary\_conninfo.



A new timeline is branched for the new primary server due to promotion, so 'latest' needs to be specified for the recovery\_target\_timeline parameter so that the old primary server (new standby server) follows the new primary server.

### 4.1.1.2 Rebuild the Standby Server

The starting of the recovered original primary server as the standby server is referred to as the "standby server rebuild".

On the original primary server, start Mirroring Controller and the instance.

Enabling automatic switch/disconnection

As the instance administrator user, execute the mc\_ctl command in start mode.

Example)

> mc\_ctl start -M D:\mcdir\inst1

#### Disabling automatic switch/disconnection

As the instance administrator user, execute the mc\_ctl command in start mode with the -F option specified.

Example)

> mc\_ctl start -M D:\mcdir\inst1 -F



After Mirroring Controller is started, automatic switch/disconnection can be enabled or disabled using the enable-failover or disable-failover mode of the mc\_ctl command.

.....

### 4.1.1.3 Failback of the Primary Server

To revert the primary server and standby server to the original server configuration after rebuilding the standby server, perform failback for the primary server.

Do this to execute the main job on the previous primary server.

Perform the following procedure:

1. Failback of the primary server

Execute the mc\_ctl command in switch mode on the primary server or the standby server.

Example)

```
> mc_ctl switch -M D:\mcdir\inst1
```

After executing the mc\_ctl command in switch mode, the status will be as follows:

Example)

2. Stop the original primary server

On the original primary server, execute the mc\_ctl command in stop mode to stop Mirroring Controller and the instance.

Example)

```
> mc_ctl stop -M D:\mcdir\inst1
```

3. Create a copy of the new primary server instance in the original primary server (new standby server)

Execute the pg\_basebackup command to create data in the new standby server by synchronizing with the new primary server.

Example)

```
> pg_basebackup -D D:\database\inst1 -X fetch --waldir=E:\transaction\inst1 --progress --verbose -R --dbname="application_name=standbyServerName" -h primaryServerHostName -p primaryServerPortNumber
```



The procedure for copying the new primary server instance to the new standby server is the same as the procedure for setting up the new standby server.

Refer to "2.5.2 Creating, Setting, and Registering the Standby Server Instance", and then perform the recovery.

#### 4. Rebuild the standby server

On the standby server, start Mirroring Controller and the instance.

#### Enabling automatic switch/disconnection

As the instance administrator user, execute the mc\_ctl command in start mode.

Example)

> mc\_ctl start -M D:\mcdir\inst1

#### Disabling automatic switch/disconnection

As the instance administrator user, execute the mc\_ctl command in start mode with the -F option specified.

Example)

> mc\_ctl start -M D:\mcdir\inst1 -F



After Mirroring Controller is started, automatic switch/disconnection can be enabled or disabled using the enable-failover or disable-failover mode of the mc\_ctl command.

# 4.1.2 Operations when the Server has Started Degrading after a Disconnection has Occurred

This section explains the operations when the server has started degrading after a disconnection has occurred.



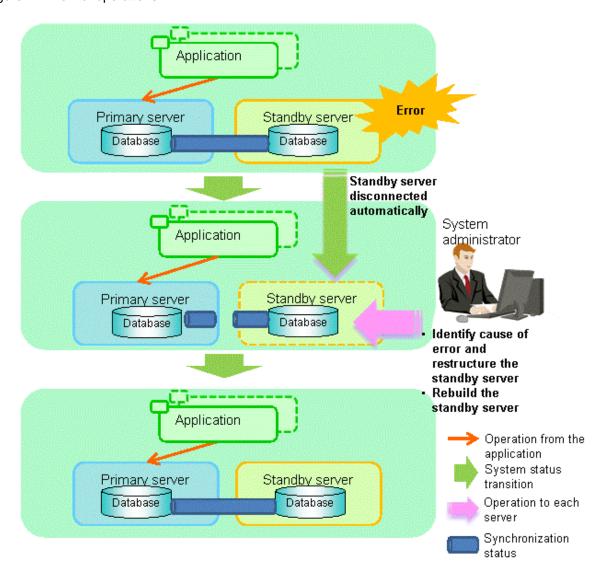
After a disconnection has occurred as a result of an abnormality on the standby server, the database will not have a multiplexed configuration until the standby server is rebuilt. Remove the cause of the error as quickly as possible, and then rebuild the standby server.

If the disconnection occurred and the server has started degrading, perform the following operations to recover the standby server and revert it to its original state:

- Identify Cause of Error and Restore the Standby Server
- Rebuild the Standby Server

The flow of these operations is shown in the figure below.

Figure 4.2 Flow of operations



### 4.1.2.1 Identify Cause of Error and Restore the Standby Server

Perform the recovery according to the following procedure:

- 1. Stop Mirroring Controller
- 2. Recovery of the Mirroring Controller management directory
- 3. Identify cause of error and perform recovery

### 4.1.2.1.1 Stop Mirroring Controller

Execute the mc\_ctl command in stop mode for the standby server on which the error occurred.

Example)

> mc\_ctl stop -M D:\mcdir\inst1

This also stops the instance that is required to perform the recovery.



If the instance does not stop, refer to "Actions in Response to Failure to Stop an Instance" in the Operation Guide, and then stop the instance.

Then, specify the -e option in the above command to forcibly stop Mirroring Controller.

### 4.1.2.1.2 Recovery of the Mirroring Controller management directory

Copy the files in the Mirroring Controller management directory from the backup data, and then perform the recovery.

### 4.1.2.1.3 Identify cause of error and perform recovery

Refer to the event log of the primary server and the standby server to identify the cause of the error, and then perform recovery.

Execute the pg\_basebackup command to perform recovery by synchronizing data in the primary server with the standby server.

Example)

> pg\_basebackup -D D:\database\inst1 -X fetch --waldir=E:\transaction\inst1 --progress --verbose -R --dbname="application\_name=standbyServerName" -h primaryServerHostName -p primaryServerPortNumber



This recovery procedure is the same as the procedure for setting up the standby server.

Refer to "2.5.2 Creating, Setting, and Registering the Standby Server Instance", and then perform the recovery.

### 4.1.2.2 Rebuild the Standby Server

Start the Mirroring Controller and the instance of the standby server, and rebuild the standby server.

Enabling automatic switch/disconnection

As the instance administrator user, execute the mc\_ctl command in start mode.

Example)

> mc\_ctl start -M D:\mcdir\inst1

#### Disabling automatic switch/disconnection

As the instance administrator user, execute the mc\_ctl command in start mode with the -F option specified.

Example)

> mc\_ctl start -M D:\mcdir\inst1 -F



**Point** 

After Mirroring Controller is started, automatic switch/disconnection can be enabled or disabled using the enable-failover or disable-failover mode of the mc\_ctl command.

## 4.1.3 Addressing Errors During Degrading Operation

This section explains how to address errors that may occur on the server on which operation is continuing during degrading operation triggered by a switch or disconnection.

#### If needing to recover from backup data

If it is necessary to recover the database using backup data due to data becoming corrupted from disk failure or user operation error, refer to the following for information on recovery to database multiplexing mode:

- Action Required when All Database Servers or Instances Stopped
- Recovering from an Incorrect User Operation

#### If a temporary error occurs

If a temporary error occurs, such as due to a high load on the server or insufficient system resources, remove the cause of the error and restart Mirroring Controller, and then refer to the following for details on recovery to database multiplexing mode:

- Operations when the Server has Started Degrading after a Switch has Occurred
- Operations when the Server has Started Degrading after a Disconnection has Occurred



Refer to "3.2.1 Starting Mirroring Controller" and "3.2.2 Stopping Mirroring Controller" for information on restarting Mirroring Controller.

## 4.2 Action Required when Automatic Switch Fails

If the system behavior is unstable, for example there are insufficient temporary system resources, the Mirroring Controller automatic switch may fail.

Perform the switch manually using one of the following methods:

- Refer to the procedures in "3.4 Manually Switching the Primary Server".
- In the standby server, execute the mc\_ctl command in switch mode with the -force option specified to forcibly perform the switch.

  Example)

> mc\_ctl switch -M D:\mcdir\inst1 --force



- Even if connection cannot be established between database servers, it is possible to fence the primary server and forcibly switch by executing the mc\_ctl command in switch mode with the --force option specified.

.....

- The primary server is not fenced in the cases below, so stop Mirroring Controller and instances of the primary server database in advance:
  - The --no-fencing option is specified when performing forced switch.
  - The heartbeat\_error\_action parameter in *serverIdentifier*.conf is set to "message" and the fencing command is not configured to be used (the fencing\_command parameter is omitted in *serverIdentifier*.conf).
  - The heartbeat\_error\_action parameter in serverIdentifier.conf is set to "fallback".



#### Recovery to database multiplexing mode

Refer to "4.1.1.2 Rebuild the Standby Server" and "4.1.1.3 Failback of the Primary Server" for information on recovery to database multiplexing mode.

## 4.3 Action Required when Automatic Disconnection Fails

If the system behavior is unstable, for example there are insufficient system resources such as available memory or free disk space, automatic disconnection using Mirroring Controller may not be possible.

Perform the disconnection manually using one of the following methods:

- Refer to the procedures in "3.5 Manually Disconnecting the Standby Server".
- In the primary server, execute the mc\_ctl command in detach mode to perform forced disconnection.

Example)

> mc\_ctl detach -M D:\mcdir\inst1



- Even if connection cannot be established between database servers, it is possible to fence the standby server and forcibly disconnect by executing the mc\_ctl command in detach mode.
- In the cases below, stop Mirroring Controller and instances of the standby server database in advance so that the standby server is not fenced:
  - The --no-fencing option is specified when performing forced disconnection.
  - The heartbeat\_error\_action parameter in *serverIdentifier*.conf is set to "message" and the fencing command is not configured to be used (the fencing\_command parameter is omitted in *serverIdentifier.conf*).
  - The heartbeat\_error\_action parameter in *serverIdentifier*.conf is set to "fallback".



#### Recovery to database multiplexing mode

Refer to "4.1.2.2 Rebuild the Standby Server" for information on recovery to database multiplexing mode.

# 4.4 Action Required when All Database Servers or Instances Stopped

This section explains what happens when all database servers or instances on the database server have stopped, so jobs cannot continue.



See

#### Recovery to database multiplexing mode

Refer to "4.1.1.2 Rebuild the Standby Server" and "4.1.1.3 Failback of the Primary Server" for information on recovery to database multiplexing mode.

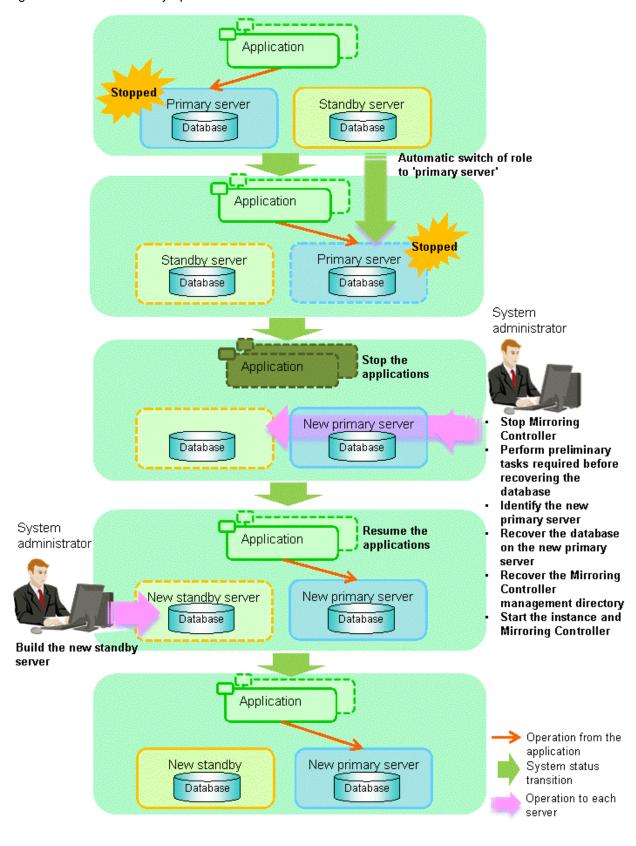
#### Overview of recovery operations

After recovering the database to the state immediately prior to the failure on a specific server comprising the database multiplexing system, restore the system.

In other words, after specifying the server on which the database is to be recovered and then recovering it as the new primary server, configure all other servers as new standby servers.

The flow of these recovery operations is shown in the figure below.

Figure 4.3 Flow of recovery operations



Perform the following procedure.

Stop applications
 Stop running applications.

#### 2. Stop Mirroring Controller

Execute the mc\_ctl command in stop mode on all servers that comprise the database multiplexing system.

Example)

> mc\_ctl stop -M D:\mcdir\inst1



#### Forcibly stop Mirroring Controller

If Mirroring Controller does not stop, execute the mc\_ctl command in stop mode with the -e option specified.

Example)

```
> mc_ctl stop -M D:\mcdir\instl -e
```

3. Perform prerequisite tasks before recovering the database

First, refer to "Actions when an Error Occurs" in the Operation Guide, and then identify the cause of the error and perform recovery of the disk on which the failure occurred, etc.

4. Identify the new primary server

Perform the following operations on all servers comprising the database multiplexing system, and check the server containing the backup data that shows the latest date. This server will become the new primary server, on which the database is to be recovered.

Example)

In the example below, the pgx\_rcvall command is executed with the -l option specified and the backup data that shows the latest date is identified.

5. Recover the database on the new primary server

Recover the database using the recovery method that uses the pgx\_rcvall command based on the backup data.

a. Perform the following operations on all servers comprising the database multiplexing system, and check the server containing the archive log and mirrored transaction log that show the latest date.

```
Example)
```

In the example below, the archive log and mirrored transaction log that show the latest date are identified.

```
> dir /OD <backupDataStorageDir>\*_wal
```

b. If the server containing the latest archive log and mirrored transaction log is different to the new primary server identified in step 4, all files and directories under the directory shown below are copied and written to the backup storage destination directory on the new primary server.

Deployment destination directory of the archive log and mirrored transaction log

```
<backupDataStorageDir>\*_wal
```

c. Execute the pgx\_rcvall command on the new primary server, specifying the backup storage destination directory of the new primary server.

```
Example)
```

In the example below, the pgx\_rcvall command is executed with the -B option specified.

```
> pgx_rcvall -B E:\backup\inst1 -D D:\database\inst1
```



Refer to "Actions when an Error Occurs" in the Operation Guide for information on the pgx\_rcvall command.

6. Recover the Mirroring Controller management directory

Copy the files in the Mirroring Controller management directory from the backup data on the new primary server, and then perform the recovery.

7. Start the instance and Mirroring Controller

Start the instance and Mirroring Controller on the new primary server.

Enabling automatic switch/disconnection

As the instance administrator user, execute the mc\_ctl command in start mode.

Example)

> mc\_ctl start -M D:\mcdir\inst1

#### Disabling automatic switch/disconnection

As the instance administrator user, execute the mc\_ctl command in start mode with the -F option specified.

Example)

> mc\_ctl start -M D:\mcdir\inst1 -F



After Mirroring Controller is started, automatic switch/disconnection can be enabled or disabled using the enable-failover or disable-failover mode of the mc\_ctl command.

8. Resume applications

Resume execution of applications.

9. Build the new standby server

Refer to "2.5 Setting Up the Standby Server" for information on building (setting up) a standby server from the primary server.



It is not necessary to repeat steps that have already been performed, such as registering to Windows services.

### 4.5 Recovering from an Incorrect User Operation

This section describes how to recover an instance when data has been corrupted due to incorrect user operation.

For example, when data has been corrupted due to incorrect user operation, such as data being unintentionally changed or deleted by an application or command, it is necessary to restore the original data on the primary server and resynchronize with the standby server.

Use the following procedure to perform recovery.

1. Identify the primary server

Execute the mc\_ctl command in status mode on each server, and search for a server for which "primary" or "none(inactivated primary)" is displayed.

2. Stop the applications and commands that caused the incorrect operation to occur

Stop applications and commands that are running on the primary server. This will minimize the impact caused by the incorrect data.

Also, if any applications used for reference by the standby server are running, stop them too.

3. Stop the instance and Mirroring Controller

Stop the instance and Mirroring Controller on both the primary server and standby server.

Example)

\$ mc\_ctl stop -a -M D:\mcdir\inst1

4. Recover the database on the primary server

Recover the database using the recovery method in which the pgx\_rcvall command uses the backup data to recover the database to a restore point prior to the time when the incorrect operation was performed.



#### See

Refer to "Recovering from an Incorrect User Operation" in the Operation Guide for information on using the pgx\_rcvall command to recover the database to a restore point, and then perform only the database recovery procedure while the instance is in a stop state.

5. Start the instance and Mirroring Controller

Start the instance and Mirroring Controller on the primary server.

Enabling automatic switch/disconnection

As the instance administrator user, execute the mc\_ctl command in start mode.

Example)

> mc\_ctl start -M D:\mcdir\inst1

#### Disabling automatic switch/disconnection

As the instance administrator user, execute the mc\_ctl command in start mode with the -F option specified.

Example)

mc\_ctl start -M D:\mcdir\inst1 -F



### Point

After Mirroring Controller is started, automatic switch/disconnection can be enabled or disabled using the enable-failover or disablefailover mode of the mc\_ctl command.

6. Build the new standby server

Refer to "2.5 Setting Up the Standby Server" for information on building (setting up) a standby server from the primary server.



It is not necessary to repeat steps that have already been performed, such as registering to Windows services.

# Chapter 5 Managing Mirroring Controller Using WebAdmin

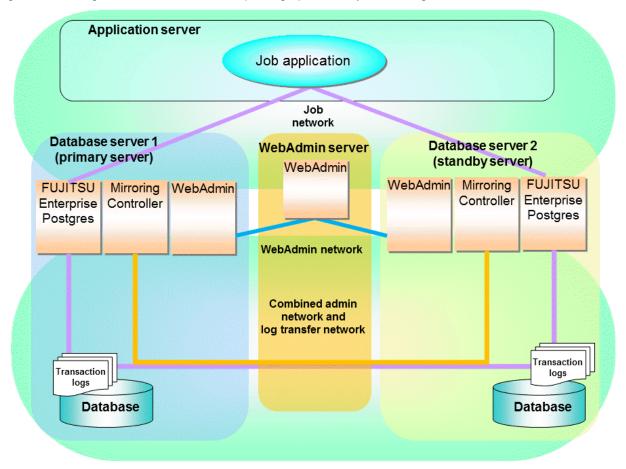
This chapter describes how to set up and manage Mirroring Controller in a streaming replication cluster using WebAdmin.

Mirroring Controller can be used to monitor a streaming replication cluster and perform automatic switching or disconnect synchronous replication when there is an error.

WebAdmin can be used to set up Mirroring Controller in an existing replication cluster. Mirroring Controller can be set up for either synchronous standby instances or asynchronous standby instances.

The configuration of the database multiplexing system built using WebAdmin is shown below:

Figure 5.1 Configuration of database multiplexing operation system using WebAdmin





- If Mirroring Controller is set up to the replication cluster using WebAdmin, the network with the host name (or IP address) specified in [Host name] will be used as the admin network and the log transfer network.
- To use a network other than the job network as the log transfer network, before building the replication cluster specify a host name other than the job network one in [Host name].



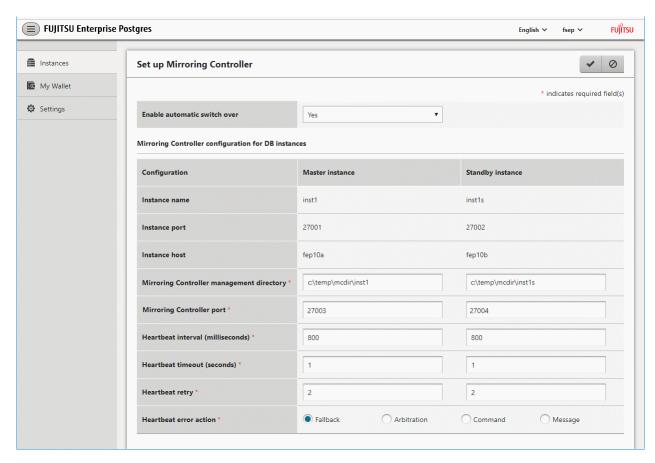
If you set up the arbitration server using WebAdmin, install WebAdmin on the arbitration server.

## **5.1 Mirroring Controller Setup**

Perform the following procedure to set up Mirroring Controller in a streaming replication cluster.

- 1. In the [Instances] tab, select the standby instance on which Mirroring Controller needs to be set up.
- 2. Click
- 3. Enter the information for the Mirroring Controller to be set up.

In the example below, Mirroring Controller is being set up for the replication cluster having master instance "inst1" and standby instance "inst1s".



The instance name, host address and port of the master and standby instances are displayed for easy reference.

Enter the following items on master instance and on standby instance fields for Mirroring Controller setup, as shown in the above screenshot:

- [Enable automatic switch over]: Toggles the automatic switch/disconnection functionality. Select "Yes". The default is "No".
- [Mirroring Controller management directory]: Directory where the Mirroring Controller configuration files will be stored. When the [Mirroring Controller management directory] is entered, WebAdmin will search the Mirroring Controller configuration files in the entered directory based on the [Data storage path] of the corresponding DB instance. If Mirroring Controller configuration files are found, the Mirroring Controller fields will be auto filled.
- [Mirroring Controller port]: Port number of Mirroring Controller. Note that if the Windows firewall feature is enabled, you must enable the port number of Mirroring Controller. Refer to "E.2 Windows Firewall Settings" for details.
- [Heartbeat interval (milliseconds)]: Number of milliseconds between two consecutive heartbeat checks. The default is "800".
- [Heartbeat timeout (seconds)]: Number of seconds for the heartbeat timeout. The default is "1".
- [Heartbeat retry]: Number of retries for heartbeat monitoring, before failover occurs. The default is "2".

- [Heartbeat error action]: Operation when a heartbeat abnormality is detected. The default is "Fallback".

When using FUJITSU Enterprise Postgres 10 and 11 instances created with previous versions, the instances will be in compatibility mode, and the "Fallback" is preselected and cannot be changed in the [Heartbeat error action] for Mirroring Controller setup.

When setting up Mirroring Controller for FUJITSU Enterprise Postgres 9.5 and 9.6 instances, the [Heartbeat error action] is not supported and therefore is not displayed.

When the [Heartbeat error action] is set to "Arbitration", the following extra items are displayed:

- [Arbitration network IP address]: IP address of the arbitration network.
- [Mirroring Controller Arbitration port]: Port number of Mirroring Controller for communicating with the arbitration server.

The [Arbitration server configuration] section is also displayed with the following items. The [Arbitration server configuration] will not be auto filled.

- [Location]: Location of the arbitration server. "Local" or "Remote" can be selected depending on your configuration.

If the arbitration server and WebAdmin server are located on the same server, you can select "Local" and the following items are displayed:

- [Arbitration management directory]: Directory where the arbitration server configuration files will be stored.
- [Arbitration server host or IP address]: Host name or IP address of the arbitration server.
- [Arbitration process port]: Port number for the arbitration process.
- [Fencing command]: Full path of the fencing command that fences a database server when an abnormality is detected.

If "Remote" is set for the item, the items below are displayed in addition to the above items.

- In the [Arbitration server configuration] section, [Operating system credential] is displayed where you can enter the following information:

[User name]: User name to access the arbitration server.

[Password]: Password to access the arbitration server.

- In the [Remote WebAdmin for Arbitration server] section, the following items are displayed:

[Remote WebAdmin address]: IP address of the remote WebAdmin installed on the arbitration server.

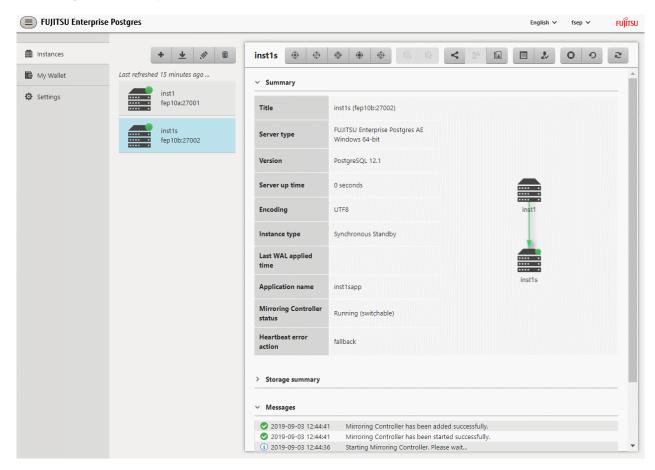
[Remote WebAdmin port]: Port number for the WebAdmin installed on the arbitration server.

When the [Heartbeat error action] is set to "Command", the following extra items are displayed:

- [Arbitration command]: Full path of the arbitration command to be executed when an abnormality is detected.
- [Fencing command]: Full path of the fencing command that fences a database server when an abnormality is detected.
- 4. Click to set up Mirroring Controller.

5. Upon successful completion, Mirroring Controller will be started on master and standby instances.

In the [Instances] tab, select standby instance. The page below is displayed, in which you can check the Mirroring Controller status. In the example below, standby instance "inst1s" is used.



After the Mirroring Controller has been set up, ([Edit Mirroring Controller] button) and ([Mirroring Controller Configuration] button) are available. These buttons are displayed only when FUJITSU Enterprise Postgres 10, 11 and 12 instances are created with FUJITSU Enterprise Postgres WebAdmin 12.

For FUJITSU Enterprise Postgres 9.5 and 9.6 instances, the [Heartbeat error action] will not be displayed.

When the [Heartbeat error action] is "Arbitration", the following information is displayed: whether the arbitration status is "online" or "offline", the arbitration server IP address and the arbitration process port.



Operating system credential (User name, Password) should not contain hazardous characters. Refer to "Appendix F WebAdmin Disallow User Inputs Containing Hazardous Characters".

# 5.2 Edit Mirroring Controller Setup

Settings made in "5.1 Mirroring Controller Setup" can be updated in either the master instance or a standby instance using WebAdmin. Perform the following procedure to edit Mirroring Controller configuration:

- 1. In the [Instances] tab, select the instance for which the Mirroring Controller configuration is to be edited.
- 2. Click .

- 3. Enter the information for the Mirroring Controller to be updated. Refer to "5.1 Mirroring Controller Setup".
- 4. Click to update the Mirroring Controller.
- 5. Upon successful completion, Mirroring Controller will be started on master and standby instances.

Editing and saving the [Edit Mirroring Controller] page will reset all other settings that are not listed on this page to default values.

### 5.3 Mirroring Controller Configuration

The information related to Mirroring Controller monitoring and control (refer to "A.4.1 Server Configuration File for the Database Servers") and the information related to arbitration and control of the Mirroring Controller arbitration process (refer to "A.4.2 Arbitration Configuration File") can be set using WebAdmin. You can view and update the configuration on either the master instance or the standby instance.

Perform the following procedure:

- 1. In the [Instances] tab, select the instance for the Mirroring Controller configuration you want to view.
- 2. Click to view the Mirroring Controller configuration.
- 3. Click to show the editing page for the Mirroring Controller configuration. The Mirroring Controller configurations defined during [Mirroring Controller Setup] are read-only on this page. Refer to "5.1 Mirroring Controller Setup".

Additionally, refer to the "Appendix A Parameters" for information about the settings and the corresponding parameter names.

The items common to all [Heartbeat error action] are:

- Target DB
- Core file path
- Remote call timeout (milliseconds)
- Agent alive timeout (seconds)
- DB instance check interval (milliseconds)
- DB instance check timeout (seconds)
- DB instance check retry
- DB instance timeout action
- Disk check interval (milliseconds)
- Disk check retry
- Tablespace directory error action
- Post-switch command
- Post-promote command

(Post-promote command is replaced in FUJITSU Enterprise Postgres 12. The Post-promote command is still valid and will be displayed when it is used in the server configuration file of Mirroring Controller.)

- Post-attach command
- Pre-detach command
- State transition command timeout (seconds)
- Check synchronous standby names validation

When the [Heartbeat error action] is set to "Arbitration", the following extra items are displayed:

- Arbitration timeout (seconds)
- Arbiter alive interval (milliseconds)

- Arbiter alive retry
- Arbiter alive timeout (seconds)
- Arbiter connect interval (milliseconds)
- Arbiter connect timeout (seconds)
- Fencing command
- Fencing command timeout (seconds)
- Shutdown detached synchronous standby

When the [Heartbeat error action] is set to "Arbitration", the [Arbitration server configuration] section is displayed with the following items:

- Core file path
- Fencing command timeout (seconds)
- Heartbeat interval (milliseconds)
- Heartbeat timeout (seconds)
- Heartbeat retry

When the [Heartbeat error action] is set to "Command", the following extra items are available:

- Fencing command timeout (seconds)
- Arbitration command timeout (seconds)
- Shutdown detached synchronous standby

When the [Heartbeat error action] is set to "Message", the following extra items are available:

- Fencing command
- Fencing command timeout (seconds)

In addition, the following configurations are provided:

- DB instance JDBC connection SSL parameters
- DB instance libpq connection SSL parameters
- 4. Click to update the Mirroring Controller configurations.

# 5.4 Stopping Mirroring Controller

Mirroring Controller can be stopped either in master instance or in standby instance using WebAdmin.

Perform the following procedure to stop Mirroring Controller.

- 1. In the [Instances] tab, select the instance where to stop Mirroring Controller.
- 2. Click
- 3. In the confirmation dialog box, click [Yes].

Mirroring Controller will be stopped on the selected instance. The Mirroring Controller status will be updated, and a confirmation message entry will be displayed in the [Message] section.

### 5.5 Starting Mirroring Controller

Mirroring Controller can be started either in master instance or in standby instance using WebAdmin.

Perform the following procedure to start Mirroring Controller.

- 1. In the [Instances] tab, select the instance where to start Mirroring Controller.
- 2. Click
- 3. In the confirmation dialog box, select the desired failover mode, and then click [Yes].

Mirroring Controller will be started on the selected instance. The Mirroring Controller status will be updated, and a confirmation message entry will be displayed in the [Message] section.

# 5.6 Disabling Failover Mode

Disabling failover mode in Mirroring Controller disables automatic switch/disconnection between master and standby instances.

Perform the following procedure to disable failover mode.

- 1. In the [Instances] tab, select the instance.
- 2. Click
- 3. In the confirmation dialog box, click [Yes].

Failover mode will be disabled in Mirroring Controller. The Mirroring Controller status will be updated and a confirmation message entry will be displayed in the [Message] section.

## 5.7 Enabling Failover Mode

Enabling failover mode in Mirroring Controller enables automatic switch/disconnection between master and standby instances.

Perform the following procedure to enable failover.

- 1. In the [Instances] tab, select the instance.
- 2. Click (F).
- 3. In the confirmation dialog box, click [Yes].

Failover mode will be enabled in Mirroring Controller. The Mirroring Controller status will be updated and a confirmation message entry will be displayed in the [Message] section.

# 5.8 Deleting Mirroring Controller Setup

Deleting Mirroring Controller setup removes its setup from master and standby instances.

- 1. In the [Instances] tab, select the instance.
- 2. Click
- 3. In the confirmation dialog box, click [Yes].

Mirroring Controller setup will be removed from the cluster. The cluster status will be updated and a confirmation message entry will be displayed in the [Message] section.

For the instances in FUJITSU Enterprise Postgres 12, WebAdmin does not delete the Mirroring Controller management directory and the configuration files.

### 5.9 Status Update after Failover

When Mirroring Controller performs a failover, standby instance will be promoted to standalone instance. The Mirroring Controller setup will be removed from both standby and master instances.

The following scenario describes one of the ways in which failover can be triggered, and the results achieved by the use of Mirroring Controller in WebAdmin.

- 1. In the [Instances] tab, select the master instance "inst1".
- 2. Click .
- 3. In the confirmation dialog box, the warning "This instance is being monitored by Mirroring Controller. Stopping the instance may result in cluster failover." is displayed.
- 4. Choose the stop mode and click [Yes].

In the server, the following takes place:

- a. The master instance is stopped.
- b. Failover is triggered in Mirroring Controller.
- c. The Mirroring Controller setup is removed from both master and standby instances
- d. Standby instance is promoted to standalone.
- 5. When the instance is refreshed in WebAdmin, the latest status of the instances will be displayed.



When failover is performed, the Mirroring Controller setup is removed from both master and standby instances. Therefore, to manage the Mirroring Controller using WebAdmin again, create the standby instance and set up Mirroring Controller.

Refer to "Creating a Standby Instance" in the Operation Guide for details.

Refer to "5.1 Mirroring Controller Setup" for details.

# 5.10 Action Required when an Error Occurs in the Combined Admin Network and Log Transfer Network

Communication errors may temporarily occur in the network used as the admin network and log transfer network due to reasons such as high load on the server or insufficient system resources. Because of this, there is a risk of causing a split-brain situation by mistake even though the server has no issues.

Split brain is a phenomenon in which both servers temporarily operate as primary servers, causing data updates to be performed on both servers.

#### How to detect split brain using WebAdmin

If the conditions below are met, split brain may occur. Refer to "Split-brain detection method" and "How to recover from a split-brain" in "Appendix D Notes on Performing Automatic Degradation Immediately after a Heartbeat Abnormality" and take the actions described.

- 1. A standby instance is selected in the [Instances] tab, and
- 2. "Standalone" is displayed in [Instance type], and
- 3. A master instance is selected in the [Instances] tab, and
- 4. "Standalone" is displayed in [Instance type].



The admin network is important because Mirroring Controllers use it to confirm the status of each server.

The log transfer network is also important to maintain the data freshness.

Therefore, use network configurations resistant to faults for these networks by using the network redundancy channel bonding feature provided by the operating system or network driver vendor.

# 5.11 Performing Automatic Degradation Using the Arbitration Server

If database multiplexing is performed using WebAdmin, it is also possible to perform automatic degradation using the arbitration server. In such cases, it is necessary to perform tasks on the database server and the arbitration server after setting up Mirroring Controller in WebAdmin.

#### Tasks on the arbitration server

Perform setup of the arbitration server using Mirroring Controller commands.

1. Set up the arbitration server.

Refer to "2.3 Setting Up the Arbitration Server" in "Chapter 2 Setting Up Database Multiplexing Mode" for information on how to set up the arbitration server.

#### Tasks on the database server

Change some of the settings after setting up Mirroring Controller in WebAdmin.

1. Set up Mirroring Controller in WebAdmin.

Refer to "5.1 Mirroring Controller Setup" for details.

2. Use WebAdmin to stop Mirroring Controller on the master and standby instances.

Refer to "5.4 Stopping Mirroring Controller" for details.

3. Edit the network configuration file of the master and standby instances, and add the arbitration server information.

The network configuration file is network.conf, which exists in the Mirroring Controller management directory specified during Mirroring Controller setup. Refer to "A.3 Network Configuration File" for details.

A definition example of network.conf is shown below.

#### Example:

The port number of the database server to be used as the arbitration network is set to "27541". The ID of the server of the Mirroring Controller arbitration process is set to "arbiter", and its port number is set to "27541".

```
dbsvm27500 192.0.2.100,192.0.3.100 27540,27541 server dbsvs27500 192.0.2.110,192.0.3.110 27540,27541 server arbiter 192.0.3.120 27541 arbiter
```



- Ensure that the port numbers set for the database server and the arbitration server do not conflict with other software. In addition, do not configure the same segment for the admin network and the arbitration network.
- If the server type is "server", two IP addresses or host names, and two port numbers need to be specified in the following order:
  - IP address or host name of the database server used as the admin network
  - IP address or host name of the database server used as the arbitration network
  - Port number of the database server used as the admin network
  - Port number of the database server used as the arbitration network

- If the server type is "arbiter", specify the IP address or host name set for the my\_address parameter and the port number set for the port parameter in arbitration.conf of the arbitration server.
- 4. Edit the server configuration file of the master and standby instances, and add the parameters required for automatic degradation using the arbitration server.

The server configuration file is *instanceName*.conf or *instancePort*.conf, which exists in the Mirroring Controller management directory specified during Mirroring Controller setup.

To perform automatic degradation using the arbitration server, set the heartbeat\_error\_action parameter to "arbitration".

Refer to "A.4.1 Server Configuration File for the Database Servers" for information on other parameters.

5. Use WebAdmin to start Mirroring Controller on the master and standby instances.

Refer to "5.5 Starting Mirroring Controller" for details.

#### Common tasks

1. Use the Mirroring Controller command to check the connection status from the database server or the arbitration server.

Refer to "2.8 Checking the Connection Status" for information on how to check the connection status.

# Appendix A Parameters

This appendix describes the configuration files and parameters required by the database multiplexing mode.



Refer to "Server Configuration" in the PostgreSQL Documentation for information on the postgresql.conf file.

# A.1 Parameters Set on the Primary Server

The content for the parameters set in the postgresql.conf file of the primary server is shown in the table below.

Table A.1 postgresql.conf file

Parameter	Value set	Explanation	
wal_level	replica or logical	Specify the output level for the transaction log.	
		Specify "logical" when logical decoding is also to be used	
max_wal_senders	2 or more	Specify "2" when building a Mirroring Controller cluster system.	
		When additionally connecting asynchronous standby servers to the cluster system, add the number of simultaneous connections from these standby servers.	
synchronous_standby_names	'standbyServerName'	Use single quotation marks (') to enclose the name that will identify the standby server. Any name can be specified.	
		Do not change this parameter while Mirroring Controller is running.	
		Do not specify multiple names to this parameter as the Mirroring Controller can manage only one standby server	
hot_standby	on	Specify whether queries can be run on the standby server	
		Specify this to execute reference jobs on the standby server.	
wal_keep_segments	Number of file segments	If a delay exceeding the value set in this parameter occurs, the WAL segment required later by the primary server may be deleted.	
		Additionally, if you stop a standby server (for maintenance, for example), consider the stop time and set a value that will not cause the WAL segment to be deleted.	
		Refer to "Estimating Transaction Log Space Requirements" in the Installation and Setup Guide for Server for information on estimating the WAL segment.	
wal_log_hints	on	When using the pg_rewind command to recover a standby server, specify this parameter or enable checksums when executing the initdb command.	
wal_sender_timeout	Timeout (milliseconds)	Specify the time period after which it is determined that the receiver process (walreceiver) of the transaction log is in an abnormal state on the primary server.	
		The specified value must be larger than the value set for the wal_receiver_status_interval parameter set in the postgresql.conf file of the standby server.	

Parameter	Value set	Explanation
		By aligning this value with the value for the database process heartbeat monitoring time, you can unify the time after which it is determined that an error has occurred.
wal_receiver_timeout	Timeout (milliseconds)	Specify the time period after which it is determined that an error has occurred when the transaction log was received on the standby server.
		By aligning this value with the value for the database process heartbeat monitoring time, you can unify the time after which it is determined that an error has occurred.
archive_mode	on	Specify the archive log mode.
archive_command	'cmd /c ""installDir\\bin\ \pgx_walcopy.cmd" "%p"  "backupDataStorageDestinationDir \\archived_wal\\%f""	Specify the command and storage destination to save the transaction log.
backup_destination	Backup data storage destination directory	Specify the name of directory where to store the backup data.
		Set the permissions so that only the instance administrator user can access the specified directory.
		Specify the same full path on all servers, so that the backup data of other servers can be used to perform recovery.
listen_addresses	Primary server IP address, host name, or "*"	Specify the IP address or host name of the primary server. Specify the IP address or corresponding host name that will be used to connect to the log transfer network.
		The content specified is also used to allow connections from client applications.
		To receive the connection and the transaction log from any client or standby server, specify "*".
		Refer to "Connections and Authentication" in the PostgreSQL Documentation for details.
connections to the instance + connections fro		The value specified is also used to restrict the number of connections from client applications and the number of connections for the management of instances.
	value	Refer to "When an Instance was Created with the initdb Command" in the Installation and Setup Guide for Server, and "Connections and Authentication" in the PostgreSQL Documentation, for details.
superuser_reserved_connection s	Add the number of simultaneous executions of mc_ctl status (*1) + 2	Specify the number of connections reserved for connections from database superusers.
		Add the number of connections from Mirroring Controller processes. Also reflect the added value in the max_connections parameter.
restart_after_crash	off	If "on" is specified, or the default value is used for this parameter, behavior equivalent to restarting FUJITSU Enterprise Postgres, including crash recovery, will be performed when some server processes end abnormally.
		However, when database multiplexing monitoring is used, a failover will occur after an error is detected when some server processes end abnormally, and the restart of those server processes is forcibly stopped. Specify "off" to

Parameter	Value set	Explanation	
		prevent behavior such as this from occurring for no apparent reason.	
synchronous_commit	on or remote_apply	Specify up to what position WAL send is to be performed before transaction commit processing returns a normal termination response to a client.  The recommended value is "on" or "remote_apply" to prevent data loss caused by operating system or server down immediately after a switch or switch.	
recovery_target_timeline	latest	Specify "latest" so that the new standby server (original primary server) will follow the new primary server when a switch occurs.  This parameter is required when the original primary server is incorporated as a new standby server after the primary server is switched.	

<sup>\*1:</sup> Number of simultaneous executions of the mc\_ctl command in the status mode.

# A.2 Parameters Set on the Standby Server

This section explains the content of the file and parameters set on the standby server. After editing postgresql.conf file, start the instance.

The content for the parameters specified in postgresql.conf file is shown in the table below.

Table A.2 postgresql.conf file

Parameter	Value set	Explanation	
wal_level	replica or logical	Specify the output level for the transaction log.	
		Specify "logical" when logical decoding is also to be used.	
max_wal_senders	2 or more	Specify "2" when building a Mirroring Controller cluster system.	
		When additionally connecting asynchronous standby servers to the cluster system, add the number of simultaneous connections from these standby servers.	
synchronous_standby_names	'primaryServerName'	Use single quotation marks (') to enclose the name that will identify the primary server. Any name can be specified.	
		This name will be required to rebuild the original primary server as the new standby server after the primary server was switched.	
		Do not change this parameter while Mirroring Controller is running.	
		Do not specify multiple names to this parameter as the Mirroring Controller can manage only one standby server.	
hot_standby	on	Specify whether queries can be run on the standby server.	
		Specify this to execute reference type jobs on the standby server.	

Parameter	Value set	Explanation	
wal_keep_segments	Number of file segments	If a delay exceeding the value set in this parameter occurs, the WAL segment required later by the primary server may be deleted.	
		Additionally, if you stop a standby server (for maintenance, for example), consider the stop time and set a value that will not cause the WAL segment to be deleted.	
		Refer to "Estimating Transaction Log Space Requirements" in the Installation and Setup Guide for Server for information on estimating the WAL segment.	
wal_log_hints	on	When using the pg_rewind command to recover a standby server, specify this parameter or enable checksums when executing the initdb command.	
wal_sender_timeout	Timeout (milliseconds)	Specify the time period after which it is determined that the receiver process (walreceiver) of the transaction log is in an abnormal state on the primary server.	
		The specified value must be larger than the value set for the wal_receiver_status_interval parameter set in the postgresql.conf file of the standby server.	
		By aligning this value with the value for the database process heartbeat monitoring time, you can unify the time after which it is determined that an error has occurred.	
wal_receiver_timeout	Timeout (milliseconds)	Specify the time period after which it is determined that an error has occurred when the transaction log was received on the standby server.	
		By aligning this value with the value for the database process heartbeat monitoring time, you can unify the time after which it is determined that an error has occurred.	
backup_destination	Backup data storage destination	Specify the name of the backup data storage directory.	
	directory	Set the permissions so that only the instance administrator user can access the specified directory.	
		Specify the same full path on all servers so that the backup data of other servers can be used to perform recovery.	
archive_mode	on	Specify the archive log mode.	
archive_command	'cmd /c ""installDir\\bin\ \pgx_walcopy.cmd" "%p"  "backupDataStorageDestinationD ir\\archived_wal\\%f""	Specify the command and storage destination to save the transaction log.	
listen_addresses	Standby server IP address, host name, or "*"	Specify the IP address or host name of the standby server. Specify the IP address or corresponding host name that will be used to connect to the log transfer network.	
		The content specified is also used to allow connections from client applications.	
		To receive the connection and the transaction log from any client or standby server, specify "*".	

Parameter	Value set	Explanation	
		Refer to "Connections and Authentication" in the PostgreSQL Documentation for details.	
max_connections	Number of simultaneous client connections to the instance + superuser_reserved_connections	The value specified is also used to restrict the number of connections from client applications and the number of connections for the management of instances.	
	value	Refer to "When an Instance was Created with the initdb Command" in the Installation and Setup Guide for Server, and "Connections and Authentication" in the PostgreSQL Documentation, for details.	
superuser_reserved_connections	Add the number of simultaneous executions of mc_ctl status (*1)	Specify the number of connections reserved for connections from database superusers.	
	+2	Add the number of connections from Mirroring Controller processes. Also reflect the added value in the max_connections parameter.	
restart_after_crash	off	If "on" is specified, or the default value is used for this parameter, behavior equivalent to restarting FUJITSU Enterprise Postgres, including crash recovery, will be performed when some server processes end abnormally.	
		However, when database multiplexing monitoring is used, a failover will occur after an error is detected when some server processes end abnormally, and the restart of those server processes is forcibly stopped. Specify "off" to prevent behavior such as this from occurring for no apparent reason.	
synchronous_commit	on or remote_apply	Specify up to what position WAL send is to be performed before transaction commit processing returns a normal termination response to a client.	
		The recommended value is "on" or "remote_apply" to prevent data loss caused by operating system or server down immediately after a switch or switch.	
primary_conninfo	'streamingReplication ConnectionDestinationInfo'	Use single quotation marks (') to enclose the connection destination information of the streaming replication.	
		The default value of this parameter is automatically set to postgresql.auto.conf in the procedure to run pg _ basebackup for instance setup.	
recovery_target_timeline	latest	Specify "latest" so that the new standby server (original primary server) will follow the new primary server when a switch occurs.	
		This parameter is required when the original primary server is incorporated as a new standby server after the primary server is switched.	

# A.3 Network Configuration File

This section explains the network configuration file (network.conf) to be defined individually for the database servers and the arbitration server. Define the same content on the primary server and standby server.

For database multiplexing mode, define the network configuration for the following in network.conf.

- Integration between Mirroring Controller processes
- Integration between a Mirroring Controller process and the Mirroring Controller arbitration process

#### Items to be defined in network.conf

#### Format:

```
serverIdentifier hostName[,hostName] portNum[,portNum] [serverType]
Or,
serverIdentifier ipAddr[,ipAddr] portNum[,portNum] [serverType]
```

Specify the server identifier, IP address or host name, port number, and server type, using a space as the delimiter.

The items are explained in the table below.

#### Table A.3 network.conf file

Item	Description		
serverIdentifier	Specify any identifier for the server. The maximum length is 64 bytes. Use ASCII characters excluding spaces and number signs (#) to specify this parameter.		
ipAddrOrHostName	Specify the IP address or its corresponding host name that will connect to the admin network that performs communication between the database servers, and to the arbitration network that performs communication between a database server and the arbitration server.		
	When specifying two IP addresses or host names delimited by a comma, do not insert a space after the comma.		
	Use ASCII characters excluding spaces to specify the host name.		
portNum	A port number cannot be specified if it exceeds the range 0 to 65535. Ensure that the port number does not conflict with other software. Do not specify an ephemeral port that may temporarily be assigned by another program. Note that the value specified in this parameter must also be set in the services file.		
	When specifying two port numbers delimited by a comma, do not insert a space after the comma.		
serverType	Specify "server" for a database server ("server" can be omitted), or "arbiter" for the arbitration server.		

#### Content to be defined on the database servers

This section explains the network.conf content to be defined on the database servers.

The content to be defined depends on the operation settings at the time a heartbeat abnormality is detected.

When automatic degradation by the arbitration server is selected

- Specify definitions related to the admin network and arbitration network.
- Specify the IP address or host name and port number according to the server type (database server or arbitration server) as shown in the table below.

Server type	IP address or host name		Port number	
	First	Second	First	Second
server	IP address or host name used as the admin network	IP address or host name used as the arbitration network (*1)	Port number used as the admin network	Port number used as the arbitration network (*1)
arbiter	IP address or host name of the arbitration server  Specify the same value as that specified in the my_address parameter of arbitration.conf on the arbitration server.	Not required	Port number on the arbitration server  Specify the same value as that specified in the port parameter of arbitration.conf on the arbitration server.	Not required

\*1: This value can be omitted from definitions not related to the local server. If it is omitted, network.conf must be created on both the primary server and standby server.

#### Example)

IPv4

```
server1 192.0.2.100,192.0.3.100 27540,27541 server
server2 192.0.2.110,192.0.3.110 27540,27541 server
arbiter 192.0.3.120 27541 arbiter
```

IPv6

```
server1 2001:258:8404:1217:250:56ff:fea7:559f,2001:258:8404:1217:250:56ff:fea8:559f 27540,27541 server server2 2001:258:8404:1217:250:56ff:fea8:55a0 27540,27541 server arbiter 2001:258:8404:1217:250:56ff:fea8:55a0 27541 arbiter
```

When operation other than automatic degradation by the arbitration server is selected

- Specify definitions related to the admin network.
- Define the same content on the primary server and standby server.
- Define lines for database servers only.
- Specify only one IP address or host name and port number.

IP address or host name		Port number	
First	Second	First	Second
IP address or host name to be used as the admin network	Not required	Port number used as the admin network	Not required

### Example)

The literal space represents a space.

IPv4

```
server1 192.0.2.100 27540
server2 192.0.2.110 27540
```

IPv6

```
server1 2001:258:8404:1217:250:56ff:fea7:559f 27540
server2 2001:258:8404:1217:250:56ff:fea7:55a0 27540
```

#### Content to be defined on the arbitration server

This section explains the network.conf content to be defined on the arbitration server.

- Specify definitions related to the arbitration network.
- Define lines for database servers only.
- For the IP address or host name, specify the same value as the second IP address or host name specified in the database server line in network.conf of the database server.

- For the port number, specify the same value as the second port number specified in the database server line in network.conf of the database server.

#### Example)

The literal space represents a space.

IPv4

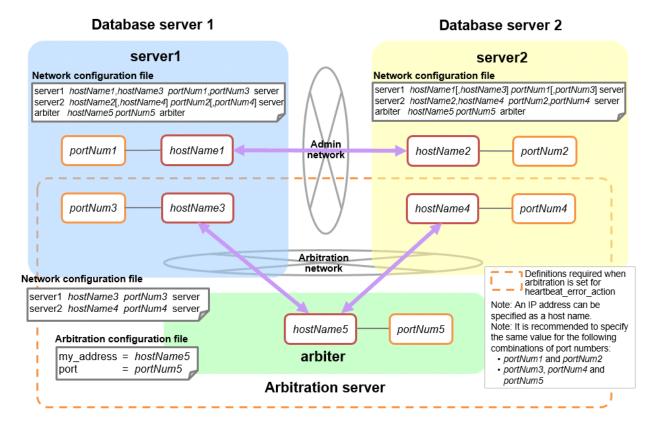
```
server1 192.0.3.100 27541
server2 192.0.3.110 27541
```

IPv6

```
server1 2001:258:8404:1217:250:56ff:fea8:559f 27541
server2 2001:258:8404:1217:250:56ff:fea8:55a0 27541
```

#### Relationship between network-related definitions

Refer to the diagram below for the relationship between the host names and IP addresses or port numbers specified in the network configuration file (network.conf) and arbitration configuration file (arbitration.conf).



# A.4 Server Configuration File

# A.4.1 Server Configuration File for the Database Servers

Define the information related to Mirroring Controller monitoring and control in the serverIdentifier.conf file. The maximum length of the server identifier is 64 bytes. Use ASCII characters excluding spaces to specify this parameter.

If the primary server and standby server environments are different, define content that is different, according to the environment.

Table A.4 serverIdentifier.conf file

Parameter	Value set	Explanation
db_instance	'dataStorageDestinationDir'  [Example]	Use halfwidth single quotation marks (') to enclose the data storage destination directory that will identify the monitoring target instance.
	db_instance = 'D:\\database1\\inst1'	Use ASCII characters, and specify "\\" as the path delimiter.
target_db	postgres or template1	Specify the name of the database to be connected to the database instance. The default is "postgres".
db_instance_username	'usernameToConnectToDbInstance'	Specify the username to connect to the database instance.
		Use ASCII characters to specify this parameter.
		Specify this parameter if the database administrator user is different from the operating system user who starts Mirroring Controller.
		Enclose the username of the database superuser in single quotation marks (').
		The maximum length of the username is 63 bytes.
		The default is the operating system user who starts Mirroring Controller.
db_instance_password	'passwordOfInstanceAdminUser'	Specify the password used when Mirroring Controller connects to a database instance, enclosed in single quotation marks ('). Use ASCII characters to specify this parameter.
		If password authentication is performed, you must specify this parameter in the settings used when Mirroring Controller connects to a database instance.
		If you specify this parameter when password authentication is not performed, the parameter will be ignored.
		If the specified value of this parameter includes ' or  write \' or \ respectively.
enable_hash_in_password	on or off	Specify on to treat the # in the db_instance_password specification as a password character, or off to treat it as a comment.
		The default is "off".
core_file_path	'coreFileOutputDir'	Specify the directory to which the core file is to be output, enclosed in single quotation marks ('). Use ASCII characters, and specify "\\" as the path delimiter.
		If this parameter is omitted, it will be assumed that the Mirroring Controller management directory was specified.
db_instance_service_name	'registeredServiceNameOfFujitsuEnt erprisePostgresInstance'	Specify the registered service name of the FUJITSU Enterprise Postgres instance in the Windows service.
		Use ASCII characters to specify this parameter, enclosed in single quotation marks (').
mc_service_name	'registeredServiceNameOfMirroring Controller'	Specify the Mirroring Controller service name registered in the Windows service. Use ASCII characters excluding forward slash (/) and backslash (\) to specify this parameter, enclosed in single quotation marks (').

Parameter	Value set	Explanation
		The maximum length of the service name is 124 bytes.
		The default is 'MirroringControllerOpen'.
event_source(*1)	'eventSourceName'	Specify the event source name to be used to identify the Mirroring Controller message in the event log, enclosed in single quotation marks ('). Use ASCII characters to specify this parameter.
		The maximum length of the event source name is 255 bytes.
		The default is 'MirroringControllerOpen'.
remote_call_timeout	Admin communication timeout	Specify the timeout value (milliseconds) of the Mirroring Controller agent process for communication between servers.
		Specify a value between 0 and 2147483647 to be less than the operation system TCP connection timeout (*2).
		The value 0 indicates that there is no timeout limit.
		The default is 70000 milliseconds (70 seconds).
agent_alive_timeout	Timeout for Mirroring Controller process heartbeat monitoring	If there is no response for at least the number of seconds specified, the Mirroring Controller process is restarted.
	(seconds)	Specify 0 or a value between 2 and 2147483647. The value 0 indicates that there is no timeout limit.
		The default is 0 seconds.
heartbeat_error_action	Operation when a heartbeat abnormality is detected using operating system or server heartbeat monitoring	arbitration: Perform automatic degradation using the arbitration server.
		command: Call a user exit to determine degradation, and perform automatic degradation if required.
		message: Notify messages.
		fallback: Perform automatic degradation unconditionally.
		The default is "arbitration".
		Set the same value on the primary server and standby server.
heartbeat_interval	Interval time for abnormality monitoring during heartbeat	Abnormality monitoring of the operating system or server is performed at the interval specified in heartbeat_interval.
	monitoring of the operating system or server (milliseconds)	If an error is detected, operation will conform to the value specified for heartbeat_error_action. If "arbitration" is specified in heartbeat_error_action, the error detection time during monitoring of the operating system or server becomes longer than when the arbitration server is not used, by up to the value specified for arbitration_timeout.
		Specify a value between 1 and 2147483647.
		The specified value is used as the default for db_instance_check_interval and disk_check_interval.
		The default is 800 milliseconds.
heartbeat_timeout	Timeout for abnormality monitoring during heartbeat monitoring of the operating system or server (seconds)	If there is no response for at least the number of seconds specified, it will be assumed that an error has occurred that requires the primary server to be switched, or the standby server to be disconnected.

Parameter	Value set	Explanation
		If an error is detected, operation will conform to the value specified for heartbeat_error_action. If "arbitration" is specified in heartbeat_error_action, the error detection time during monitoring of the operating system or server becomes longer than when the arbitration server is not used, by up to the value specified for arbitration_timeout.
		Specify a value between 1 and 2147483647.
		The specified value is used as the default for db_instance_check_timeout.
		The default is 1 second.
heartbeat_retry	Number of retries for abnormality monitoring during heartbeat monitoring of the operating system or server (number of times)	Specify the number of retries to be performed when an error has been detected that requires the primary server to be switched, or the standby server to be disconnected. If an error is detected in succession more than the specified number of times, switch or disconnection will be performed.
		If an error is detected, operation will conform to the value specified for heartbeat_error_action. If "arbitration" is specified in heartbeat_error_action, the error detection time during monitoring of the operating system or server becomes longer than when the arbitration server is not used, by up to the value specified for arbitration_timeout.
		Specify a value between 0 and 2147483647.
		The specified value is used as the default for db_instance_check_retry and disk_check_retry.
		The default is 2 times.
db_instance_check_interval	Database process heartbeat monitoring interval (milliseconds)	Heartbeat monitoring of the database process is performed at the interval specified in db_instance_check_interval.
		This parameter setting is also used for abnormality monitoring of streaming replication.
		Specify a value between 1 and 2147483647.
		The default is the value set for heartbeat_interval.
db_instance_check_timeout	Database process heartbeat monitoring timeout (seconds)	If there is no response for at least the number of seconds specified, it will be assumed that an error has occurred that requires the primary server to be switched, or the standby server to be disconnected.
		Specify a value between 1 and 2147483647.
		The default is the value set for heartbeat_timeout.
db_instance_check_retry	Number of retries for database process heartbeat monitoring (number of times)	Specify the number of retries to be performed when an error has been detected that requires the primary server to be switched, or the standby server to be disconnected. If an error is detected in succession more than the specified number of times, switch or disconnection will be performed. However, if it detects that the database process is down, it will immediately switch or disconnect regardless of the setting of this parameter.  This parameter setting is also used for abnormality monitoring of streaming replication.

Parameter	Value set	Explanation
		Specify a value between 0 and 2147483647.
		The default number of retries is the value set for heartbeat_retry.
db_instance_timeout_action	none, message, or failover	Specify the behavior for no-response monitoring of the instance.
		none: Do not perform no-response monitoring.
		message: Notify messages if an error is detected during no- response monitoring.
		failover: Perform automatic degradation if an error is detected during no-response monitoring.
		The default is "failover".
disk_check_interval	Interval time for disk abnormality monitoring (milliseconds)	Abnormality monitoring of disk failure is performed at the interval (milliseconds) specified in disk_check_interval. If the file cannot be created, it will be assumed that an error has occurred that requires the primary server to be switched, or the standby server to be disconnected.
		Specify a value between 1 and 2147483647.
		The default is the value set for heartbeat_interval.
disk_check_retry	Number of retries for disk abnormality monitoring (number of times)	Specify the number of retries to be performed when an error has been detected that requires the primary server to be switched, or the standby server to be disconnected.
		If an error is detected in succession more than the specified number of times, switch or disconnection will be performed.
		Specify a value between 0 and 2147483647.
		The default number of retries is the value set for heartbeat_retry.
tablespace_directory_error_ action	message or failover	Specify the behavior to be implemented if an error is detected in the tablespace storage directory.
		message: Notify messages.
		failover: Perform automatic degradation.
		The default is "failover".
arbiter_alive_interval	Interval time for monitoring connection to the Mirroring Controller	A heartbeat is sent to the Mirroring Controller arbitration process at the specified interval.
	arbitration process (milliseconds)	Specify a value between 1 and 2147483647.
		The default is 16000 milliseconds.
		This parameter does not need to be set for operation that does not use the arbitration server.
arbiter_alive_timeout	Timeout for monitoring connection to the Mirroring Controller arbitration process (seconds)	If the heartbeat does not respond within the specified number of seconds, the Mirroring Controller arbitration process is determined to have been disconnected, a message is output, and reconnection is attempted.
		Specify a value between 1 and 2147483647.
		The default is 20 seconds.

Parameter	Value set	Explanation
		This parameter does not need to be set for operation that does not use the arbitration server.
arbiter_alive_retry	Number of retries for monitoring connection to the Mirroring Controller arbitration process (number of times)	Specify the number of heartbeat retries to be performed if an error is detected in the heartbeat to the Mirroring Controller arbitration process. If the heartbeat does not respond within the specified number of retries, the Mirroring Controller arbitration process is determined to have been disconnected.
		Specify a value between 0 and 2147483647.
		The default is 0 times.
		This parameter does not need to be set for operation that does not use the arbitration server.
arbiter_connect_interval	Attempt interval for connection to the Mirroring Controller arbitration process (milliseconds)	Reconnection is attempted at the specified interval if connection fails at startup of the Mirroring Controller process or if the Mirroring Controller arbitration process is disconnected.
		Specify a value between 1 and 2147483647.
		The default is 16000 milliseconds.
		This parameter does not need to be set for operation that does not use the arbitration server.
arbiter_connect_timeout	Timeout for connection to the Mirroring Controller arbitration process (seconds)	If reconnection at startup of the Mirroring Controller process or after disconnection of the Mirroring Controller arbitration process does not succeed within the specified number of seconds, connection to the Mirroring Controller arbitration process is determined to have failed and reconnection is attempted.
		Specify a value between 1 and 2147483647.
		The default is 20 seconds.
		This parameter does not need to be set for operation that does not use the arbitration server.
fencing_command	'fencingCmdFilePath' [Setting example]	Specify the full path of the fencing command that fences a database server where an error is determined to have occurred.
	fencing_command = 'c:\\mc\ \fencing_dir\\execute_fencing.bat'	Specify "\\" as the delimiter.
	(	Enclose the path in single quotation marks (').
		Specify the path using less than 260 bytes.
		Any multibyte characters must use the same encoding as the operating system.
		This parameter must be specified when "command" is set for heartbeat_error_action.
fencing_command_timeout	Fencing command timeout (seconds)	If the command does not respond within the specified number of seconds, fencing is determined to have failed and a signal (SIGTERM) is sent to the fencing command execution process.
		Specify a value between 1 and 2147483647.
		The default is 20 seconds.

Parameter	Value set	Explanation
arbitration_timeout	Arbitration processing timeout in the Mirroring Controller arbitration process (seconds)	The specified value must be at least equal to the value of fencing_command_timeout in the arbitration configuration file, which is the heartbeat monitoring time of the operating system or server.
		If there is no response for at least the number of seconds specified, the primary server will not be switched and the standby server will not be disconnected. Therefore, perform degradation manually.
		If the heartbeat_interval, heartbeat_timeout, and heartbeat_retry values are specified in arbitration.conf for the arbitration server, use the arbitration server values to design arbitration_timeout.
		Specify a value between 1 and 2147483647.
		The default is 30 seconds.
		This parameter does not need to be set for operation that does not use the arbitration server.
arbitration_command	'arbitrationCmdFilePath' [Setting example]	Specify the full path of the arbitration command to be executed when an abnormality is detected by heartbeat monitoring of the operating system or server.
	arbitration_command = 'c:\\mc\ \arbitration_dir\	Specify "\\" as the delimiter.
	\arbitration_dir\ \execute_arbitration_command.bat'	Enclose the path in single quotation marks (').
		Any multibyte characters must use the same encoding as the operating system.
		Specify the path using less than 260 bytes.
		This parameter must be specified when "command" is set for heartbeat_error_action.
arbitration_command_timeo ut	Arbitration command timeout (seconds)	If the arbitration command does not respond within the specified number of seconds, it is determined that execution of the arbitration command has failed and a signal (SIGTERM) is sent to the arbitration command execution process.
		Specify a value between 1 and 2147483647.
		The default is 30 seconds.
		This parameter can be specified only when "command" is set for heartbeat_error_action.
shutdown_detached_synchr onous_standby	on or off	Specify whether to forcibly stop the instance on the standby server when the standby server is disconnected.
		on: Stop the instance.
		off: Do not stop the instance.
		If "on" is specified and the pre-detach command was created, the pre-detach command is executed and then the instance is stopped.
		The default is "off".
post_switch_command	'postSwtchCmdFilePath' [Setting example]	Specify the full path of the command to be called by Mirroring Controller after a new primary server is promoted during a failover of the primary server.

Parameter	Value set	Explanation
	post_switch_command = 'c:\\mc\	Specify "\\" as the delimiter.
	\status_change\ \execute_post_switch.bat'	Enclose the path in single quotation marks (').
	(execute_post_switch.but	Specify the path using less than 260 bytes.
post_attach_command	'postAttachCmdFilePath'  [Setting example]  post_attach_command = 'c:\\mc\ \status_change\ \execute_post_attach.bat'	Specify the full path of the command to be called by Mirroring Controller after the standby server is attached to the cluster system.  Specify "\\" as the delimiter.  Enclose the path in single quotation marks (').  Specify the path using less than 260 bytes.
pre_detach_command	'preDetachCmdFilePath'  [Setting example]  pre_detach_command = 'c:\\mc\ \status_change\ \execute_pre_detach.bat'	Specify the full path of the command to be called by Mirroring Controller before the standby server is disconnected from the cluster system.  Specify "\\" as the delimiter.  Enclose the path in single quotation marks (').  Specify the path using less than 260 bytes.
status_change_command_ti meout	State transition command timeout (seconds)	Specify the timeout value of the post-switch command, post-attach command, and pre-detach command. If the command does not respond within the specified number of seconds, a signal (SIGTERM) is sent to the execution process of the status change command.  Specify a timeout between 1 and 2147483647.  The default is 20 seconds.
check_synchronous_standb y_names_validation	on or off	Specify whether Mirroring Controller is to periodically check during operations whether the synchronous_standby_names parameter in postgresql.conf was changed by an incorrect user operation.  However, it is not recommended to enable this parameter, because performing this check causes Mirroring Controller to use the CPU of the database server redundantly and execute SQL statements at high frequency.  This parameter is compatible with operations in FUJITSU Enterprise Postgres 9.6 or earlier.  The default is "off".
db_instance_ext_pq_connin fo	'libpqConnectionSSLParmToConnect ToDbinstance'	Specify, in key-value form, the connection parameter for libpq that Mirroring Controller adds when connecting to a database. The connection parametersyou can specify are those related to SSL. Use ASCII characters to specify this parameter.  The connection parameter specified in this parameter must also be specified in the db_instance_ext_jdbc_conninfo.
db_instance_ext_jdbc_conni nfo	'JDBCConnectionSSLParmToConne ctToDbinstance'	Specify, in URI form, the connection parameter for JDBC that Mirroring Controller adds when connecting to a database. The connection parametersyou can specify are those related to SSL. Use ASCII characters to specify this parameter.  The connection parameter specified in this parameter must also be specified in the db_instance_ext_pq_conninfo.

- \*1: By using an event source name that is similar to the specified event\_source parameter of the postgresql.conf file, the Mirroring Controller output content can be referenced transparently, so log reference is easy.
- \*2: The operating system TCP connection timeout period is determined by the TCP global parameters initial RTO and maximum SYN retransmits. The remote\_call\_timeout parameter must be set to a value that is shorter than the timeout period for the operating system TCP connection timeout, so change either parameter as necessary.

The availability of some parameters depends on the value set for the heartbeat\_error\_action parameter that sets the operation to be performed if heartbeat monitoring of the operating system or server detects a heartbeat abnormality.

Table A.5 Parameter availability depending on the value set for the heartbeat error action parameter

Possmeter	Value set			
Parameter	arbitration	command	message	fallback
arbiter_alive_interval	Y	N	N	N
arbiter_alive_timeout	Y	N	N	N
arbiter_alive_retry	Y	N	N	N
arbiter_connect_interval	Y	N	N	N
arbiter_connect_timeout	Y	N	N	N
arbitration_timeout	Y	N	N	N
arbitration_command	N	R	N	N
arbitration_command_timeout	N	Y	N	N
fencing_command	Y	R	Y	N
fencing_command_timeout	Y	Y	Y	N
shutdown_detached_synchronous_standby	Y	Y	N	N

R: Required

Y: Can be specified

N: Cannot be specified

# A.4.2 Arbitration Configuration File

In arbitration.conf, define the information related to arbitration and control of the Mirroring Controller arbitration process.

#### Linux

Table A.6 arbitration.conf file (Linux)

Parameter	Value set	Description
port	Port number of the Mirroring Controller arbitration process	The specified value must not exceed the range 0 to 65535. Ensure that the port number does not conflict with other software. Do not specify an ephemeral port that may temporarily be assigned by another program.  For the port number of the arbitration server to be specified in network.conf on the database server, specify the same value as the port number specified in this parameter.
my_address	'ipAddrOrHostNameThatAcceptsConn ectionFromMirroringControllerProces sOnDbServer' [Setting example] my_address = '192.0.3.120'	For the IP address or host name of the arbitration server to be specified in network.conf on the database server, specify the same value as the IP address or host name specified in this parameter.  IPv4 and IPv6 addresses can be specified.

Parameter	Value set	Description
		Specify the IP address or host name, enclosed in single quotation marks (').
core_file_path	'coreFileOutputDir'	Specify the directory to which the core file is to be output, enclosed in single quotation marks ('). Use ASCII characters to specify this parameter.
		If this parameter is omitted, it will be assumed that the Mirroring Controller arbitration process management directory was specified.
syslog_facility	Specify LOCAL0, LOCAL1, LOCAL2, LOCAL3, LOCAL4, LOCAL5, LOCAL6, or LOCAL7.	When the import of logs to the syslog is enabled, the value of this parameter will be used for "facility" of the syslog.
		The default is "LOCAL0".
syslog_ident	'programName'	Specify using single quotation marks (') to enclose the program name used to identify the Mirroring Controller arbitration process message in the system log. Use ASCII characters excluding spaces to specify this parameter.
		The default is 'MirroringControllerArbiter'.
fencing_command	'fencingCmdFilePath'	Specify the full path of the fencing command that
	[Setting example]	fences a database server where an error is determined to have occurred.
	fencing_command = '/arbiter/	Enclose the path in single quotation marks (').
	fencing_dir/execute_fencing.sh'	Specify the path using less than 1024 bytes.
fencing_command_timeout	Fencing command timeout (seconds)	If the command does not respond within the specified number of seconds, fencing is determined to have failed and a signal (SIGTERM) is sent to the fencing command execution process.
		Specify a value between 1 and 2147483647.
		The default is 20 seconds.
heartbeat_interval(*1)	Interval time for heartbeat monitoring of the operating system or server (milliseconds)	The heartbeat monitoring of the database server is checked at the specified interval and arbitration is performed.
		Specify a value between 1 and 2147483647.
		The default is the value specified in serverIdentifier.conf of the database server.
		Specify this parameter to perform optimization taking into account differences in the line load to the admin network and the reduction in the time it takes to degrade.
heartbeat_timeout	Timeout for heartbeat monitoring of the operating system or server (seconds)	If there is no response for at least the number of seconds specified, it will be assumed that an error has occurred that requires the primary server or standby server to be fenced.
		Specify a value between 1 and 2147483647.
		The default is the value specified in serverIdentifier.conf of the database server.

Parameter	Value set	Description
		Specify this parameter to perform optimization taking into account differences in the line load to the admin network and the reduction in the time it takes to degrade.
heartbeat_retry	Number of retries for heartbeat monitoring of the operating system or server (number of times)	Specify the number of retries to be performed when an error has been detected that requires the primary server or standby server to be fenced.  If an error is detected in succession more than the specified number of times, fencing will be performed.  Specify a value between 0 and 2147483647.  The default is the value specified in serverIdentifier.conf of the database server.  Specify this parameter to perform optimization taking into account differences in the line load to the admin network and the reduction in the time it takes to degrade.

<sup>\*1:</sup>Refer to "2.11.4 Tuning for Optimization of Degradation Using Abnormality Monitoring" for information on the tuning parameters for operating system or server abnormality monitoring when using an arbitration server.

#### Windows

Table A.7 arbitration.conf file (Windows)

Parameter	Value set	Description
port	Port number of the Mirroring Controller arbitration process	The specified value must not exceed the range 0 to 65535. Ensure that the port number does not conflict with other software. Do not specify an ephemeral port that may temporarily be assigned by another program.
		For the port number of the arbitration server to be specified in network.conf on the database server, specify the same value as the port number specified in this parameter.
my_address	'ipAddrOrHostNameThatAcceptsConn ectionFromMirroringControllerProces sOnDbServer' [Setting example] my_address = '192.0.3.120'	For the IP address or host name of the arbitration server to be specified in network.conf on the database server, specify the same value as the IP address or host name specified in this parameter.  IPv4 and IPv6 addresses can be specified.  Specify the IP address or host name, enclosed in single quotation marks (').
core_file_path	'coreFileOutputDir	Specify the directory to which the core file is to be output, enclosed in single quotation marks ('). Use ASCII characters, and specify "\\" as the path delimiter.  If this parameter is omitted, it will be assumed that the Mirroring Controller arbitration process management directory was specified.
service_name	'registeredServiceNameOfMirroringCo ntrollerArbitrationProcess'	Specify the Mirroring Controller arbitration process service name to be registered as a Windows service, enclosed in single quotation

Parameter	Value set	Description
		marks ('). Use ASCII characters excluding forward
		slash (/) and backslash (\) to specify this parameter.
		The service name is up to 124 bytes.
ayant sayıra	'eventSourceName'	The default is 'MirroringControllerArbiter'.  Specify the event source name to be used to
event_source	eventsourceivanie	identify the Mirroring Controller arbitration process message in the event log, enclosed in single quotation marks ('). Use ASCII characters to specify this parameter.
		The maximum length of the event source name is 255 bytes.
		The default is 'MirroringControllerArbiter'.
fencing_command	'fencingCmdFilePath' [Setting example]	Specify the full path of the fencing command that fences a database server where an error is determined to have occurred.
	fencing_command = 'c:\\arbiter\	Specify "\\" as the delimiter.
	\fencing_dir\\execute_fencing.bat'	Enclose the path in single quotation marks (').
		Specify the path using less than 260 bytes.
		Any multibyte characters must use the same encoding as the operating system.
fencing_command_timeout	Fencing command timeout (seconds)	If the command does not respond within the specified number of seconds, fencing is determined to have failed and a signal (SIGTERM) is sent to the fencing command execution process.
		Specify a value between 1 and 2147483647.
		The default is 20 seconds.
heartbeat_interval(*1)	Interval time for heartbeat monitoring of the operating system or server (milliseconds)	The heartbeat monitoring of the database server is checked at the specified interval and arbitration is performed.
		Specify a value between 1 and 2147483647.
		The default is the value specified in serverIdentifier.conf of the database server.
		Specify this parameter to perform optimization taking into account differences in the line load to the admin network and the reduction in the time it takes to degrade.
heartbeat_timeout	Timeout for heartbeat monitoring of the operating system or server (seconds)	If there is no response for at least the number of seconds specified, it will be assumed that an error has occurred that requires the primary server or standby server to be fenced.
		Specify a value between 1 and 2147483647.
		The default is the value specified in serverIdentifier.conf of the database server.
		Specify this parameter to perform optimization taking into account differences in the line load to the admin network and the reduction in the time it takes to degrade.

Parameter	Value set	Description	
heartbeat_retry	Number of retries for heartbeat monitoring of the operating system or server (number of times)	Specify the number of retries to be performed when an error has been detected that requires the primary server or standby server to be fenced.	
		If an error is detected in succession more than the specified number of times, fencing will be performed.	
		Specify a value between 0 and 2147483647.	
		The default is the value specified in serverIdentifier.conf of the database server.	
		Specify this parameter to perform optimization taking into account differences in the line load to the admin network and the reduction in the time it takes to degrade.	

<sup>\*1:</sup>Refer to "2.11.4 Tuning for Optimization of Degradation Using Abnormality Monitoring" for information on the tuning parameters for operating system or server abnormality monitoring when using an arbitration server.

# Appendix B Supplementary Information on Building the Primary Server and Standby Server on the Same Server

The primary server and standby server can be pseudo-configured on the same server for system testing, for example. Out of consideration for performance and reliability, do not use this type of configuration for any other purposes. For this reason, do not use this type of configuration in a production environment.

Note that the setup and operations is the same as if the primary and standby servers are built on different servers.

This appendix provides supplementary information explaining how to configure the primary server and standby server on the same server.



Even if automatic degradation by an arbitration server is set when the primary server and standby server are configured on the same server, there will be no effect of it.

# **B.1 Backup Data Storage Destination Directory**

It is not a problem if the same backup data storage destination directory is used on the primary server and standby server.

# **B.2** Registering Service Names and Event Source Names in the Windows Service

Ensure that the following names of resources to be registered on the operating system by FUJITSU Enterprise Postgres and Mirroring Controller are not duplicated between the primary server and standby server:

- Service name registered in the Windows service
- Event source name

# B.3 How to Execute the mc ctl Command

When executing the mc\_ctl command, specify the server identifier in the --local-server option in order to identify the operation destination server.

Below is an example of starting Mirroring Controller of the server "server1" defined in the network.conf file. For mc\_ctl command operations using another mode, also specify the --local-server option.

Define two server identifiers for the same IP address with different port numbers in the network.conf file.

Example)

```
server1 192.0.2.100 27540 server2 192.0.2.100 27541
```

Ensure that the port numbers of both primary server and standby server do not conflict with any other software.

Enabling automatic switch/disconnection

Start Mirroring Controller of the server "server1":

Example)

```
> mc_ctl start -M D:\mcdir\inst1 --local-server server1
```

Stop Mirroring Controller of the server "server1":

Example)

```
> mc_ctl stop -M D:\mcdir\inst1 --local-server server1
```

Disabling automatic switch/disconnection

Start Mirroring Controller of the server "server1":

Example)

```
> mc_ctl start -M D:\mcdir\instl -F --local-server server1
```

Stop Mirroring Controller of the server "server1":

Example)

> mc\_ctl stop -M D:\mcdir\inst1 --local-server server1



To specify the mc\_ctl command with register mode, for registering to the Windows service, and the mc\_ctl command with unregister mode, for unregistering from the Windows service, add the --local-server option.

# Appendix C User Commands

This appendix describes three categories of commands:

- Fencing command
- Arbitration command
- State transition commands

This appendix describes each category of user command.

# C.1 Fencing Command

#### **Format**

The syntax for calling the fencing command from the Mirroring Controller process or the Mirroring Controller arbitration process is described below.

Fencing command of the database server

fencingCmd executionMode mcDegradationOper cmdServerId targetServerId primarycenter

Fencing command of the arbitration server

fencingCmd executionMode mcDegradationOper targetServerId

#### Input

Fencing command of the database server

Execution mode

monitor: Detect issues via automatic monitoring of the Mirroring Controller process

command: Mirroring Controller command execution (switch mode or detach mode of the mc\_ctl command)

Degradation operation to be performed by Mirroring Controller

switch: Switch detach: Disconnect

cmdServerId

ID of the database server that called the command

targetServerId

ID of the database server to be fenced

primarycenter

Fixed value

Fencing command of the arbitration server

Execution mode

monitor: Detect issues via automatic monitoring of the Mirroring Controller process

command: Mirroring Controller command execution (switch mode or detach mode of the mc\_ctl command)

Degradation operation to be performed by Mirroring Controller

switch: Switch detach: Disconnect

targetServerId

ID of the database server to be fenced

#### Output

#### Return value

0: Mirroring Controller will continue the degradation process.

Other than 0: Mirroring Controller will cancel the degradation process.

#### Description

Identifies the database server targeted for fencing based on the input server identifier, and implements the process that isolates it from the cluster system.

#### **Notes**

- The command is executed by the operating system user who started Mirroring Controller or the Mirroring Controller arbitration process. Therefore, if the command is to be executed by a specific operating system user, change the executing user of the command accordingly.
- The operating system user who started Mirroring Controller or the Mirroring Controller arbitration process must have execution privileges to the command. Otherwise, the degradation process will be canceled.
- From a security point of view, set the access privileges as necessary so that the fencing command is not overwritten and unauthorized operations are not performed by unintended operating system users.
- If the fencing command returns a value other than 0, Mirroring Controller will cancel the degradation process, so it is necessary for the user to check the status of the server, and switch or disconnect it manually.
- Before executing the fencing command, check if the server is already fenced, to avoid the command terminating abnormally.
- If the command times out, Mirroring Controller will stop the command, output an error message, and cancel the degradation process.

#### Information

The fencing command can be implemented by simply stopping the operating system or server. For example, if stopping the power for the database server, it is possible to use a utility to control the hardware control board in environments equipped with boards compatible with IPMI hardware standard.

#### Linux

Below is a sample script of a fencing command that powers off the database server using the IPMI tool.

Sample shell script

/installDir/share/mcarb\_execute\_fencing.sh.sample

#### Windows

Below is a sample script of a fencing command that powers off the database server using IPMIUTIL.

Sample shell script

installDir\share\mcarb\_execute\_fencing.bat.sample

#### C.2 Arbitration Command

#### **Format**

The syntax for calling the arbitration command from the Mirroring Controller process is described below.

arbitrationCmd cmdServerId targetServerId primarycenter

#### Input

cmdServerId

ID of the database server that called the command

targetServerId

ID of the database server to arbitrate

primarycenter

Fixed value

#### Output

Return value

0: The database server to arbitrate has an issue, and Mirroring Controller will continue the degradation process.

Other than 0: The database server to arbitrate is normal, and Mirroring Controller will cancel the degradation process.

#### Description

Identifies the database server to arbitrate based on the input server identifier, and checks the status of the server.

#### **Notes**

- The command is executed by the operating system user who started Mirroring Controller.
- The operating system user who started Mirroring Controller must have execution privileges to the command. Otherwise, the command will not be called, and the degradation process will be canceled.
- If the command times out, Mirroring Controller will stop the command, output an error message, and cancel the degradation process.

### C.3 State Transition Commands

State transition commands include the three types of user commands below. Any of the commands can be implemented by Mirroring Controller in conjunction with database server status transitions.

- Post-switch command
- Pre-detach command
- Post-attach command

#### C.3.1 Post-switch Command

#### **Format**

The syntax for calling the post-switch command from the Mirroring Controller process is described below.

postswitchCmd serverIdentifier primarycenter

#### Input

serverldentifier

ID of the database server (new primary server) that was switched

primarycenter

Fixed value

#### Output

Return value

None

#### Notes

- The command is executed by the operating system user who started Mirroring Controller.
- The operating system user who started Mirroring Controller must have execution privileges to the command. Otherwise, the command will not be called.
- If the command times out, Mirroring Controller will stop the command, output an error message, and cancel the process.

#### C.3.2 Pre-detach Command

#### **Format**

The syntax for calling the pre-detach command from the Mirroring Controller process is described below.

predetachCmd cmdServerId serverRole targetServerId primarycenter

#### Input

#### cmdServerId

ID of the database server that called the command

#### Server role

Role of the database server that called the command

primary: Primary standby: Standby

#### targetServerId

ID of the standby server to be disconnected from the cluster system

primarycenter

Fixed value

#### Output

Return value

None

#### **Notes**

- The command is executed by the operating system user who started Mirroring Controller.
- The operating system user who started Mirroring Controller must have execution privileges to the command. Otherwise, the command will not be called, however, Mirroring Controller will output an error message and continue the process.
- If the command times out, Mirroring Controller will stop the command, output an error message, and cancel the process.

#### C.3.3 Post-attach Command

#### **Format**

The syntax for calling the post-attach command from the Mirroring Controller process is described below.

postattachCmd cmdServerId serverRole targetServerId primarycenter

#### Input

#### cmdServerId

ID of the database server that called the command

#### Server role

Role of the database server that called the command

primary: Primary standby: Standby

#### targetServerId

ID of the standby server to be attached to the cluster system

primarycenter

Fixed value

#### Output

Return value

None

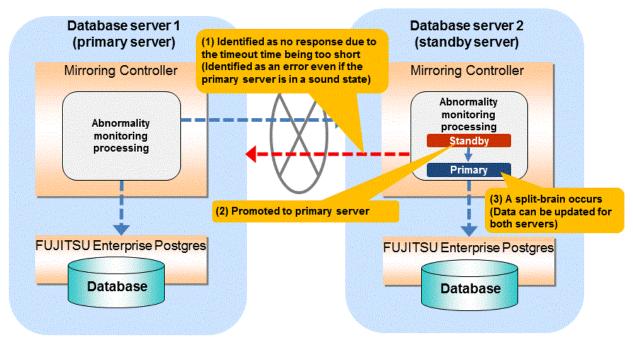
#### **Notes**

- The command is executed by the operating system user who started Mirroring Controller.
- The operating system user who started Mirroring Controller must have execution privileges to the command. Otherwise, the command will not be called.
- If the command times out, Mirroring Controller will stop the command, output an error message, and cancel the process.

# Appendix D Notes on Performing Automatic Degradation Immediately after a Heartbeat Abnormality

The type of issue below occurs if automatic degradation is performed unconditionally after an issue is detected during heartbeat monitoring of an operating system or server, and heartbeat monitoring was not properly tuned.

#### • If the timeout time is too short



#### If the timeout time is too long Database server 1 Database server 2 (primary server) (standby server) Mirroring Controller Mirroring Controller (1) The operating system or server crashes or is unresponsive Abnormality ity itoring monitoring p cessi processing Failure occurs (2) Errors cannot be detected in realtime due to the timeout time being too long FUJITSU Enterprise Postgres FUJITSU Enterprise Postgres **Database Database** (3) Operations cannot Application server continue until the server is switched Job application

#### Notes on monitoring when the operating system or server crashes or is unresponsive

As illustrated in the diagram above, timeout is used to monitor whether the operating system or server crashes or is unresponsive. Therefore, if tuning has not been performed correctly, there is a risk of a split-brain mistakenly occurring even if the server is in a sound state.

Split-brain is a phenomenon in which both servers temporarily operate as primary servers, causing data updates to be performed on both servers.

#### Split-brain detection method

It can be confirmed that split-brain occurs under the following conditions:

- 1. When the mc\_ctl command is executed in status mode on both servers, the "host\_role" of both servers is output as "primary", and
- 2. The following message is output to the system log of one of the servers:

```
promotion processing completed (MCA00062)
```

#### How to recover from a split-brain

Use the procedure described below. Note that the new primary server is the server that was confirmed in step 2 of the aforementioned detection method.

- 1. Stop all applications that are running on the old and new primary servers.
- Investigate and recover the database.Investigate the update results that have not been reflected to the new primary server from the database of the old primary server, and apply to the new primary server as necessary.
- 3. Stop the old primary server instance and the Mirroring Controller.

- 4. Resume the applications that were stopped in step 1.
- Recover the old primary server.
   While referring to "2.5 Setting Up the Standby Server", build (set up) the old primary server as the new standby server, from the new primary server.

# Appendix E Supplementary Procedure on Configuring for Operation in Database Multiplexing Mode

This appendix explains a supplementary procedure on the configuration required for operation in database multiplexing mode.

# **E.1 Security Policy Settings**

This section explains how to configure the security settings to enable an operating system user account designated as an instance administrator user to log on as a service.

1. Displaying the [Local Security Policy] window

In Windows, select [Administrative Tools], and then click [Local Security Policy].

- 2. Setting up security
  - 1. In the [Local Security Policy] window, select [Security Settings], select [Local Policies], and then click [User Rights Assignment].
  - 2. Under [Policy] in the [User Rights Assignment] window, double-click [Log on as a service].
  - 3. In the [Log on as a service Properties] window, set the following:
    - a. Select the [Local Security Setting] tab.
    - b. On the [Local Security Setting] tab, click [Add User or Group].
    - c. In the [Select Users or Groups] window, enter the operating system user account of the instance administrator user in [Enter the object names to select].
    - d. Click [OK].
  - 4. In the [Log on as a service Properties] window, click [OK].
  - 5. From the [Local Security Policy] tree, click [Local Policies], and then double-click [Security Options].
  - Scroll down and double-click [User Account Control: Behavior of the elevation prompt for administrators in Admin Approval Mode].
  - 7. From the drop-down menu, select the "Elevate without prompting" in the [Local Security Setting] tab.
  - 8. Click [OK].

# E.2 Windows Firewall Settings

This section explains how to enable the port number used by Mirroring Controller, if the Windows firewall feature is enabled.

#### Windows Server(R) 2019:

- 1. In the [Windows Defender Firewall] window, click [Advanced settings] on the left side of the window.
- 2. In the [Windows Defender Firewall with Advanced Security] window, click [Inbound Rules] on the left side of the window.
- 3. Click [New Rule] on the right side of the window.
- 4. In the [New Inbound Rule Wizard] window, select [Port], and then click [Next].
- 5. Select [TCP] and [Specific local ports], then specify the port number defined in the network definition file, and then click [Next].
- 6. Select [Allow the connection], and then click [Next].
- 7. Select the profiles for which this rule applies, and then click [Next].
- 8. In [Name], specify the desired name, and then click [Finish].
- 9. In the [Windows Defender Firewall with Advanced Security] window, check if the added rule is enabled under [Inbound Rules] in the center of the window.

#### Other than above:

- 1. In the [Windows Firewall] window, click [Advanced settings] on the left side of the window.
- 2. In the [Windows Firewall with Advanced Security] window, click [Inbound Rules] on the left side of the window.
- 3. Click [New Rule] on the right side of the window.
- 4. In the [New Inbound Rule Wizard] window, select [Port], and then click [Next].
- 5. Select [TCP] and [Specific local ports], then specify the port number defined in the network definition file, and then click [Next].
- 6. Select [Allow the connection], and then click [Next].
- 7. Select the profiles for which this rule applies, and then click [Next].
- 8. In [Name], specify the desired name, and then click [Finish].
- 9. In the [Windows Firewall with Advanced Security] window, check if the added rule is enabled under [Inbound Rules] in the center of the window.

# Appendix F WebAdmin Disallow User Inputs Containing Hazardous Characters

WebAdmin considers the following as hazardous characters, which are not allowed in user inputs.

| (pipe sign)
& (ampersand sign)
; (semicolon sign)
\$ (dollar sign)
% (percent sign)
@ (at sign)
' (single apostrophe)
" (quotation mark)
\' (backslash-escaped apostrophe)
\" (backslash-escaped quotation mark)
\$ (triangular parenthesis)
() (parenthesis)
+ (plus sign)
CR (Carriage return, ASCII 0x0d)
LF (Line feed, ASCII 0x0a)

, (comma sign)
\ (backslash)

# Appendix G Collecting Failure Investigation Data

If the cause of an error that occurs while building the environment or during operations is unclear, data must be collected for initial investigation.

This appendix describes how to collect data for initial investigation.

Use FJQSS (Information Collection Tool) to collect data for initial investigation.



See

Refer to the following manual for information on how to use FJQSS.

- In the [Apps] menu, select [FJQSS (Information Collection Tool)], and then click [FJQSS User's Guide].



#### Using a database server to collect data

- When using FJQSS to collect data for initial investigation, a window will be displayed for you to set the following environment variables:

- Environment variable required for collecting data from the database

Refer to "Collecting Failure Investigation Data" in the Operation Guide.

- MCCONTROLDIR

Set the Mirroring Controller management directory.

Refer to "1.2.1 Mirroring Controller Resources" for information on the Mirroring Controller management directory.

- FJQSS operations must be performed by the instance administrator user.

#### Using an arbitration server to collect data

#### Linux

- When using FJQSS to collect data for initial investigation, you must set the following environment variables:
  - ARBCONTROLDIR

Set the Mirroring Controller arbitration process management directory.

Refer to "1.2.1 Mirroring Controller Resources" for information on the Mirroring Controller arbitration server management directory.

- ARBUSER

Specify the operating system user who started the Mirroring Controller arbitration process.

- FSEP\_SA\_HOME

Set the Server Assistant installation directory.

#### Windows

- When using FJQSS to collect data for initial investigation, a window will be displayed for you to set the following environment variables:
  - ARBCONTROLDIR

Set the Mirroring Controller arbitration process management directory.

Refer to "1.2.1 Mirroring Controller Resources" for information on the Mirroring Controller arbitration server management directory.

# <u>Index</u>

[A]	[E]
Action Required when a Heartbeat Abnormality is Detected 70	Encryption of Transaction Logs Transferred to the Standby
Action Required when All Database Servers or Instances Stopped	Server14
99	
Action Required when an Error Occurs in the Database	[F]
Multiplexing Mode	Failback of the Primary Server94
Action Required when Automatic Disconnection Fails 98	Fencing Command
Action Required when Automatic Switch Fails98	n. n
Action Required when Server Degradation Occurs	[H]
Addressing Errors During Degrading Operation97	How to Execute the mc_ctl Command
Application Connection Server Settings44	[1]
arbitration.conf file (Linux)	Identify cause of error and perform recovery91,97
arbitration.conf file (Windows)131	Identify Cause of Error and Restore the Standby Server90
Arbitration Command	Identify Cause of Error and Restructure the Standby Server 96
Arbitration Configuration File129	If Performing the Referencing Job on the Synchronous Standby
Arbitration Server Maintenance	Server
Arbitration Server Process9	If Prioritizing the Main Job on the Primary Server
Arbitration Server Resources9	Installation
Authentication of the Standby Server13	instantion
•	[M]
[B]	Manually Disconnecting the Standby Server70
Backing up Database Multiplexing Mode Information62	Manually Switching the Primary Server69
Backup Data Storage Destination Directory	Matching the system times
Backup Operation62	Mirroring Controller Resources
(0)	Monitoring Mirroring Controller Messages71
[C]	Monitoring Using Database Multiplexing Mode
Changes in Operation	
Changes Required when the Standby Server is Stopped80	[N]
Changing from Database Multiplexing Mode to Single Server	network.conf file119
Mode	Network Configuration File118
Changing from Single Server Mode to Database Multiplexing	Notes on CPU Architecture and Products10
Mode	Notes on Performing Automatic Degradation Immediately after a
Changing Parameters	Heartbeat Abnormality141
Changing to Database Multiplexing Mode when the Arbitration	[0]
Server is Used for Automatic Degradation	Operations in Database Multiplexing Mode64
Checking the Behavior	Operations when the Server has Started Degrading after a
Checking the Connection Status	Disconnection has Occurred99
Checking the Connection Status on a Database Server42	Operations when the Server has Started Degrading after a Switch
Checking the Connection Status on the Arbitration Server43	has Occurred88
Checking the Database Multiplexing Mode Status	Overview of Database Multiplexing Mode
Checking the Status of the Arbitration Server	Overview of Database Multiplexing Mode
Checking the Status of the Database Server	[P]
Configuring ICMP	Parameters
Configuring the Arbitration Server	Parameters Set on the Primary Server114
Constitute Statistics and Project Statistics of Project Statistics and Project Statistics of Project Statistic	Parameters Set on the Standby Server116
Creating, Setting, and Registering the Primary Server Instance30	Post-attach Command
Creating, Setting, and Registering the Standby Server Instance37	Post-switch Command
Creating Applications	postgresql.conf file114,116
Creating a User Exit for a Database Server	Pre-detach Command
Creating a User Exit for the Arbitration Server24	Preparatory Tasks for the Output of Error Logs to the Event Log
[D]	
Database Backup Operation	Preparing for Setup17
Database Server Processes	Preparing the Arbitration Server18
Deciding on Operation when a Heartbeat Abnormality is	Preparing the Backup Disk17
Detected	Preparing the Database Server

Preparing to Output Error Logs to the Event Log (Windows). 18
[R]
Rebuild the Standby Server
Recovery of the Mirroring Controller management directory
Redundancy of the Admin and Log Transfer Networks10
Referencing on the Standby Server
Registering Service Names and Event Source Names in the
Windows Service
Rolling Updates
Rolling Opulies
[S]
Security in Database Multiplexing11
Security Policy Settings
Security Policy Settings (Windows)19
ServerConfiguration File
Server Configuration File for the Database Servers121
serverIdentifier.conf file
Server Maintenance
Setting Automatic Start and Stop of Mirroring Controller and
Multiplexed Instances
Setting Automatic Start and Stop of the Mirroring Controller
Arbitration Process
Setting Up Database Multiplexing Mode
Setting Up Database Multiplexing Mode on the Primary Server25
Setting Up Database Multiplexing Mode on the Standby Server35
Setting Up the Arbitration Server
Setting Up the Primary Server25
Setting Up the Standby Server35
Setup
Starting and Stopping Mirroring Controller
Starting and Stopping the Mirroring Controller Arbitration
Process
Starting Mirroring Controller
Starting Mirroring Controller on the Primary Server
Starting Mirroring Controller on the Standby Server
Starting the Mirroring Controller Arbitration Process 24,64
State Transition Commands
Stop Mirroring Controller
Stopping for Maintenance
Stopping Mirroring Controller
Stopping the Mirroring Controller Arbitration Process
Supplementary Information on Building the Primary Server and
Standby Server on the Same Server
Database Multiplexing Mode
System Configuration for Database Multiplexing Mode
System Configuration for Database Multiplexing Mode
ſΠ
Tuning44
Tuning for Optimization of Degrading Operation Using
Abnormality Monitoring45
Tuning to Stabilize Queries on the Standby Server44
Tuning to Stabilize Queries on the Standby Server (when
Performing Frequent Updates on the Primary Server)44
Tuning to Stabilize the Database Multiplexing Mode44

[U]
Uninstalling in Database Multiplexing Mode86
Jsers who perform setup and operations on the arbitration server
Jsers who perform setup and operations on the database server 15
[W]
What is Database Multiplexing Mode1
Vindows Firewall Settings144



# FUJITSU Enterprise Postgres 12



# Cluster Operation Guide PRIMECLUSTER

## **Preface**

#### Purpose of this document

This document explains the items required to operate FUJITSU Enterprise Postgres on a cluster system that uses PRIMECLUSTER.

#### Intended readers

This document is aimed at people who install and operate FUJITSU Enterprise Postgres on a cluster system that uses PRIMECLUSTER. Readers of this document are also assumed to have general knowledge of:

- PRIMECLUSTER
- PostgreSQL
- SQL
- Linux

#### Structure of this document

This document is structured as follows:

#### Chapter 1 Overview of Failover Operation

Provides an overview of failover operation.

#### Chapter 2 Setting Up Failover Operation

Explains how to set up failover operation.

#### **Chapter 3 Failover Operation**

Explains the procedures involved in failover operation.

#### Chapter 4 Procedures Required after a Failover Error

Explains the procedures required after an abnormality occurs during failover operation.

#### Appendix A Creating Resources and Creating/Modifying Cluster Applications

Explains how to create and modify cluster applications that include FUJITSU Enterprise Postgres resources.

#### Appendix B Command Reference

Explains for details in command.

#### **Export restrictions**

Exportation/release of this document may require necessary procedures in accordance with the regulations of your resident country and/or US export control laws.

#### Issue date and version

First edition: July 2020

#### Copyright

Copyright 2018-2020 FUJITSU LIMITED

# **Contents**

Chapter 1 Overview of Failover Operation	
1.1 Definition of Failover Operation.	
1.2 System Configuration of Failover Operation	
Chapter 2 Setting Up Failover Operation	
2.1 Installing PRIMECLUSTER and FUJITSU Enterprise Postgres	
2.2 Configuring PRIMECLUSTER	
2.3 Creating a GDS Volume	
2.4 Creating a File System	
2.4.1 Creating a File System.	
2.4.2 Editing Configuration Files	
2.4.3 Checking	
2.5 Creating an Operating System User to Start FUJITSU Enterprise Postgres	
2.6 Mounting the File System	
2.7 Creating a Symfoware Database Cluster	
2.8 Registering Resource Information for the FUJITSU Enterprise Postgres Database Cluster	
2.9 Configuring Storage Data Protection Using Transparent Data Encryption	
2.10 Unmounting the File System	
2.11 Creating a Cluster Application	1
2.12 Application Development	1
2.12.1 Precautions when Developing Applications	12
2.12.2 Corrective Action with Application when State Transition Occurs	
2.13 Checking Operation	14
Chapter 3 Failover Operation	
3.1 Adding Tablespaces	
3.2 Modifying Database Cluster Resources	
3.3 Operations under Normal Circumstances	
3.4 Operation at State Transition	
3.5 Maintenance Tasks	
3.5.1 Simultaneous Stopped Node Maintenance	19
3.5.1.1 Stopping the Cluster Applications on Both Nodes for Node Maintenance	
3.5.1.2 Maintenance Tasks on Both Nodes	20
3.5.1.3 Restarting the Cluster Applications Stopped on Both Nodes for Node Maintenance	
3.5.2 Mutual Switch Maintenance	20
3.5.2.1 Stopping the Cluster Applications on the Standby Node for Node Maintenance	
3.5.2.2 Standby Node Maintenance.	2
3.5.2.3 Restarting the Cluster Applications Stopped on the Standby Node for Node Maintenance	2
3.5.2.4 Mutual Switch.	2
3.5.2.5 Standby Node Maintenance after the Mutual Switch	2
Chapter 4 Procedures Required after a Failover Error	
4.1 Identifying the Cause of an Error	
4.2 Stopping the Cluster Applications for Node Maintenance	
4.3 Starting the Database Cluster	
4.4 Recovery Tasks	
4.5 Stopping the Database Cluster	
4.6 Restarting the Cluster Applications Stopped for Node Maintenance	23
Appendix A Creating Resources and Creating/Medifying Chater Applications	2
Appendix A Creating Resources and Creating/Modifying Cluster Applications	
A.1 Creating Resources and Cluster Applications.	
A.1.1 Preparing for Creation	
A.1.2 Creating a Cluster Application and Configuring Its Attributes	
A.1.4 Creating EUUTSU Enterprise Postgres Postgres	
A.1.4 Creating FUJITSU Enterprise Postgres Resources.  A.1.5 Creating Gls or Takeover Network Resources.	
A. L. A. Clearing, Cits of Takeover Inciwork Resources	1

A.1.6 Creating Fsystem Resources.	34
A.1.7 Creating Resources and Finalizing Cluster Application Creation	
A.2 Modifying Cluster Applications	36
A.2.1 Moving to the Cluster Application Edit Window	36
A.2.2 Adding a Created Shared Disk Class to the Cluster Application	
A.2.3 Adding a Created File System to the Cluster Application	38
A.2.4 Reflecting Modifications Made to the Cluster Application.	38
Appendix B Command Reference	40
B.1 pgx_pclrsc	40
Index	42

# Chapter 1 Overview of Failover Operation

This chapter provides an overview of failover operation in FUJITSU Enterprise Postgres.

# 1.1 Definition of Failover Operation

When an abnormality occurs on any of the multiple server devices that make up the cluster system (hereafter referred to as "nodes"), the failover feature transfers a job operating on such a node to another one. Failover can reduce the time for which jobs are stopped when an abnormality occurs, and allows jobs to continue while the node on which the abnormality occurred is recovered.

In FUJITSU Enterprise Postgres, failover can be integrated with PRIMECLUSTER. In such a configuration, the shared disk (GDS) provided by PRIMECLUSTER is used. A server in the cluster system receives the application processing, acting as the active server (active node).

For example, if the active node fails, another server (standby node) inherits the shared disk content and becomes the active node. Using the PRIMECLUSTER feature, the applications to be run on the active node for the database can also be linked with the database and switched. Therefore, jobs can be continued even during recovery of a node where an abnormality has occurred. There is also no need to send the updated content of the database on the active node to the standby node, so the processing performance will be the same as that of a non-cluster system.

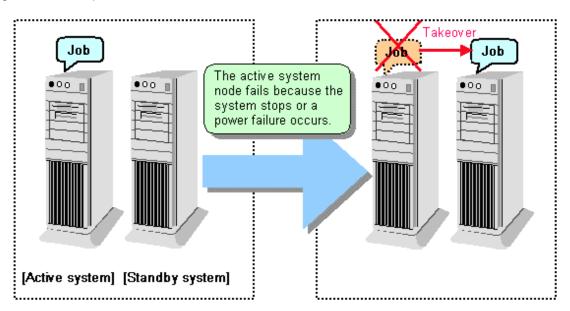


When failover operation is integrated with PRIMECLUSTER, the disk is not referenced by the standby node, and therefore failover operation cannot be used for load distribution.

In addition, failover operation cannot be performed concurrently with database multiplexing.

Hereinafter, failover operation integrated with PRIMECLUSTER is referred to as "failover operation".

Figure 1.1 Concept of failover



#### Feature of failover operation

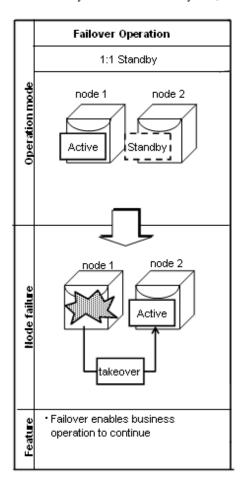
Failover operation in Symfoware Server supports the standby feature.

When an abnormality occurs on a node, the standby feature starts the instance that is the standby system at the time of the switch, and transfers the job to it.

#### Operation mode of failover operation

The operation mode supported in failover operation in Symfoware Server is 1:1 standby.

1:1 standby is a mode in which one active system or standby system operates on one node. Because failover allows jobs to continue when an abnormality occurs in the active system, this mode allows highly reliable systems to be constructed.



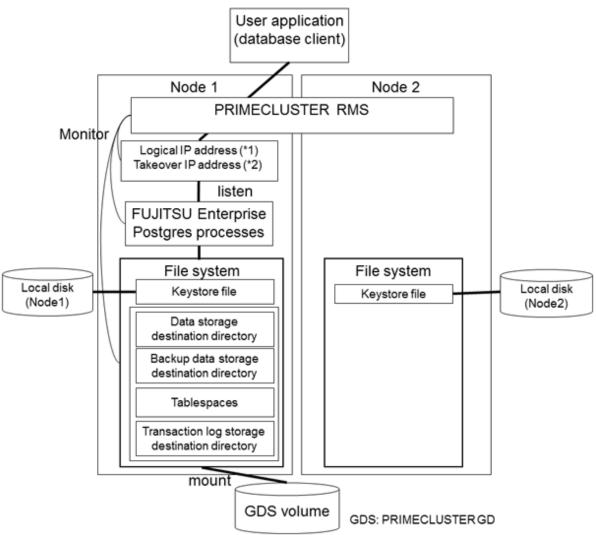
# 1.2 System Configuration of Failover Operation

In the illustration below, the file system is on a GDS volume (shared disk), mounted on node 1 only.

The following FUJITSU Enterprise Postgres resources are located in this file system:

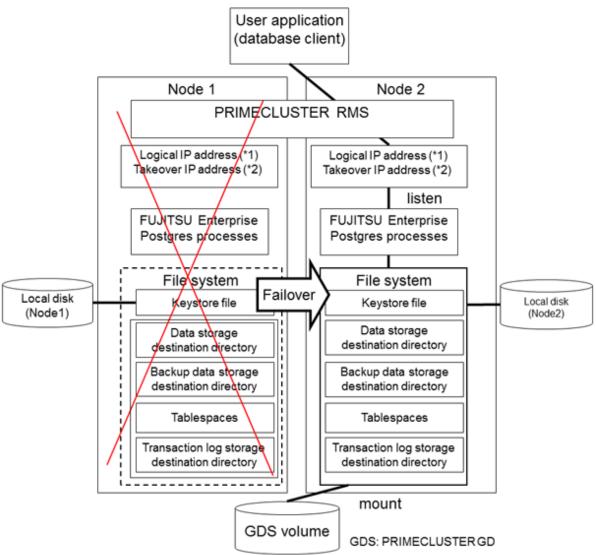
- Data storage destination directory
- Tablespaces
- Backup data storage destination directory
- Transaction log storage destination directory

A FUJITSU Enterprise Postgres process waits for connections of user applications that access the database, on a logical IP address activated on node 1 only or on a takeover IP address.



- \*1: The function of PRIMECLUSTERGLS (Global Link Services)
- \*2: The function of PRIMECLUSTER Takeover network resource

When an error is detected on node 1, PRIMECLUSTER mounts the file system on node 2, and activates the logical IP address or the takeover IP address on it. The FUJITSU Enterprise Postgres process is also started on node 2. Since the logical IP address or the takeover IP address remains unchanged, the user application can access FUJITSU Enterprise Postgres process on node 2 simply by re-establishing the database connection.



- \*1: The function of PRIMECLUSTER GLS (Global Link Services)
- \*2: The function of PRIMECLUSTER Taking over network resource

# Chapter 2 Setting Up Failover Operation

Use the procedure in the table below to perform setup:

Step	Work	Reference	
	Active system	Standby system	
1	Installing PRIMECLUSTER and	I FUJITSU Enterprise Postgres	2.1 Installing PRIMECLUSTER and FUJITSU Enterprise Postgres
2	Configuring PRIMECLUSTER		2.2 Configuring PRIMECLUSTER
3	Creating a GDS volume (*1)		2.3 Creating a GDS Volume
4	Creating a file system (*1)		2.4 Creating a File System
5	Creating a user for the operating system	to start FUJITSU Enterprise Postgres	2.5 Creating an Operating System User to Start FUJITSU Enterprise Postgres
6	Mounting the file system		2.6 Mounting the File System
7	Creating a Symfoware database cluster(*1)		2.7 Creating a Symfoware Database Cluster
8	Registering resource information for the FUJITSU Enterprise Postgres database cluster		2.8 Registering Resource Information for the FUJITSU Enterprise Postgres Database Cluster
9	Configuring storage data protection using transparent data encryption (*1)		2.9 Configuring Storage Data Protection Using Transparent Data Encryption
10	Unmounting the file system		2.10 Unmounting the File System
11	Creating a cluster application		2.11 Creating a Cluster Application
12	Application development		2.12 Application Development
13	Checking operation		2.13 Checking Operation

<sup>\*1:</sup> Some settings and operations must also be performed on the standby node. Refer to the relevant sections for details.

Note: GDS: PRIMECLUSTER GD

The sections that follow explain each step.

## 2.1 Installing PRIMECLUSTER and FUJITSU Enterprise Postgres

Refer to the manual for each product, and then install each product.



#### See

- Refer to the PRIMECLUSTER Installation and Administration Guide for information on how to install PRIMECLUSTER.
- Refer to the Installation and Setup Guide for Server for information on how to install FUJITSU Enterprise Postgres.

## 2.2 Configuring PRIMECLUSTER

Configure the initial settings for the cluster.

If using PRIMECLUSTER GL (hereafter referred to as "GLS"), configure the GLS settings.



#### See

Refer to the PRIMECLUSTER Installation and Administration Guide for information on steps required before setup.

## 2.3 Creating a GDS Volume

Each of the following FUJITSU Enterprise Postgres resources must be located in a file system on a shared disk:

- Data storage destination directory
- Tablespaces
- Backup data storage destination directory
- Transaction log storage destination directory



- Store data storage destination directories or tablespaces and backup data storage destination directories for these in different GDS classes to guard against file system corruption and the MONITORONLY attribute settings at the time of the Gds resource and Fsystem resource creation.
- It is also recommended that database storage destination directories and transaction log storage destination directories be placed in different groups or different single disks, so that the I/O load is distributed.
- Create a directory for each resource on each node. Ensure that the directories to be mounted have the same path on all nodes.



### See

- Refer to the PRIMECLUSTER Installation and Administration Guide for information on GDS volume creation.

- Refer to "Preparing Directories to Deploy Resources" in the Installation and Setup Guide for Server for information on how to create directories.

## 2.4 Creating a File System

Create a file system.



- Refer to the PRIMECLUSTER Installation and Administration Guide for information on how to create a file system on a GDS volume.

### 2.4.1 Creating a File System

Use the operating system or file system command to create a file system on the GDS volume.

Example)

The example below creates a file system using ext4.

# mkfs -t ext4 /dev/sfdsk/class0001/dsk/volume0001

## 2.4.2 Editing Configuration Files

Edit /etc/fstab and /etc/fstab.pcl on each node.

- /etc/fstab

Comment out each line below by prepending it with "#".

```
#/dev/sfdsk/class0001/dsk/volume0001 /mnt/swdsk1 ext4 noauto 0 0
#/dev/sfdsk/class0001/dsk/volume0002 /mnt/swdsk2 ext4 noauto 0 0
```

- /etc/fstab.pcl

Copy the entries above, and then replace "#" in each line with "#RMS#", as follows:

```
#RMS#/dev/sfdsk/class0001/dsk/volume0001 /mnt/swdsk1 ext4 noauto 0 0
#RMS#/dev/sfdsk/class0001/dsk/volume0002 /mnt/swdsk2 ext4 noauto 0 0
```



Different files are specified depending on the version of the PRIMECLUSTER. Check the version of PRIMECLUSTER being used and edit appropriate files. Refer to the PRIMECLUSTER Installation and Administration Guide for details.

## 2.4.3 Checking

Start the GDS volume, and ensure that each file system can be mounted on each node.

```
# mount -t ext4 /dev/sfdsk/class0001/dsk/volume0001 /mnt/swdsk1
# mount -t ext4 /dev/sfdsk/class0001/dsk/volume0002 /mnt/swdsk2
# umount /mnt/swdsk1
# umount /mnt/swdsk2
```

## 2.5 Creating an Operating System User to Start FUJITSU Enterprise Postgres

Create an operating system user (an OS user that will become the instance administrator) who is to start FUJITSU Enterprise Postgres.



Ensure that the name, UID, and GID of the operating system user (an OS user that will become the instance administrator) who is to start FUJITSU Enterprise Postgres match on all nodes that make up the cluster system.

## 2.6 Mounting the File System

- 1. Start the GDS volume if it has not already started.
- Mount all file systems containing the resources required to create database clusters, such as for the data storage destination directory, as shown below:

# mount -t ext4 /dev/sfdsk/class0001/dsk/volume0001 /mnt/swdsk1



Create directories in the file system for storing each resource. Each directory must meet certain conditions to be able to create the FUJITSU Enterprise Postgres database cluster. Check the notes in "Preparing Directories to Deploy Resources" in the Installation and Setup Guide for Server.

## 2.7 Creating a Symfoware Database Cluster

1. Create the database cluster.



- If the configuration requires password authentication, create a .pgpass file in the home directory of the operating system user who is to start FUJITSU Enterprise Postgres, and then specify a password for the template1 database as this is required. This will be used when PRIMECLUSTER checks the operation of the FUJITSU Enterprise Postgres database server. If authentication fails, a fault will be assumed to have occurred in the database server.
- The value specified for the port parameter in postgresql.conf needs to be set in the /etc/services files on all nodes.
- Since postgresql.conf is stored in the data directory, it becomes a common configuration file for each node in the shared disk. Therefore, it must be taken into account that the path set for postgresql.conf must be valid for the FUJITSU Enterprise Postgres server processes started on each node. For example, if a local directory is specified in the core\_directory parameter, that local directory must have been created on all nodes.
- To execute the FUJITSU Enterprise Postgres commands, you must configure the environment variables on each node.
- Use the initdb command to create a database cluster.
- Specify "off" for the restart\_after\_crash parameter in postgresql.conf.
  If "on" is specified, or the default value is used for the restart\_after\_crash parameter, behaviors equivalent to restarting the FUJITSU Enterprise Postgres, including crash recovery, will be performed when some server processes go down. When linking with PRIMECLUSTER, however, a failover will occur after an error is detected as soon as some server processes go down, and server processes restarting will be forced to stop. This is to inhibit complex behaviors such as processes being canceled in the middle of restarting for no particular meaningful purpose.
- 2. Check starting, connection, and stopping at each node. The procedure is as follows.



Queries are issued to monitor the FUJITSU Enterprise Postgres from PRIMECLUSTER. With this connection process, if a password is requested, the monitoring process cannot be continued. For this reason, when checking the connection, ensure that connections can be made without entering a password.



When monitoring FUJITSU Enterprise Postgres from PRIMECLUSTER, the name of the application that connects to instances is "pgx\_wch\_svprocess".

1. Check the starting, connection, and stopping at node 1.

Use the pg\_ctl command to start and stop the node.

For connection, execute the following to ensure that connections can be made without entering a password:

```
su - OsUser
psql -d templatel -p portNum [ -U dbUser]
```

OsUser. Operating system user to start the FUJITSU Enterprise Postgres portNum: Port number specified when the database cluster was created

dbUser. Database superuser

If the database superuser is specified at the time of executing the initdb command, or if the database superuser has been changed after the initdb command was executed, specify the user for the -U option. If the database superuser is not specified or changed, omit this.

- 2. Unmount all file systems containing the resources required to create database clusters (such as the data storage destination directory) on node 1, and stop the GDS volume.
- 3. Start the GDS volume on node 2, and mount all file systems containing the resources required to create database clusters (such as the data storage destination directory).
- 4. Check starting, connection, and stopping at node 2.

Perform the same check as step 1.

- 5. Unmount all file systems containing the resources required to create database clusters (such as the data storage destination directory) on node 2, and stop the GDS volume.
- 6. Start the GDS volume on node 1, and mount all file systems containing the resources required to create database clusters (such as the data storage destination directory).



- Refer to "Using the initdb Command" of "Creating an Instance" in the Installation and Setup Guide for Server" for information on how to create a database cluster.

- Refer to "The Password File" in "Client Interfaces" in the PostgreSQL Documentation for information on files for which passwords are set.

# 2.8 Registering Resource Information for the FUJITSU Enterprise Postgres Database Cluster

1. Use the pgx\_pclrsc command to register the database cluster in PRIMECLUSTER - the following is the simplest execution example: The database cluster does not need to be started when executing the command.

# pgx\_pclrsc -a -c databaseClusterName -u OsUser -D pgData -w workDir -n memberNodes

Item	Explanation	
databaseClusterName	Specify the FUJITSU Enterprise Postgres database cluster name.	
	The database cluster name is case-sensitive.	
	The database cluster name must be as follows:	
	- Up to 16 bytes, and	
	- The first byte must be an ASCII letter, and	
	<ul> <li>The other bytes must be ASCII alphanumeric characters, or underscores (_).</li> </ul>	

Item	Explanation
OsUser	Specify the operating system user (an OS user that will become the instance administrator) who can start the FUJITSU Enterprise Postgres database server.
pgData	Specify the absolute path of the data storage destination directory specified during database cluster creation.
workDir	Specify the directory for storing the temporary files required for FUJITSU Enterprise Postgres resource monitoring and state transition, and for storing the trace logs required in case of problems.
	If a parent directory that does not exist is specified, create the directory. The newly created directory path assumes the <i>OsUser</i> specified with the -u option as the owner, and 0700 as the privilege
	Required directory size is twice the size that is specified in thetrace-max-file-size option. If the default value is specified in thetrace-max-file-size option, 10240 KB x 2 is required for the directory size.
memberNodes	Specify the names of all nodes that make up the cluster system.
	Use a comma (,) to separate each node name.
	Suffix each cluster node name with the literal "RMS".
	Specification example:
	-n node1RMS,node2RMS

Other than the above, the following options can be specified:

Option	Explanation
db-user=name	Specify the FUJITSU Enterprise Postgres database superuser. The user specified in the -u option is the default.
response-timeout=seconds	In cluster operations, queries are regularly issued to the server to perform error detection and state monitoring.
	Specify a timeout time for queries issued for the heartbeat monitoring of the FUJITSU Enterprise Postgres database server.
	For queries, "SELECT 1" for the database "template1" is used. If 0 is specified, queries will not time out. The default is 0.
timeout-retry-count=count	Specify the upper limit for the number of retries when queries for monitoring time out.
	The default is 6 times.
trace-max-file-size	Specify the maximum size of the trace file. The default is 10240 KB.

Option	Explanation	
watch-interval=seconds	Specify the monitoring interval (the interval with which the queries are issued to the server) for a monitoring program to monitor the FUJITSU Enterprise Postgres database server. The default is 3 seconds.	



Refer to the "B.1 pgx\_pclrsc" for information on the pgx\_pclrsc command.

## 2.9 Configuring Storage Data Protection Using Transparent Data Encryption

If using transparent data encryption, refer to "Protecting Storage Data Using Transparent Data Encryption" in the Operation Guide to configure it. Take note on the following when configuring it:

### Location of the keystore file

Do not specify a shared disk for the keystore\_location parameter in postgresql.conf. Instead, specify a local directory with the same path on all nodes that comprise the cluster application.

If a shared disk is specified, startup from the cluster application will fail. This is because the keystore aks file, which is generated when automatic opening of the keystore is enabled, is obfuscated so that it can only be read on the node where it was generated, and therefore it cannot be shared across the shared disk.

### Distribution of the keystore file

The master encryption key settings must be common across all nodes. For this reason, configure the master encryption key settings on one node, and then copy the keystore file to each node. Also, if the master encryption key or passphrase is changed, you must also copy the keystore file to each node.

### Enabling automatic opening of the keystore

Automatic opening of the keystore must be enabled in all nodes that comprise the cluster application. Note that, when enabling the automatic opening of the keystore, only the directory in which the keystore file is stored will be referenced, therefore tasks such as starting and stopping the GDS volume are not required.

If it is not enabled, then startup from the cluster application will fail. This is because the startup process will not finish, since it waits for the manual input of the passphrase that opens the keystore.

## 2.10 Unmounting the File System

1. Unmount all file systems containing the resources required to create database clusters (such as the data storage destination directory).

# umount /mnt/swdsk1

2. Stop the GDS volume.

## 2.11 Creating a Cluster Application

Refer to "Appendix A Creating Resources and Creating/Modifying Cluster Applications".

## 2.12 Application Development

This section explains points to consider when creating applications during cluster operations.

## 2.12.1 Precautions when Developing Applications

To specify the IP address specified by the application, specify the IP address specified when creating the Gls resource or creating the takeover network resource. When state transition occurred, the operation can continue only by re-execute without changing the application.

Refer to "A.1.5 Creating Gls or Takeover Network Resources" for details.

### 2.12.2 Corrective Action with Application when State Transition Occurs

When the application running on the another node such as an application server, access the Symfoware Server, the following describes the error information to be returned when an abnormality or the like occurs in the node where FUJITSU Enterprise Postgres operates, and the actions.

#### JDBC driver

State		Error information(*1)	Action
Node failure or FUJITSU Enterprise Postgres	Failure occurs during access	57P01 08006	After the switch is complete, reestablish the connection, or re-execute
system failure	Accessed during node/system failure	08001	the application.
Switch to the standby node	Switched during access	57P01 08006	
	Accessed during switch	08001	

<sup>\*1:</sup> Return value of the getSQLState() of SQLException.

### **ODBC** driver

State		Error information(*1)	Action
Node failure or FUJITSU Enterprise Postgres	Failure occurs during access	57P01 08S01	After the switch is complete, reestablish the connection, or re-execute
system failure	Accessed during node/system failure	08001	the application.
Switch to the standby node	Switched during access	57P01 08S01	
	Accessed during switch	08001	

<sup>\*1:</sup> Return value of SQLSTATE.

### .NET Data Provider

State		Error information	Action
Node failure	Failure occurs during	57P01 (*1)	After the switch is
or	access	Empty string (*1)	complete, reestablish the

State		Error information	Action
FUJITSU Enterprise Postgres system failure		NullReferenceException is generated.	connection, or re-execute the application.
	Accessed during node/ system failure	Empty string (*1)	
Switch to the standby node	Switched during access	57P01 (*1) Empty string (*1) NullReferenceException is generated.	
	Accessed during switch	Empty string (*1)	

 $<sup>{\</sup>bf *1:}\ This\ is\ the\ return\ value\ of\ the\ PostgresException\ attribute\ SqlState.$ 

### C library(libpq)

State		Error information	Action
Node failure or FUJITSU Enterprise Postgres system failure	Failure occurs during access	PGRES_FATAL_ERROR(* 1) 57P01(*2) NULL(*2)	After the switch is complete, reestablish the connection, or reexecute the application.
	Accessed during node/system failure	CONNECTION_BAD(*3)	
Switch to the standby node	Switched during access	PGRES_FATAL_ERROR(* 1) 57P01(*2) NULL(*2)	
	Accessed during switch	CONNECTION_BAD(*3)	

<sup>\*1:</sup> Return value of PQresultStatus().

### Embedded SQL in C

State		Error information(*1)	Action
Node failure or FUJITSU Enterprise Postgres system failure	Failure occurs during access	57P01 57P02 YE000 26000 40001	After the switch is complete, reestablish the connection, or reexecute the application.
	Accessed during node/system failure	08001	

 $<sup>*2:</sup> Return\ value\ of\ PQresultErrorField()\ PG\_DIAG\_SQLSTATE.$ 

<sup>\*3:</sup> Return value of PQstatus().

State		Error information(*1)	Action
Switch to the standby node	Switched during	57P01	
	access	57P02	
		YE000	
		26000	
		40001	
	Accessed during switch	08001	

<sup>\*1:</sup> Return value of SQLSTATE

### **Embedded SQL in COBOL**

State		Error information(*1)	Action
Node failure or FUJITSU Enterprise Postgres system failure	Failure occurs during access	57P01 57P02 YE000	After the switch is complete, reestablish the connection, or reexecute the
		26000 40001	application.
	Accessed during node/system failure	08001	
Switch to the standby node	Switched during access	57P01 57P02 YE000 26000 40001	
	Accessed during switch	08001	

<sup>\*1:</sup> Return value of SQLSTATE.

## 2.13 Checking Operation

To ensure that the environment settings have been configured correctly, start, switch, and stop from Web-Based Admin View, and check the behavior.

To do a failover test, follow the procedure below.

 $1. \ \, Stop \ the \ FUJITSU \ Enterprise \ Postgres \ server \ processes \ by \ pg\_ctl \ command \ with \ immediate \ mode.$ 

Example)

```
# pg_ctl stop -m immediate
```

2. Check whether switching was done correctly.



FUJITSU Enterprise Postgres uses the su(1) command to periodically monitor its resources.

Therefore, in a RHEL environment, after a cluster system is built, operating system messages relating to the su(1) command will be periodically output to /var/log/messages during cluster system operation.

Refer to the relevant manual of your operating system for information on how to control these messages.

## Chapter 3 Failover Operation

This chapter explains the procedures involved in failover operation.

## 3.1 Adding Tablespaces

This section explains how to add tablespaces to a new file system.

This procedure is not required when you are adding tablespaces to an existing file system.

Perform the following steps when using a new shared disk class:

- Stop RMS on all nodes.
- Perform setup as described from "2.3 Creating a GDS Volume" to "2.4 Creating a File System".
- Modify the cluster application (refer to "A.2 Modifying Cluster Applications" for details).

Perform the following steps when using a new shared disk class:

- Stop RMS on all nodes.
- Perform the setup as described in "2.4 Creating a File System".

If a new GDS volume is to be used, the GDS volume must be created in advance.

- Modify the cluster application (refer to "A.2 Modifying Cluster Applications" for details).

## 3.2 Modifying Database Cluster Resources

This section explains the following operations:

- Displaying database cluster resource information
- Modifying database cluster resource content
- Deleting a database cluster resource

### Displaying database cluster resource information

Execute the pgx\_pclrsc command as shown below (note that if -c is not specified, the command lists all registered database cluster names):

```
# pgx_pclrsc -p -c databaseClusterName
```

### Modifying database cluster resource content

- 1. Stop RMS.
- 2. Execute the pgx\_pclrsc command as shown below (this is the simplest example):

```
# pgx_pclrsc -d -c databaseClusterName
# pgx_pclrsc -a -c databaseClusterName -u OsUser -D pgData -w workDir -n memberNodes
```



If you are modifying resource content, first delete the FUJITSU Enterprise Postgres resource from the cluster application, and then re-create it.



Refer to "B.1 pgx\_pclrsc" for information on the pgx\_pclrsc command.

### Deleting a database cluster resource

- 1. Stop RMS.
- 2. Delete the FUJITSU Enterprise Postgres resource from the cluster application.
- 3. Execute the pgx\_pclrsc command as shown below:

# pgx\_pclrsc -d -c databaseClusterName



The directory specified in the -w option of the pgx\_pclrsc command during registration will not be deleted. Use operating system commands to delete it if necessary.

## 3.3 Operations under Normal Circumstances

### Starting and stopping

For cluster operations, perform starting and stopping from the cluster application. If starting and stopping are performed by using the pg\_ctl command or WebAdmin during failover operation, the cluster application will misjudge that services have gone down, resulting in unexpected behaviors.

### **Switching**

There are two ways to switch between the active node and the standby node, as shown below.

Under normal circumstances, switch using mutual switch. In an emergency, for example when there is no response from the active node, perform forced switch. Forced switch should only be used in emergency scenarios because the differences in FUJITSU Enterprise Postgres stop modes when the active node is offline will cause the statistics to be initialized, and the load will increase as crash recovery is performed after the switch.

- Mutual switch

Refer to "3.5.2.4 Mutual Switch".

FUJITSU Enterprise Postgres is stopped with the "fast" mode of the pg\_ctl command.

- Forced switch

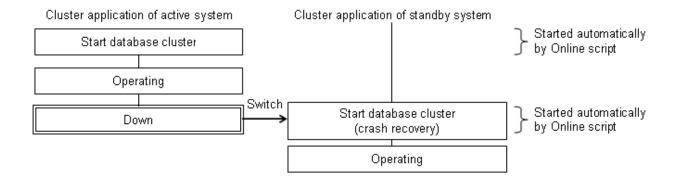
Right-click the cluster application object, and select [Switch] from the menu. From the switchable nodes displayed in the drop-down menu, select a target node, and switch the cluster application to that node.

FUJITSU Enterprise Postgres is stopped with the "immediate" mode of the pg\_ctl command.

## 3.4 Operation at State Transition

The workflow for operation at state transition is shown below.

The start operation during failover operation will be performed automatically by the Online script.



#### Non-transferrable feature

- All transactions being processed by the cluster application on the active system will be rolled back.
- The statistics FUJITSU Enterprise Postgres has collected will be initialized.

The statistics will be initialized in any of the following cases. Refer to the notes outlined in "Starting and Stopping an Instance" in the Operation Guide for details.

- When the node goes down because of a kernel panic, for example
- When forced switch is performed

Statistics are initialized because the "immediate" mode of the pg\_ctl command is used to stop FUJITSU Enterprise Postgres.

Refer to "3.3 Operations under Normal Circumstances" for details.

- When an error occurs in the cluster application resources

Statistics are initialized because the "immediate" mode of the pg\_ctl command is used to stop FUJITSU Enterprise Postgres to perform an immediate switch.

### **Errors during state transition**

If an error occurs in a FUJITSU Enterprise Postgres script during cluster application state transition, the following message will be displayed in the switchlog of Cluster Admin (the registered database cluster name will be displayed in the inst1 part shown below):

```
2017-05-14 11:08:22.060:(DET, 3): ERROR: FAULT REASON: Resource <Fsep_instl> transitioned to a Faulted state due to a script failure.
```

To display the switchlog, select [Tools] >> [View switchlog] in the [Cluster Admin] window.

If the resource displayed in the error message is a FUJITSU Enterprise Postgres resource name, a message indicating the cause will be output either before the message above or to the system log.

## 3.5 Maintenance Tasks

When you need to perform tasks such as configuration changes, patch application, and hardware parts replacement that may require restarting, there are two ways of node maintenance. Performing maintenance on both nodes by stopping both active and standby nodes, and performing maintenance on one node at a time by leaving the active node running to perform maintenance on the standby node first, and then switch the active node to perform maintenance on the other node. In this section, the former is referred to as simultaneous stopped node maintenance, whereas the latter is referred to as mutual switch maintenance.

Refer to the tasks outlined below in either "Figure 3.1 Workflow for simultaneous stopped node maintenance" or "Figure 3.2 Workflow for mutual switch maintenance", and perform your maintenance tasks accordingly.

Figure 3.1 Workflow for simultaneous stopped node maintenance

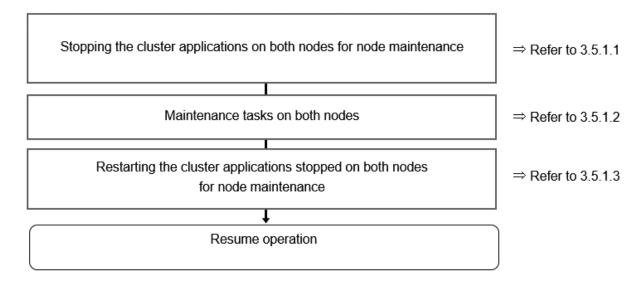
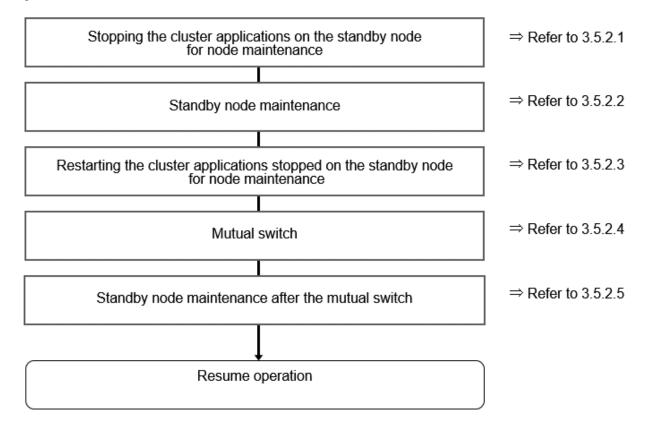


Figure 3.2 Workflow for mutual switch maintenance



## 3.5.1 Simultaneous Stopped Node Maintenance

This section explains the procedure to perform maintenance on both active and standby nodes simultaneously.

### 3.5.1.1 Stopping the Cluster Applications on Both Nodes for Node Maintenance

Follow the steps shown below to stop the cluster application on the active and standby nodes:

1. Start the [Cluster Admin] window.

### 2. Stop the cluster application.

On the RMS tree in the [Cluster Admin] window, right-click the standby cluster application to be stopped, and select [Offline] from the pop-up menu to stop it.

Stop the cluster application on both the active and standby nodes.

### 3. Stop RMS.

On the RMS tree in the [Cluster Admin] window, right-click the system node where the cluster application to be stopped exists, and select [Shutdown RMS] from the pop-up menu to stop RMS. To perform maintenance without stopping RMS, execute the hvutil - m on command to change the cluster application to maintenance mode.



- If applying the FUJITSU Enterprise Postgres update patches, RMS on the node (onto which the patches are applied) must be stopped. This is because the FUJITSU Enterprise Postgres monitoring process that runs when linked with PRIMECLUSTER must be stopped.
- By stopping RMS, all cluster applications on the system node where RMS is stopped will stop. When you have multiple cluster applications configured, ensure that stopping all other cluster applications will not cause any problems, before stopping RMS.

### 3.5.1.2 Maintenance Tasks on Both Nodes

Perform maintenance tasks such as configuration changes, patch application, hardware and parts replacement that may require restarting.

# 3.5.1.3 Restarting the Cluster Applications Stopped on Both Nodes for Node Maintenance

Follow the steps shown below to restart the cluster application on the active and standby nodes.

- 1. Start the [Cluster Admin] window.
- 2. If RMS has not been started, select [Tools] >> [Start RMS] in the [Cluster Admin] window, and select the node to be started and click [Ok].

Note that if the cluster applications were stopped for node maintenance using the hvutil -m on command, execute the hvutil -m off command instead of starting RMS.

3. Start the cluster application.

On the RMS tree in the [Cluster Admin] window, right-click the standby cluster application, and from the pop-up menu, if you are starting the node as the active node, select [Online], or if you are starting it as the standby node, select [Standby]. If the state is [Fault], then select [Clear fault].

This operation is not required if the cluster application has been configured to start automatically when RMS is started.

### 3.5.2 Mutual Switch Maintenance

This section explains the procedure to perform maintenance on the standby node while running the active node at the same time.

### 3.5.2.1 Stopping the Cluster Applications on the Standby Node for Node Maintenance

Follow the steps shown below to stop the cluster application on the standby node only:

1. Edit the /opt/SMAW/SMAWRrms/bin/hvenv.local file, and add "export HV\_RCSTART=0".

If the hvenv.local file does not exist, then create it.

```
# vi /opt/SMAW/SMAWRrms/bin/hvenv.local
Add export HV_RCSTART=0
```

2. On the RMS tree in the [Cluster Admin] window, right-click the cluster application on the standby node to be stopped, and select [Offline] from the pop-up menu to stop it.

3. On the RMS tree in the [Cluster Admin] window, right-click the system node (where the cluster application on the standby node that was stopped exists), and select [Shutdown RMS]. At this stage, select [Stop all Apps] as the option, and stop RMS.



- Ensure that the steps up to step 3 have been completed properly before performing the maintenance tasks. If the active node goes down at the point when steps 1 and 2 have been completed, FUJITSU Enterprise Postgres will be started on the supposedly stopped standby node, and it becomes the active node instead. As a result, this will compete with the maintenance tasks and unexpected behavior may occur.
- If applying the FUJITSU Enterprise Postgres update patches, RMS on the node (onto which the patches are applied) must be stopped. This is because the FUJITSU Enterprise Postgres monitoring process that runs when linked with PRIMECLUSTER must be stopped.
- By stopping RMS, all cluster applications on the system node where RMS is stopped will stop. When you have multiple cluster applications configured, ensure that stopping all other cluster applications will not cause any problems, before stopping RMS.

### 3.5.2.2 Standby Node Maintenance

Perform maintenance tasks such as patch application and hardware parts replacement.

# 3.5.2.3 Restarting the Cluster Applications Stopped on the Standby Node for Node Maintenance

Follow the steps shown below to restart the cluster application on the standby node only (that was stopped for maintenance):

1. Edit the /opt/SMAW/SMAWRrms/bin/hvenv.local file, and delete "export HV\_RCSTART=0".

```
# vi /opt/SMAW/SMAWRrms/bin/hvenv.local
Delete export HV_RCSTART=0
```

- 2. Select [Tools]-[Start RMS] in the [Cluster Admin] window, and select the node to be started and click [Ok].
- 3. On the RMS tree in the [Cluster Admin] window, right-click the cluster application to be started as the standby node, and select [Standby] from the pop-up menu to start it. If the state is [Fault], then select [Clear fault].

### 3.5.2.4 Mutual Switch

To perform mutual switch between the active and standby nodes, perform the steps shown below:

- 1. Select and right-click the cluster application on the active node, and select [Offline] from the pop-up menu to stop it.
- 2. Select and right-click the cluster application on the standby node, and select [Online] from the pop-up menu to start it.
- 3. Select and right-click the cluster application that was previously on the active node, and select [Standby] from the pop-up menu to start it as standby.

### 3.5.2.5 Standby Node Maintenance after the Mutual Switch

To perform maintenance tasks on the new standby node after mutual switch, perform the steps from "3.5.2.1 Stopping the Cluster Applications on the Standby Node for Node Maintenance" to "3.5.2.3 Restarting the Cluster Applications Stopped on the Standby Node for Node Maintenance".

## Chapter 4 Procedures Required after a Failover Error

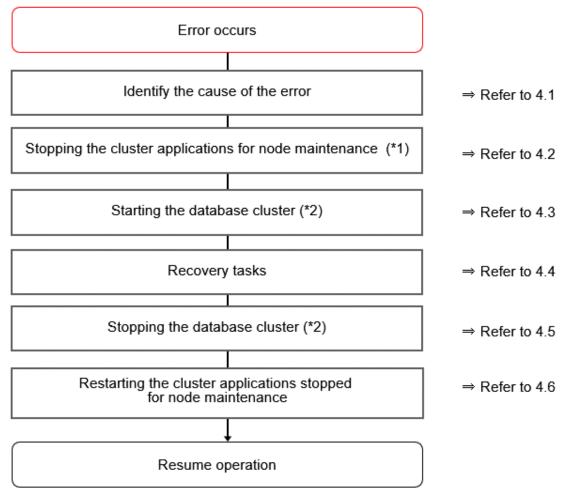
This chapter explains how to perform recovery when an abnormality occurs during failover operation.

When an error occurs during failover operation, refer to the system log to identify the cause and stop the RMS to perform recovery, separately from the RMS management. Normal operations can be resumed by restarting the RMS once the recovery is completed.

Note that the database cluster may be started or stopped without using cluster applications when performing recovery.

Refer to the tasks outlined below in "Figure 4.1 Operation flow when a failover error occurs", and perform recovery.

Figure 4.1 Operation flow when a failover error occurs



<sup>\*1)</sup> If the machine is already stopped when the error occurred, this step is not necessary.

## 4.1 Identifying the Cause of an Error

When an error occurs during failover operation, refer to the system log to identify the cause. When cluster applications fail to start, refer to the server log file to identify the cause.

The server log file is a file created under the directory specified in the -w option when executing the pgx\_pclrsc command. This file has almost the same content as the file specified in the -l option when executing the pg\_ctl command to start the database cluster, and errors that are not output to the system log are output to these files, including startup errors.

<sup>\*2)</sup> Depending on the recovery tasks, start or stop the database cluster as required.

## 4.2 Stopping the Cluster Applications for Node Maintenance

When an error occurs during failover operation, stop the RMS to perform recovery, separately from the RMS management. Refer to "3.5 Maintenance Tasks" for information on how to stop cluster applications for node maintenance.

Note that this task is not required in cases such as when the RMS is stopped because the machine is stopped due to an error.

## 4.3 Starting the Database Cluster

To perform recovery when an error occurs, it is necessary to start the database cluster without using cluster applications.

Follow the procedure below to start the database cluster if required.

Check if the database cluster needs to be started by referring to the recovery tasks written in "Actions When an Error Occurs" in the Operation Guide.

- 1. Start the GDS volume.
- 2. Mount the necessary file systems such as those where the data storage destination directory is placed.
- 3. Start the database cluster by using the pg\_ctl command. Refer to "Using Server Commands" in "Starting an Instance and Creating a Database" in the Operation Guide for information on how to start the database cluster.

## 4.4 Recovery Tasks

Refer to "Actions When an Error Occurs" in the Operation Guide for information on how to perform recovery. If there are two ways available to perform recovery; using WebAdmin or commands, use commands.

## 4.5 Stopping the Database Cluster

If a database cluster is started, stop the database cluster.

- 1. Stop the database cluster using the pg\_ctl command. Refer to "Using Server Commands" in "Starting an Instance and Creating a Database" in the Operation Guide" for information on how to stop the database cluster.
- 2. Unmount the file systems such as those where the data storage destination directory is placed.
- 3. Stop the GDS volume.

# 4.6 Restarting the Cluster Applications Stopped for Node Maintenance

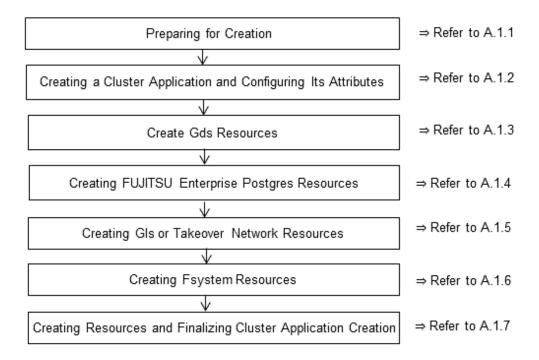
Refer to "3.5 Maintenance Tasks" to restart the cluster applications stopped for node maintenance.

## Appendix A Creating Resources and Creating/Modifying Cluster Applications

This appendix explains how to create and modify cluster applications that include FUJITSU Enterprise Postgres resources in PRIMECLUSTER.

## A.1 Creating Resources and Cluster Applications

Use the following procedure to create resources and a cluster application:



## A.1.1 Preparing for Creation

1. Execute the hvgdsetup command for all shared disk classes to be used by all file systems to be registered in the cluster application.

```
# /opt/SMAW/SMAWRrms/bin/hvgdsetup -a class0001
```

- 2. From the Web-Based Admin View, stop RMS on all nodes.
- 3. Execute the hvw command.

# /opt/SMAW/SMAWRrms/bin/hvw



Refer to the PRIMECLUSTER Installation and Administration Guide for information on the hygdsetup command.

## A.1.2 Creating a Cluster Application and Configuring Its Attributes

1. Select [Application-Create].

pcl-vm13: Main configuration menu, current configuration: config
No RMS active in the cluster
1) HELP 10) Configuration-Remove

2) QUIT 11) Configuration-Freeze 3) Application-Create 12) Configuration-Thaw 4) Application-Edit 13) Configuration-Edit-Global-Settings 5) Application-Remove 14) Configuration-Consistency-Report 6) Application-Clone 15) Configuration-ScriptExecution 7) Configuration-Generate 16) RMS-CreateMachine 8) Configuration-Activate 17) RMS-RemoveMachine 9) Configuration-Copy Choose an action: 3

### 2. Select [STANDBY].

Creation: Application type selection menu:

1) HELP 7) SCALABLE

2) QUIT 8) STANDBY

3) RETURN

4) OPTIONS

5) DEMO

6) GENERIC
Application Type: 8

#### 3. Select [Machines+Basics].

Settings of turnkey wizard "STANDBY" (APP1:not yet consistent)

1) HELP

4) REMOVE+EXIT

2) NO-SAVE+EXIT

5) ApplicationName=APP1

3) SAVE+EXIT

6) Machines+Basics(-)

Choose the setting to process: 6



Although ApplicationName can be set to any name within the range of constraints of PRIMECLUSTER, do not set it to the database cluster name specified when registering a database cluster to PRIMECLUSTER using the pgx\_pclrsc command. If the same name is set, RMS will fail to start.

.....

### 4. Select [AdditionalMachine].

Machines+Basics (appl:consistent) 1) HELP 14) (AutoStartUp=no) 2) -15) (AutoSwitchOver=No) 3) SAVE+EXIT 16) (PreserveState=no) 4) REMOVE+EXIT 17) (PersistentFault=0) 5) AdditionalMachine 18) (ShutdownPriority=) 6) AdditionalConsole 19) (OnlinePriority=) 7) Machines[0]=pcl-vm13RMS 20) (StandbyTransitions=) 8) (PreCheckScript=) 21) (LicenseToKill=no) 9) (PreOnlineScript=) 22) (AutoBreak=yes) 10) (PostOnlineScript=) 23) (AutoBreakMaintMode=no) 11) (PreOfflineScript=) 24) (HaltFlag=no) 12) (OfflineDoneScript=) 25) (PartialCluster=0) 13) (FaultScript=) 26) (ScriptTimeout=) Choose the setting to process: 5

### 5. Select a standby node.

1) HELP
2) RETURN
3) pcl-vml3RMS
4) pcl-vml4RMS
Choose a machine for this application: 4

6. Ensure that all nodes that make up the cluster application are displayed in [Machines]. Note that the item displayed in Machines[0] needs to be the active node, and Machines[1] to be the standby node. If these are not set correctly, modify the setting.

```
Machines+Basics (appl:consistent)
1) HELP
                                       15) (AutoStartUp=no)
2) -
                                       16) (AutoSwitchOver=No)
3) SAVE+EXIT
                                       17) (PreserveState=no)
4) REMOVE+EXIT
                                      18) (PersistentFault=0)
5) AdditionalMachine
                                      19) (ShutdownPriority=)
                                      20) (OnlinePriority=)
6) AdditionalConsole
7) Machines[0]=pcl-vm13RMS
                                      21) (StandbyTransitions=)
8) Machines[1]=pcl-vm14RMS
                                      22) (LicenseToKill=no)
9) (PreCheckScript=)
                                       23) (AutoBreak=yes)
10) (PreOnlineScript=)
                                       24) (AutoBreakMaintMode=no)
11) (PostOnlineScript=)
                                       25) (HaltFlag=no)
12) (PreOfflineScript=)
                                       26) (PartialCluster=0)
13) (OfflineDoneScript=)
                                       27) (ScriptTimeout=)
14) (FaultScript=)
Choose the setting to process:
```

7. Configure attributes for cluster applications from step 8.

Refer to the PRIMECLUSTER Installation and Administration Guide for information on the attributes for cluster applications.

Configure the attributes as shown below:

Attribute	Value
AutoStartUp	Yes
(Automatically starts cluster applications when the RMS is started.)	
AutoSwitchOver	- HostFailure
(Automatically switches nodes when a fault	- ResourceFailure
occurs in cluster applications.)	- ShutDown
PersistentFault	0 or 1
(Selects whether to keep the Faulted state of the cluster application after restarting the RMS.) (*1)	
OnlinePriority	0 or 1
(Selects whether to use the same online nodes as those used before all clusters were restarted or follow the configured order for nodes, when all clusters were restarted.) (*2)	
StandbyTransitions	- ClearFaultRequest
(Timing to perform standby state transition.)	- StartUp
	- SwitchRequest
HaltFlag	Yes
(Selects whether to terminate the node forcibly when another error occurs during the Fault process.)	

<sup>\*1:</sup> If 0 is specified, the Faulted state is automatically canceled when the RMS is started. If 1 is specified, the Faulted state is kept.

<sup>\*2:</sup> If 0 is specified, cluster applications will become Online in accordance with the node setting order when all clusters are restarted. If 1 is specified, cluster applications will become Online on nodes where the cluster applications were Online before the restart.

### 8. Select [AutoStartUp].

```
Machines+Basics (appl:consistent)
1) HELP
                                         15) (AutoStartUp=no)
2) -
                                         16) (AutoSwitchOver=No)
3) SAVE+EXIT
                                         17) (PreserveState=no)
4) REMOVE+EXIT
                                         18) (PersistentFault=0)
5) AdditionalMachine
                                        19) (ShutdownPriority=)
                                        20) (OnlinePriority=)
6) AdditionalConsole
                                   21) (StandbyTransitions=)
22) (LicenseToKill=no)
7) Machines[0]=pcl-vm13RMS
8) Machines[1]=pcl-vm14RMS
9) (PreCheckScript=)
                                         23) (AutoBreak=yes)
                                        24) (AutoBreakMaintMode=no)25) (HaltFlag=no)
10) (PreOnlineScript=)
11) (PostOnlineScript=)
                                        26) (PartialCluster=0)
12) (PreOfflineScript=)
13) (OfflineDoneScript=)
                                        27) (ScriptTimeout=)
14) (FaultScript=)
Choose the setting to process: 15
```

### 9. Select [yes].

```
1) HELP
2) RETURN
3) yes
4) no
Set the AutoStartUp mode: 3
```

### 10. Select [AutoSwitchOver].

```
Machines+Basics (appl:consistent)
1) HELP
                                      15) (AutoStartUp=yes)
2) -
                                      16) (AutoSwitchOver=No)
3) SAVE+EXIT
                                      17) (PreserveState=no)
4) REMOVE+EXIT
                                     18) (PersistentFault=0)
5) AdditionalMachine
                                     19) (ShutdownPriority=)
6) AdditionalConsole
                                     20) (OnlinePriority=)
7) Machines[0]=pcl-vm13RMS
                                     21) (StandbyTransitions=)
8) Machines[1]=pcl-vm14RMS
                                     22) (LicenseToKill=no)
9) (PreCheckScript=)
                                     23) (AutoBreak=yes)
10) (PreOnlineScript=)
                                     24) (AutoBreakMaintMode=no)
11) (PostOnlineScript=)
                                     25) (HaltFlag=no)
12) (PreOfflineScript=)
                                     26) (PartialCluster=0)
                                     27) (ScriptTimeout=)
13) (OfflineDoneScript=)
14) (FaultScript=)
Choose the setting to process: 16
```

## 11. Select [HOSTFAILURE], then [RESOURCEFAILURE], then [SHUTDOWN], and then select [SAVE+RETURN] when the following window is displayed:

```
Set flags for AutoSwitchOver: Currently set: HOSTFAILURE, RESOURCEFAILURE, SHUTDOWN (HRS)

1) HELP
4) DEFAULT
7) NOT:RESOURCEFAILURE(R)

2) -
5) NO(N)
8) NOT:SHUTDOWN(S)

3) SAVE+RETURN
6) NOT:HOSTFAILURE(H)

Choose one of the flags: 3
```

### 12. Select [PersistentFault].

```
Machines+Basics (appl:consistent)

1) HELP

2) -

3) SAVE+EXIT

4) REMOVE+EXIT

5) AdditionalMachine

6) AdditionalConsole
```

```
7) Machines[0]=pcl-vm13RMS
8) Machines[1]=pcl-vm14RMS
9) (PreCheckScript=)
10) (PreOnlineScript=)
11) (PostOnlineScript=)
12) (PreOfflineScript=)
13) (OfflineDoneScript=)
14) (FaultScript=)
15) (AutoStartUp=yes)
16) (AutoSwitchOver=HostFailure | ResourceFailure | ShutDown)
17) (PreserveState=no)
18) (PersistentFault=0)
19) (ShutdownPriority=)
20) (OnlinePriority=)
21) (StandbyTransitions=)
22) (LicenseToKill=no)
23) (AutoBreak=yes)
24) (AutoBreakMaintMode=no)
25) (HaltFlag=no)
26) (PartialCluster=0)
27) (ScriptTimeout=)
Choose the setting to process: 18
```

### 13. Select [0] or [1] (the example below selects [1]).

```
1) HELP
2) RETURN
3) 0
4) 1
Enable persistent fault feature for this application: 4
```

### 14. Select [OnlinePriority].

```
Machines+Basics (appl:consistent)
1) HELP
2) -
3) SAVE+EXIT
4) REMOVE+EXIT
5) AdditionalMachine
6) AdditionalConsole
7) Machines[0]=pcl-vm13RMS
8) Machines[1]=pcl-vm14RMS
9) (PreCheckScript=)
10) (PreOnlineScript=)
11) (PostOnlineScript=)
12) (PreOfflineScript=)
13) (OfflineDoneScript=)
14) (FaultScript=)
15) (AutoStartUp=yes)
16) (AutoSwitchOver=HostFailure ResourceFailure ShutDown)
17) (PreserveState=no)
18) (PersistentFault=1)
19) (ShutdownPriority=)
20) (OnlinePriority=)
21) (StandbyTransitions=)
22) (LicenseToKill=no)
23) (AutoBreak=yes)
24) (AutoBreakMaintMode=no)
25) (HaltFlag=no)
26) (PartialCluster=0)
27) (ScriptTimeout=)
Choose the setting to process: 20
```

15. Select [0] or [1] (the example below selects [1]).

```
1) HELP
2) RETURN
3) 0
4) 1
Enable Online Priority (Active Standby) feature for this application: 4
```

16. Select [StandbyTransitions].

```
Machines+Basics (appl:consistent)
1) HELP
2) -
3) SAVE+EXIT
4) REMOVE+EXIT
5) AdditionalMachine
6) AdditionalConsole
7) Machines[0]=pcl-vm13RMS
8) Machines[1]=pcl-vm14RMS
9) (PreCheckScript=)
10) (PreOnlineScript=)
11) (PostOnlineScript=)
12) (PreOfflineScript=)
13) (OfflineDoneScript=)
14) (FaultScript=)
15) (AutoStartUp=yes)
16) (AutoSwitchOver=HostFailure|ResourceFailure|ShutDown)
17) (PreserveState=no)
18) (PersistentFault=1)
19) (ShutdownPriority=)
20) (OnlinePriority=1)
21) (StandbyTransitions=)
22) (LicenseToKill=no)
23) (AutoBreak=yes)
24) (AutoBreakMaintMode=no)
25) (HaltFlag=no)
26) (PartialCluster=0)
27) (ScriptTimeout=)
Choose the setting to process: 21
```

17. Select [CLEARFAULTREQUEST], then [STARTUP], then [SWITCHREQUEST], and then select [SAVE+RETURN] when the following window is displayed:

```
Set flags for StandbyTransitions: Currently set: ClearFaultRequst,StartUp,SwitchRequest (CTW)

1) HELP
5) NOT:CLEARFAULTREQUEST(C)
2) - 6) NOT:STARTUP(T)
3) SAVE+RETURN
7) NOT:SWITCHREQUEST(W)
4) NO(N)
Choose one of the flags: 3
```

18. Select [HaltFlag].

```
Machines+Basics (appl:consistent)
1) HELP
2) -
3) SAVE+EXIT
4) REMOVE+EXIT
5) AdditionalMachine
6) AdditionalConsole
7) Machines[0]=pcl-vm13RMS
8) Machines[1]=pcl-vm14RMS
9) (PreCheckScript=)
10) (PreOnlineScript=)
11) (PostOnlineScript=)
```

```
12) (PreOfflineScript=)
13) (OfflineDoneScript=)
14) (FaultScript=)
15) (AutoStartUp=yes)
16) (AutoSwitchOver=HostFailure ResourceFailure ShutDown)
17) (PreserveState=no)
18) (PersistentFault=1)
19) (ShutdownPriority=)
20) (OnlinePriority=1)
21) (StandbyTransitions=ClearFaultRequest|StartUp|SwitchRequest)
22) (LicenseToKill=no)
23) (AutoBreak=yes)
24) (AutoBreakMaintMode=no)
25) (HaltFlag=no)
26) (PartialCluster=0)
27) (ScriptTimeout=)
Choose the setting to process: 25
```

### 19. Select [yes].

```
1) HELP
2) RETURN
3) yes
4) no
Set the Halt mode: 3
```

### 20. Ensure that all attributes are displayed as set in the previous steps, and then select [SAVE+EXIT].

```
Machines+Basics (appl:consistent)
1) HELP
2) -
3) SAVE+EXIT
4) REMOVE+EXIT
5) Additional Machine
6) AdditionalConsole
7) Machines[0]=pcl-vm13RMS
8) Machines[1]=pcl-vm14RMS
9) (PreCheckScript=)
10) (PreOnlineScript=)
11) (PostOnlineScript=)
12) (PreOfflineScript=)
13) (OfflineDoneScript=)
14) (FaultScript=)
15) (AutoStartUp=yes)
16) (AutoSwitchOver=HostFailure | ResourceFailure | ShutDown)
17) (PreserveState=no)
18) (PersistentFault=1)
19) (ShutdownPriority=)
20) (OnlinePriority=1)
21) (StandbyTransitions=ClearFaultRequest|StartUp|SwitchRequest)
22) (LicenseToKill=no)
23) (AutoBreak=yes)
24) (AutoBreakMaintMode=no)
25) (HaltFlag=yes)
26) (PartialCluster=0)
27) (ScriptTimeout=)
Choose the setting to process: 3
```

### 21. The following window will be displayed:

```
Settings of turnkey wizard "STANDBY" (APP1:not yet consistent)

1) HELP

10) Enterprise-Postgres(-)

2) -

11) Symfoware(-)

3) SAVE+EXIT

12) Procedure:SystemState3(-)
```

```
4) - 13) Procedure:SystemState2(-)
5) ApplicationName=APP1 14) Gls:Global-Link-Services(-)
6) Machines+Basics(appl) 15) IpAddresses(-)
7) CommandLines(-) 16) LocalFileSystems(-)
8) Procedure:Application(-) 17) Gds:Global-Disk-Services(-)
9) Procedure:BasicApplication(-)
Choose the setting to process:
```

### A.1.3 Creating Gds Resources

1. Select [Gds:Global-Disk-Services].

```
Settings of turnkey wizard "STANDBY" (APP1:not yet consistent)
1) HELP
                                       10) Enterprise-Postgres(-)
2) -
                                       11) Symfoware(-)
3) SAVE+EXIT
                                       12) Procedure:SystemState3(-)
4) -
                                       13) Procedure:SystemState2(-)
5) ApplicationName=APP1
                                      14) Gls:Global-Link-Services(-)
6) Machines+Basics(appl)
                                      15) IpAddresses(-)
7) CommandLines(-)
                                      16) LocalFileSystems(-)
                                       17) Gds:Global-Disk-Services(-)
8) Procedure:Application(-)
9) Procedure:BasicApplication(-)
Choose the setting to process:17
```

2. Select [AdditionalDiskClass].

```
Volume management (Gds_APP1:not yet consistent)

1) HELP

4) REMOVE+EXIT

7) (StandbySupport=no)

2) -

5) AdditionalDiskClass

8) (AutoRecover=no)

3) SAVE+EXIT

6) (ClassNameFilter=)

9) (Timeout=1800)

Choose the setting to process: 5
```

3. Select a class.

```
1) HELP
2) RETURN
3) FREECHOICE
4) class0001
Choose the disk class: 4
```

4. Follow the steps above for all shared disk classes to be used by all file systems to be registered in the cluster application. When doing so, configure the MONITORONLY attribute as shown in the following table:

Purpose of file system	MONITORONLY attribute
Data storage destination directory	No
Backup data storage destination directory	Yes
Table space	No
Transaction log storage destination directory	No

To set the MONITORONLY attribute to "Yes", select "MONITORONLY".

```
Set a flag for the disk class: class0001
Currently set:
1) HELP 5) MONITORONLY(M)
2) -
```

```
3) SAVE+RETURN
4) DEFAULT
Choose additionally one of the flags:
```

5. To configure the disk class, select [StandbySupport] and specify "no" - this means that upon failure, the disk will switch according to the purpose of the class.

```
Volume management (Gds_APP1:consistent)

1) HELP
6) DiskClasses[0]=class0001
2) - 7) (ClassNameFilter=)
3) SAVE+EXIT
8) (StandbySupport=no)
4) REMOVE+EXIT
9) (AutoRecover=no)
5) AdditionalDiskClass
10) (Timeout=1800)
Choose the setting to process: 8
```

6. Select [SAVE+EXIT].

```
Volume management (Gds_APP1:consistent)

1) HELP
6) DiskClasses[0]=class0001
2) - 7) (ClassNameFilter=)
3) SAVE+EXIT
8) (StandbySupport=no)
4) REMOVE+EXIT
9) (AutoRecover=no)
5) AdditionalDiskClass
10) (Timeout=1800)
Choose the setting to process: 3
```

### A.1.4 Creating FUJITSU Enterprise Postgres Resources

1. Select [Enterprise-Postgres].

```
Settings of turnkey wizard "STANDBY" (APP1:consistent)
1) HELP
                                       10) Enterprise-Postgres(-)
2) -
                                       11) Symfoware(-)
3) SAVE+EXIT
                                       12) Procedure:SystemState3(-)
4) -
                                       13) Procedure:SystemState2(-)
5) ApplicationName=APP1
                                       14) Gls:Global-Link-Services(-)
6) Machines+Basics(appl)
                                       15) IpAddresses(-)
7) CommandLines(-)
                                       16) LocalFileSystems(-)
8) Procedure:Application(-)
                                       17) Gds:Global-Disk-Services(Gds APP1)
9) Procedure:BasicApplication(-)
Choose the setting to process:10
```

2. Select [AdditionalResource].

```
Resource (Fsep_APP1:not yet consistent)

1) HELP

4) REMOVE+EXIT

2) -

5) AdditionalResource

3) SAVE+EXIT

Choose the setting to process: 5
```

3. Select the database cluster name specified in the pgx\_pclrsc command in "2.8 Registering Resource Information for the FUJITSU Enterprise Postgres Database Cluster".

```
1) HELP
2) RETURN
3) FREECHOICE
4) inst1
Choose the resource: 4
```

4. Select [SAVE+EXIT].

```
Resource (Fsep_APP1:consistent)

1) HELP

6) ResourceName[0]=inst1

2) -

7) Flags[0]=T900
```

```
3) SAVE+EXIT
8) OnlineTimeout[0]=3600
4) REMOVE+EXIT
9) OfflineTimeout[0]=1800
5) AdditionalResource
10) FaultScript[0]=''
Choose the setting to process: 3
```

### A.1.5 Creating GIs or Takeover Network Resources

Create Gls or takeover network resources to use takeover network resources. This section provides an example for configuring Gl resources.

......



- Refer to the PRIMECLUSTER Installation and Administration Guide for information on creating takeover network resources.
- Refer to the PRIMECLUSTER Installation and Administration Guide for information on creating Gl resources.
- 1. Select [Gls:Global-Link-Services].

```
Settings of turnkey wizard "STANDBY" (APP1:consistent)
1) HELP
                                      10) Enterprise-Postgres(Fsep_APP1)
2) -
                                      11) Symfoware(-)
3) SAVE+EXIT
                                      12) Procedure:SystemState3(-)
4) -
                                      13) Procedure:SystemState2(-)
5) ApplicationName=APP1
                                      14) Gls:Global-Link-Services(-)
6) Machines+Basics(app1)
                                      15) IpAddresses(-)
7) CommandLines(-)
                                      16) LocalFileSystems(-)
8) Procedure:Application(-)
                                      17) Gds:Global-Disk-Services(Gds_APP1)
9) Procedure:BasicApplication(-)
Choose the setting to process:14
```

### 2. Select [AdditionalTakeoverIpaddress].

```
Gls (Gls_APP1:not yet consistent)

1) HELP

4) REMOVE+EXIT

2) -

5) AdditionalTakeoverIpaddress

3) SAVE+EXIT

6) (Timeout=60)

Choose the setting to process: 5
```

### 3. Select a takeover IP.

```
1) HELP
2) RETURN
3) FREECHOICE
4) takeoverip
Choose a takeover IP address for Gls: 4
```

### 4. Select [SAVE+RETURN].

### 5. Select [SAVE+EXIT].

```
Gls (Gls_APP1:consistent)

1) HELP

5) AdditionalTakeoverIpaddress

2) -

6) TakeoverIpaddress[0]=N,takeoverip

3) SAVE+EXIT

7) (Timeout=60)
```

```
4) REMOVE+EXIT Choose the setting to process: 3
```

### A.1.6 Creating Fsystem Resources

1. Select [LocalFileSystem].

```
Settings of turnkey wizard "STANDBY" (APP1:consistent)
1) HELP
                                       10) Enterprise-Postgres(Fsep_APP1)
2) -
                                       11) Symfoware(-)
3) SAVE+EXIT
                                       12) Procedure:SystemState3(-)
4) -
                                       13) Procedure:SystemState2(-)
5) ApplicationName=APP1
                                       14) Gls:Global-Link-Services(Gls_APP1)
6) Machines+Basics(appl)
                                      15) IpAddresses(-)
7) CommandLines(-)
                                       16) LocalFileSystems(-)
8) Procedure:Application(-)
                                       17) Gds:Global-Disk-Services(Gds_APP1)
9) Procedure:BasicApplication(-)
Choose the setting to process:16
```

2. Select [AdditionalMountPoint].

```
File systems (Lfs_APP1:not yet consistent)

1) HELP
4) REMOVE+EXIT
7) (Timeout=180)

2) -
5) AdditionalMountPoint

3) SAVE+EXIT
6) (Filter=)

Choose the setting to process: 5
```

3. Select the mount point of the file system.

For all mount points assigned to the data storage destination directory, the backup data storage destination directory, the tablespace directory, and the transaction log storage destination directory, repeat Steps 3 to 5. Additionally, to register multiple file systems for monitoring, repeat Steps 3 to 5 before proceeding to configure settings for file systems used for the job.

```
1) HELP 5) /mnt/swdsk1
2) RETURN 6) /mnt/swdsk2
3) FREECHOICE
4) ALL
Choose a mount point: 5
```

4. Configure the MONITORONLY attribute.

Purpose of the file system	MONITORONLY attribute
Data storage destination directory	No
Backup data storage destination directory	Yes
Table space	No
Transaction log storage destination directory	No

Select MONITORONLY to specify Yes to the MONITORONLY attribute.

```
Set flags for mount point: /mnt/swdsk1 Currently set: LOCAL, AUTORECOVER (LA)

1) HELP 4) DEFAULT 7) SHARE(S)

2) - 5) SYNC(Y) 8) MONITORONLY(M)

3) SAVE+RETURN 6) NOT: AUTORECOVER(A)

Choose one of the flags: 8
```

### 5. Select [SAVE+RETURN].

```
Set flags for mount point: /mnt/swdskl Currently set:

LOCAL,AUTORECOVER, MONITORONLY (LAM)

1) HELP
4) DEFAULT
7) SHARE(S)

2) -
5) SYNC(Y)
8) NOT: MONITORONLY(M)

3) SAVE+RETURN
6) NOT: AUTORECOVER(A)

Choose one of the flags: 3
```

6. After configuring all mount points, select [SAVE+EXIT].

```
File systems (Lfs_APP1:consistent)

1) HELP

6) MountPoints[0]=LA:/mnt/monitor1

2) -

7) MountPoints[1]=LA:/mnt/monitor2

3) SAVE+EXIT

8) (Filter=)

4) REMOVE+EXIT

9) (Timeout=180)

5) AdditionalMountPoint

Choose the setting to process: 3
```

#### 7. Select [SAVE+EXIT].

```
Settings of turnkey wizard "STANDBY" (APP1:consistent)
1) HELP
                                        10) Enterprise-Postgres(Fsep_APP1)
2) -
                                        11) Symfoware(-)
3) SAVE+EXIT
                                        12) Procedure:SystemState3(-)
4) -
                                       13) Procedure:SystemState2(-)
5) ApplicationName=APP1
                                       14) Gls:Global-Link-Services(Gls_APP1)
6) Machines+Basics(appl)
                                       15) IpAddresses(-)
                                       16) LocalFileSystems(Lfs_APP1)
7) CommandLines(-)
8) Procedure:Application(-)
                                       17) Gds:Global-Disk-Services(Gds_APP1)
9) Procedure:BasicApplication(-)
Choose the setting to process:3
```



Refer to the PRIMECLUSTER Installation and Administration Guide for information on creating Fsystem resources.

## A.1.7 Creating Resources and Finalizing Cluster Application Creation

### 1. Select [Configuration-Generate].

```
pcl-vm13: Main configuration menu, current configuration: config
No RMS active in the cluster
1) HELP
                                       10) Configuration-Remove
2) QUIT
                                       11) Configuration-Freeze
3) Application-Create
                                       12) Configuration-Thaw
4) Application-Edit
                                       13) Configuration-Edit-Global-Settings
5) Application-Remove
                                      14) Configuration-Consistency-Report
 6) Application-Clone
                                      15) Configuration-ScriptExecution
7) Configuration-Generate
                                       16) RMS-CreateMachine
 8) Configuration-Activate
                                       17) RMS-RemoveMachine
9) Configuration-Copy
Choose an action: 7
```

#### 2. Select [Configuration-Activate].

If processing is successful, "The activation has finished successfully." will be displayed. If this message is not displayed, there is an issue with the information or configuration that was set. In such a case, review the changed information.

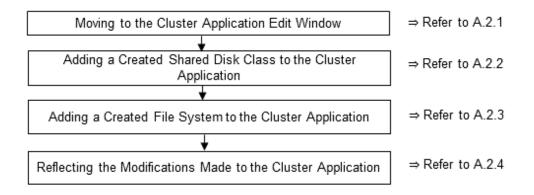
```
pcl-vm13: Main configuration menu, current configuration: config
No RMS active in the cluster
1) HELP 10) Configuration-Remove
```

- 2) QUIT 11) Configuration-Freeze 3) Application-Create 12) Configuration-Thaw 4) Application-Edit 13) Configuration-Edit-Global-Settings 5) Application-Remove 14) Configuration-Consistency-Report 6) Application-Clone 15) Configuration-ScriptExecution 7) Configuration-Generate 16) RMS-CreateMachine 8) Configuration-Activate 17) RMS-RemoveMachine 9) Configuration-Copy Choose an action: 8
- 3. Select [QUIT].

pcl-vm13: Main configuration menu, current configuration: config No RMS active in the cluster 1) HELP 10) Configuration-Remove 2) OUIT 11) Configuration-Freeze 3) Application-Create 12) Configuration-Thaw 4) Application-Edit 13) Configuration-Edit-Global-Settings 5) Application-Remove 14) Configuration-Consistency-Report 6) Application-Clone 15) Configuration-ScriptExecution 7) Configuration-Generate 16) RMS-CreateMachine 8) Configuration-Activate 17) RMS-RemoveMachine 9) Configuration-Copy Choose an action: 2

## A.2 Modifying Cluster Applications

Follow the steps shown below to add a new file system to a cluster application. Refer to the PRIMECLUSTER Installation and Administration Guide for information on other modifications on cluster applications. Stop the RMS on all nodes when modifying cluster applications.



## A.2.1 Moving to the Cluster Application Edit Window

1. Execute the hvw command.

# /opt/SMAW/SMAWRrms/bin/hvw

2. Select [Application-Edit].

pcl-vml3: Main configuration menu, current configuration: config

No RMS active in the cluster

1) HELP
10) Configuration-Remove
2) QUIT
11) Configuration-Freeze
3) Application-Create
12) Configuration-Thaw
4) Application-Edit
13) Configuration-Edit-Global-Settings
5) Application-Remove
14) Configuration-Consistency-Report
6) Application-Clone
15) Configuration-ScriptExecution

```
7) Configuration-Generate 16) RMS-CreateMachine
8) Configuration-Activate 17) RMS-RemoveMachine
9) Configuration-Copy
Choose an action: 4
```

#### 3. Select the user application to be used.

```
Edit: Application selection menu (restricted):

1) HELP

2) QUIT

3) RETURN

4) OPTIONS

5) APP1

Application Name: 5
```

## A.2.2 Adding a Created Shared Disk Class to the Cluster Application

### 1. Select [Gds:Global-Disk-Services].

```
Settings of turnkey wizard "STANDBY" (APP1:consistent)
1) HELP
                                       10) Enterprise-Postgres(Fsep_APP1)
2) READONLY
                                        11) Symfoware(-)
3) SAVE+EXIT
                                        12) Procedure:SystemState3(-)
4) -
                                       13) Procedure:SystemState2(-)
5) ApplicationName=APP1
                                       14) Gls:Global-Link-Services(Gls_APP1)
6) Machines+Basics(app1)
                                       15) IpAddresses(-)
7) CommandLines(-)
                                       16) LocalFileSystems(Lfs_APP1)
                                       17) Gds:Global-Disk-Services(Gds APP1)
8) Procedure:Application(-)
9) Procedure:BasicApplication(-)
Choose the setting to process:17
```

### 2. Select [AdditionalDiskClass].

```
Volume management (Gds_APP1:consistent)

1) HELP

6) DiskClasses[0]=class0001

2) NO-SAVE+EXIT

7) (ClassNameFilter=)

3) SAVE+EXIT

8) (StandbySupport=no)

4) REMOVE+EXIT

9) (AutoRecover=no)

5) AdditionalDiskClass

10) (Timeout=1800)

Choose the setting to process: 5
```

### 3. Select the created class.

```
1) HELP
2) RETURN
3) FREECHOICE
4) class0001
5) class0002
Choose the disk class: 5
```

### 4. Select [SAVE+EXIT].

```
Volume management (Gds_APP1:consistent)

1) HELP
7) DiskClasses[1]=class0002
2) NO-SAVE+EXIT
8) (ClassNameFilter=)
3) SAVE+EXIT
9) (StandbySupport=no)
4) REMOVE+EXIT
10) (AutoRecover=no)
5) AdditionalDiskClass
11) (Timeout=1800)
6) DiskClasses[0]=class0001
Choose the setting to process: 3
```

5. Select [SAVE+EXIT] to quit or [LocalFileSystems] to add local file systems.

```
Settings of turnkey wizard "STANDBY" (APP1:consistent)
1) HELP
                                       10) Enterprise-Postgres(Fsep_APP1)
2) -
                                       11) Symfoware(-)
3) SAVE+EXIT
                                       12) Procedure:SystemState3(-)
4) -
                                       13) Procedure:SystemState2(-)
5) ApplicationName=APP1
                                       14) Gls:Global-Link-Services(Gls_APP1)
6) Machines+Basics(app1)
                                       15) IpAddresses(-)
7) CommandLines(-)
                                       16) LocalFileSystems(Lfs_APP1)
8) Procedure:Application(-)
                                       17) Gds:Global-Disk-Services(Gds_APP1)
9) Procedure:BasicApplication(-)
Choose the setting to process:
```

## A.2.3 Adding a Created File System to the Cluster Application

1. Select [LocalFileSystems].

```
Settings of turnkey wizard "STANDBY" (APP1:consistent)
1) HELP
                                        10) Enterprise-Postgres (Fsep APP1)
2) -
                                        11) Symfoware(-)
3) SAVE+EXIT
                                        12) Procedure:SystemState3(-)
4) –
                                        13) Procedure:SystemState2(-)
5) ApplicationName=APP1
                                        14) Gls:Global-Link-Services(Gls_APP1)
6) Machines+Basics(appl)
                                        15) IpAddresses(-)
7) CommandLines(-)
                                       16) LocalFileSystems(Lfs_APP1)
8) Procedure:Application(-)
                                       17) Gds:Global-Disk-Services(Gds_APP1)
9) Procedure:BasicApplication(-)
Choose the setting to process:16
```

2. Select [AdditionalMountPoint].

```
File systems (Lfs_APP1:consistent)

1) HELP

7) MountPoints[1]=LA:/mnt/swdsk2

2) NO-SAVE+EXIT

8) (Filter=)

3) SAVE+EXIT

9) (Timeout=180)

4) REMOVE+EXIT

5) AdditionalMountPoint

6) MountPoints[0]=LA:/mnt/swdsk1

Choose the setting to process: 5
```

3. After registering the file systems to be used, select [SAVE+EXIT].

```
File systems (Lfs_APP1:consistent)

1) HELP

8) MountPoints[2]=LA:/mnt/swdsk3

2) NO-SAVE+EXIT

9) (Filter=)

3) SAVE+EXIT

10) (Timeout=180)

4) REMOVE+EXIT

5) AdditionalMountPoint

6) MountPoints[0]= LA:/mnt/swdsk1

7) MountPoints[1]=LA:/mnt/swdsk2

Choose the setting to process:
```

## A.2.4 Reflecting Modifications Made to the Cluster Application

1. Select [SAVE+EXIT].

```
Settings of turnkey wizard "STANDBY" (APP1:consistent)

1) HELP

10) Enterprise-Postgres(Fsep_APP1)

2) -

11) Symfoware(-)

3) SAVE+EXIT

12) Procedure:SystemState3(-)

4) -

13) Procedure:SystemState2(-)
```

```
5) ApplicationName=APP1 14) Gls:Global-Link-Services(Gls_APP1)
6) Machines+Basics(app1) 15) IpAddresses(-)
7) CommandLines(-) 16) LocalFileSystems(Lfs_APP1)
8) Procedure:Application(-) 17) Gds:Global-Disk-Services(Gds_APP1)
9) Procedure:BasicApplication(-)
Choose the setting to process:3
```

### 2. Select [RETURN].

```
Edit: Application selection menu (restricted):

1) HELP

2) QUIT

3) RETURN

4) OPTIONS

5) APP1

Application Name: 3
```

### 3. Select [Configuration-Generate].

```
pcl-vml3: Main configuration menu, current configuration: config
No RMS active in the cluster
1) HELP
                                       10) Configuration-Remove
2) QUIT
                                       11) Configuration-Freeze
3) Application-Create
                                       12) Configuration-Thaw
4) Application-Edit
                                      13) Configuration-Edit-Global-Settings
5) Application-Remove
                                      14) Configuration-Consistency-Report
                                      15) Configuration-ScriptExecution
6) Application-Clone
                                      16) RMS-CreateMachine
7) Configuration-Generate
                                      17) RMS-RemoveMachine
8) Configuration-Activate
9) Configuration-Copy
Choose an action: 7
```

### 4. Select [Configuration-Activate].

If processing is successful, "The activation has finished successfully." will be displayed. If this message is not displayed, there is an issue with the information or configuration that was set. In such a case, review the changed information.

```
pcl-vm13: Main configuration menu, current configuration: config
No RMS active in the cluster
1) HELP
                                       10) Configuration-Remove
2) QUIT
                                      11) Configuration-Freeze
3) Application-Create
                                      12) Configuration-Thaw
4) Application-Edit
                                      13) Configuration-Edit-Global-Settings
 5) Application-Remove
                                     14) Configuration-Consistency-Report
 6) Application-Clone
                                      15) Configuration-ScriptExecution
                                      16) RMS-CreateMachine
 7) Configuration-Generate
                                      17) RMS-RemoveMachine
 8) Configuration-Activate
 9) Configuration-Copy
Choose an action: 8
```

### 5. Select [QUIT].

pcl-vm13: Main configuration menu,	current configuration: config
No RMS active in the cluster	
1) HELP	10) Configuration-Remove
2) QUIT	11) Configuration-Freeze
3) Application-Create	12) Configuration-Thaw
4) Application-Edit	13) Configuration-Edit-Global-Settings
5) Application-Remove	14) Configuration-Consistency-Report
6) Application-Clone	15) Configuration-ScriptExecution
7) Configuration-Generate	16) RMS-CreateMachine
8) Configuration-Activate	17) RMS-RemoveMachine
9) Configuration-Copy	
Choose an action: 2	

## Appendix B Command Reference

This appendix explains for details in command.

## B.1 pgx\_pclrsc

#### Name

pgx\_pclrsc -- Register, unregister, or display FUJITSU Enterprise Postgres database cluster with PRIMECLUSTER as resource of PRIMECLUSTER.

#### **Synopsis**

#### Description

pgx\_pclrsc is an utility for registering or unregistering FUJITSU Enterprise Postgres database cluster with PRIMECLUSTER, or displaying settings of FUJITSU Enterprise Postgres database cluster registered with PRIMECLUSTER. Execute the command on a member node of the cluster. Only the super user(root) can execute this command.

#### **Options**

-a

--add

Register the specified database cluster with PRIMECLUSTER. If the specified database cluster has already been registered with PRIMECLUSTER, then terminate abnormally.

-c name

--db-cluster-name=name

Specify database cluster name. name is case-sensitive, and must be within 16 bytes, and an initial letter must be ascii alphabet, and following letters must be ascii alphabet or ascii digit or underscore(\_).

-d

--delete

Unregister the specified database cluster from PRIMECLUSTER.

-D directory

--pgdata=directory

Specify an absolute path of data directory.

--db-user=name

Specify a database super user. Default is the user specified with '-u'.

-n nodes

--member-nodes=nodes

Specify names of all member nodes of the cluster. Specify a cluster node name with uffix "RMS" to the name. Separate names with comma(,). (e.g. -n node1RMS,node2RMS)

-p --print

Print a list of database clusters registered with PRIMECLUSTER. If -c option is specified, then print settings of the database cluster.

#### --response-timeout=seconds

Specify timeout of the query for health check. It's used with a count specified with '--timeout-retry-num'. The query is "SELECT 1" to the database "template1". If '0', wait infinitely. (default: '0')

#### --timeout-retry-count=count

Specify a limit count of retrying query when timeout is occurred. If retry counter is over the limit, then PRIMECLUSTER considers status of the database server as FAULT. If a query doesn't timeout once, retry counter is reset. (default: '6')

#### --trace-max-file-size=size

Specify max size(KB) of trace file. (default: '10240')

#### -u osuser

#### --os-user=osuser

Specify an OS user who can start/stop FUJITSU Enterprise Postgres database server.

#### -w directory

#### --work-dir=directory

Specify a directory for temporary data and trace files. It's used for starting, stopping or checking FUJITSU Enterprise Postgres database server. If the directory doesn't exist, then create it. Owner of the directory created by the command is set to the user specified with '-u'. Permission of the directory created by the command is set to 0700.

#### --watch-interval=seconds

Specify an interval. It's used for the amount of time between health checks. (default: '3')

#### **Diagnostics**

0: On success

otherwise: On error

#### **Notes**

Before unregistering a resource, stop RMS of PRIMECLUSTER.

#### **Example**

The simplest example of registering a resource.

# pgx\_pclrsc -a -c inst1 -u postgres -D /mnt/swdsk1/pgdata -w /var/tmp/work -n node1RMS,node2RMS

## Index

[A]
Adding a Created File System to the Cluster Application38
Adding a Created Shared Disk Class to the Cluster Application37
Adding Tablespaces16
[0]
[C]
Checking Operation
Command Reference
Configuring PRIMECLUSTER6 Configuring Storage Data Protection Using Transparent Data
Encryption
Creating a Cluster Application and Configuring Its Attributes 24
Creating a File System
Creating a GDS Volume
Creating an Operating System User to Start FUJITSU Enterprise
Postgres
Creating a Symfoware Database Cluster
Creating Fsystem Resources
Creating FUJITSU Enterprise Postgres Resources
Creating Gds Resources
Creating Gls or Takeover Network Resources
Creating Resources and Cluster Applications
Creating Resources and Creating/Modifying Cluster
Applications24
Creating Resources and Finalizing Cluster Application Creation
[D]
Definition of Failover Operation1
Deleting a database cluster resource17
Displaying database cluster resource information16
(E)
[E]
Editing Configuration Files
Errors during state transition
[F]
Failover Operation
Feature of failover operation
[1]
Identifying the Cause of an Error22
Installing PRIMECLUSTER and FUJITSU Enterprise Postgres6
(M)
[M]
Maintenance Tasks
Modifying Cluster Applications
Modifying database cluster resource content
Mounting the File System
Mutual Switch Maintenance
vidual Switch Maintenance
[N]
Non-transferrable feature

[O]	
Operation at State Transition	17
Operation mode of failover operation	2
Operations under Normal Circumstances	17
Overview of Failover Operation	
[P]	
Preparing for Creation	
Procedures Required after a Failover Error	22
ID1	
[R]	22
Recovery Tasks	
Reflecting Modifications Made to the Cluster Application	
Registering Resource Information for the FUJITSU Enterpo	rise
Postgres Database Cluster	9
Restarting the Cluster Applications Stopped for Node	
Maintenance	23
[S]	
Setting Up Failover Operation	
Simultaneous Stopped Node Maintenance	
Starting the Database Cluster	23
Stopping the Cluster Applications for Node Maintenance	23
Stopping the Database Cluster	23
System Configuration of Failover Operation	2
_	
[U]	
Unmounting the File System	11



# FUJITSU Enterprise Postgres 12



# Connection Manager User Guide

## **Preface**

#### Purpose of this document

This document describes the Connection Manager features of FUJITSU Enterprise Postgres.

#### Intended readers

This document is aimed at people who use the Connection Manager features.

Readers of this document are also assumed to have general knowledge of:

- FUJITSU Enterprise Postgres
- PostgreSQL
- Linux

#### Structure of this document

This document is structured as follows:

#### **Chapter 1 Connection Manager Features**

Explains the features and Mechanisms of the Connection Manager.

#### Chapter 2 Setting Up

Explains setting up the Connection Manager.

#### Chapter 3 Using from an Application

Explains how to use the Connection Manager from an application.

#### Appendix A System Views

Explains the system view of Connection Manager.

#### **Export restrictions**

If this document is to be exported or provided overseas, confirm legal requirements for the Foreign Exchange and Foreign Trade Act as well as other laws and regulations, including U.S. Export Administration Regulations, and follow the required procedures.

#### Issue date and version

```
Edition 1.1: June 2025
Edition 1.0: July 2020
```

#### Copyright

Copyright 2020-2025 FUJITSU LIMITED

## **Revision History**

Revision	Location	Version
Change article about specifying connection destinations in libpq.	3.3.1 Specifying a Connection Destination	Edition 1.1

## Contents

Chapter 1 Connection Manager Features	1
1.1 Heartbeat Monitoring Feature	
1.1.1 Difference from TCP keepalive	1
1.1.2 Mechanism of Heartbeat Monitoring Feature	2
1.2 Transparent Connection Support Feature	
1.2.1 Mechanism of Connections using Transparent Connection Support Feature	
Objection C. Cattley Uni	
Chapter 2 Setting Up	
2.1 Setting Up the Client Side	
2.1.1 Creating a Directory for the conmgr Process	
2.1.2 Configuring conmgr.conf	
2.2 Setting Up the Server Side	
2.2.1 Configuring postgresql.conf	
2.2.2 Introducing the watchdog extension	9
2.3 Removing Setup	9
Chapter 3 Using from an Application	10
	• • • • • • • • • • • • • • • • • • • •
3.1 Connection Method.	10
3.1 Connection Method	
3.1 Connection Method	10
3.2 How to Detect Instance Errors	10
3.2 How to Detect Instance Errors 3.3 How to Use in libpq	10 10
3.2 How to Detect Instance Errors 3.3 How to Use in libpq	
3.2 How to Detect Instance Errors 3.3 How to Use in libpq 3.3.1 Specifying a Connection Destination 3.3.2 Using the Asynchronous Connection Method 3.3.3 Using an Asynchronous Communication Method	
3.2 How to Detect Instance Errors 3.3 How to Use in libpq 3.3.1 Specifying a Connection Destination 3.3.2 Using the Asynchronous Connection Method 3.3.3 Using an Asynchronous Communication Method 3.3.4 Behavior of PQhost() or PQhostaddr() or PQport()	
3.2 How to Detect Instance Errors 3.3 How to Use in libpq 3.3.1 Specifying a Connection Destination 3.3.2 Using the Asynchronous Connection Method 3.3.3 Using an Asynchronous Communication Method	
3.2 How to Detect Instance Errors.  3.3 How to Use in libpq	
3.2 How to Detect Instance Errors.  3.3 How to Use in libpq	
3.2 How to Detect Instance Errors.  3.3 How to Use in libpq	

## Chapter 1 Connection Manager Features

The Connection Manager provides the following features:

#### Heartbeat monitoring feature

Detects kernel panics between the server running the client and the server running the PostgreSQL instance(hereinafter referred to as instance), physical server failures, and inter-server network link downs, and notifies the client or instance. The client is notified as an error event through the SQL connection, and the instance will be notified in the form of a force collection of SQL connections with clients that are out of service.

#### Transparent connection support feature

When an application wants to connect to an instance of an attribute in a set of instances configured for replication, it can connect to that instance without being aware of which server it is running on.



- The available client drivers for Connection Manager are libpq (C language library) and ECPG (embedded SQL in C).
- The Connection Manager is only available for Linux server and Linux client.

Each function is described below.

## 1.1 Heartbeat Monitoring Feature

Describes the Connection Manager's heartbeat monitoring feature.



The Connection Manager does not monitor for delays, such as CPU busy occurring in the postmaster process or in the backend processes to which the application connects directly, or for no response, such as due to a software bugs. It also does not monitor application downtime or unresponsiveness. To detect these, use various timeout features provided by PostgreSQL or the client drivers.

## 1.1.1 Difference from TCP keepalive

A peer of TCP connections cannot automatically detect a link down or server down.

There are two main methods to detect it. One is the operating system (Not all operating systems support it) TCP keepalive feature, and the other is the keepalive-equivalent timeout function implemented at the application layer. Connection Manager's heartbeat monitoring capabilities are categorized as the latter.

The operating system TCP keepalive feature has the following disadvantages, but the Connection Manager's heartbeat monitoring feature does not:

- The keepalive does not work when the TCP layer cannot receive an acknowledgement (ACK) and retransmits the packet repeatedly. This means that it is not possible to detect a down (For example, if a network goes down,) before sending some data and receiving ACK from the other side. There is also a parameter to interrupt retransmissions, which is not supported by some operating systems. The Connection Manager's heartbeat monitoring feature does not have this disadvantage because it is timeout monitoring at the application layer.
- The periodic packets for keepalive are sent per-TCP socket. If a instance accepts too many (For example, a few thousand clients) SQL connections, the load on the instance side cannot be ignored. The Connection Manager's heartbeat monitoring feature greatly reduces the load by allowing packets to be sent to the instance on a per-server basis on which the client runs.

## 1.1.2 Mechanism of Heartbeat Monitoring Feature

On the client side, the user must start one monitoring process using the cm\_ctl command for the set of the instances to be monitored. This process, called the "conmgr process", can only be started by a user who is not an administrator (e.g. superuser(root) on Linux). An instance set is a collection of one or more instances that make up replication. One configuration file (conmgr.conf) for each conmgr process is used to set the information about the set of the instances being monitored and the parameters for monitoring.

On the server side, by installing PostgreSQL's EXTENSION that is called "watchdog", the postmaster will start two processes as background workers at instance startup.

One is the process for sending and receiving packets to and from the conmgr process for heartbeat monitoring. It is called "watchdog process". The other is the process for forcibly terminating SQL connections of the clients for which the watchdog process detects a failure on heartbeat monitoring. It is called "terminator process". SQL connections that do not use Connection Manager is also terminated, because the terminator process terminates them by IP address as key.

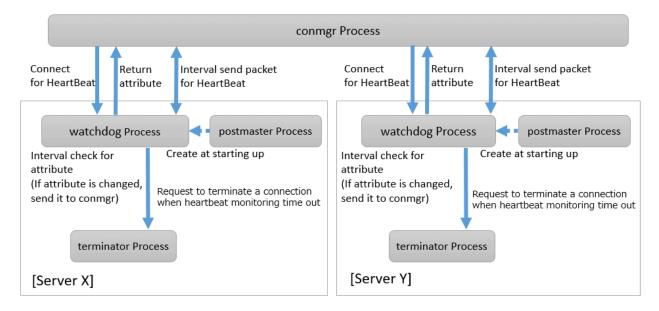


#### **System Configuration Notes**

For replication, it is recommended that the instance that connects to the upstream instance of replication and the conmgr process that regards the upstream instance as an instance to be monitored for heartbeart (specified in backend\_host parameter or backend\_hostaddr parameter that is a configuration parameter of conmgr process) be not placed on the same server. This is because if the conmgr process stops normally or abnormally, the terminator process in the upstream instance will also kill the replication connection. The replication connection will reconnect automatically even if it is forcibly disconnected, so replication will continue without any problems. However, this can be a problem when the replication load is high or on systems that are sensitive to replication delays.

Note that the replication connection have different monitoring feature than the Connection Manager, so there is no need to monitor the Connection Manager for heartbeat. Refer to PostgreSQL documentation for details.

The process relationship is as follows:





Refer to "cm ctl" in the Reference for information on cm ctl command.

## 1.2 Transparent Connection Support Feature

The features similar to Connection Manager's transparent connection support feature can be found in PostgreSQL's libpq and other client drivers.

Using libpq as an example, the connection parameter to use that feature is target\_session\_attrs parameter. If this parameter is used not through Connection Manager, libpq will attempt to find the required instance by connecting sequentially to all instances of the set of instance requested by the host parameter or hostaddr parameter. In the worst case, libpq may find the promoted primary at the connection to the last instance of instance set. This means that you cannot predict how long it will take to complete the switch.

However, when commbined with the Connection Manager, the conmgr process obtains its attributes via the watchdog process from all servers in a set of servers in advance, so that the connections to that server can be initiated as soon as the application requests it.

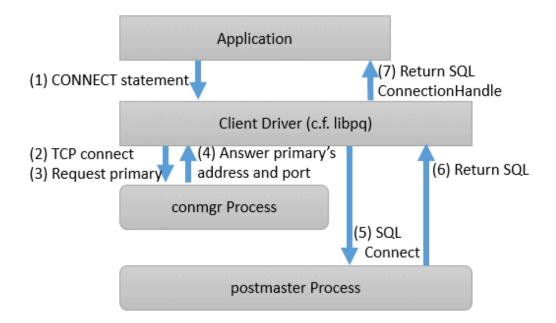
## 1.2.1 Mechanism of Connections using Transparent Connection Support Feature

A connection using this mechanism actually consists of two steps, but from the perspective of the application, it looks like a single SQL connection. In the application's connection string, specify the IP address or host name (In most cases it is "localhost") and port number where the connect process listens, and target\_session\_attrs parameter. You do not need to explicitly state that the connection is to the connect process. This is because the client driver can automatically determine whether the connection is to a instance or a connect process.

In the first phase of the connection, the client driver receives a connection request from the application and connects to the location specified in the connection string. Initially, it uses the protocol PostgreSQL requests, and if it learns in the middle that the connection is to a conmgr process, it asks the conmgr process for the IP address and port number that the instance with the attributes specified in the connection parameter target\_session\_attrs is listening for. If the destination is a backend process rather than a conmgr process, the connection process completes immediately and continues to send and receive data for normal SQL execution. The first stage of processing falls within the scope of timeout monitoring for SQL connection processing by each client driver. For example, the connection\_timeout parameter of libpq.

In the second phase of the connection, the client driver connects to the instance using the IP address and port number from the conmgr process. Thereafter, the client driver and the instance directly send and receive the data for SQL execution. This ensures that the Connection Manager does not affect the performance of the SQL execution.

When the client driver is waiting to receive data after the second stage is completed, it monitors the reception of data to the two sockets obtained at each stage of the connection. This allows the client driver to know when, for example, the conmgr process notifies the client of a network link down.



## Chapter 2 Setting Up

Describes setting up the Connection Manager.

## 2.1 Setting Up the Client Side

On the client side, configure settings for the conmgr process.

## 2.1.1 Creating a Directory for the conmgr Process

You need one conmgr process for each set of instances that you want to configure for replication. Assign a dedicated directory to each conmgr process. This directory must assign read, execute, and write permissions for the user who starts the conmgr process.

This directory is specified when you run the cm\_ctl command, which starts and stops the conmgr process. To specify a directory in the cm\_ctl command, set it in the environment variable CMDATA or specify it in the -D option.

## 2.1.2 Configuring conmgr.conf

Place the configuration file conmgr.conf in the directory for the conmgr process.

#### Syntax for conmgr.conf

- In conmgr.conf, after the symbol(#) are considered comments.
- The parameter name = value" is a set of settings and must be written on one line.
- Set the value in a format that matches the type of each parameter. The types and formats are:
  - integer: Numeric type. Express as a sequence of numbers in decimal number.
  - string: String type. You can also include spaces by enclosing them in quotation marks('). If you include quotation marks, escape them.
  - enum: Enumeration type. Possible values are determined.

#### **Parameters to Set**

port (integer)

Specify the port number on which the conmgr process listens for connections from the applications,

The value must be greater than or equal to 1 and less than or equal to 65535. The default is 27546. You must restart conmgr process for this parameter change to take effect.

backend\_host\* (string)

Specify the host name or IP address of the instance.

You can also use IPv6 address. If you specify the IP address directly, you can save time by using backend\_hostaddr parameter. If backend\_host parameter and backend\_hostaddr parameter are both specified, backend\_hostaddr parameter is used. You must restart conmgr process for this parameter change to take effect.

To distinguish multiple instances, append a zero-based number immediately after the parameter name, such as backend\_host0, backend\_host1,... This number is called the instance number. A parameter identified by the same instance number configures the settings of a single instance. If you want to exclude some instances from your replication configuration, you can simply remove the settings for that instance.



Refer to "System Configuration Notes" in "1.1.2 Mechanism of Heartbeat Monitoring Feature" for details.

For example, if two instances are listening on "host name:host0, port number:5432" and "host name:host1, port number:2345", write as follows.

```
backend_host0='host0'
```

backend\_port0=5432

backend\_host1='host1'

backend\_port1=2345

You can also mix different instance number settings:

backend\_host0='host0'

backend\_host1='host1'

backend\_port0=5432

backend\_port1=2345

It does not matter if the instance number is missing as in the following (instance number 1):

backend\_host0='host0'

backend\_host2='host2'

backend\_port0=5432

backend\_port2=2345

If the host name is omitted even if the port is specified, as in the following instance number 1, it is regarded as a missisng number.

backend\_host0='host0'

backend\_host2='host2'

backend\_port0=5432

backend\_port1=5555

backend\_port2=2345

#### backend\_hostaddr\*(string)

Same as backend\_host parameter except no name resolution is used.

#### backend\_port\* (integer)

Specify the port number the postmaster of the instance will listen on.

The value must be greater than or equal to 1 and less than or equal to 65535. The default is 27500. Append the instance number as you would for backend\_host parameter. You must restart conneg process for this parameter change to take effect.

#### watchdog\_port\* (integer)

Specify the port number on which the watchdog process listens.

The conmgr process connects to this port, but the user application does not. you must set it to the same value as watchdog.port parameter in postgresql.conf. The value must be greater than or equal to 1 and less than or equal to 65535. The default is 27545. Append the instance number as you would for backend\_host parameter. You must restart conmgr process for this parameter change to take effect.

#### heartbeat\_interval (integer)

Specify the interval at which heartbeat packets are sent for heartbeat monitoring.

Used in conjunction with heartbeat\_timeout parameter. Connection Manager heartbeat monitoring always continues to send packets periodically from both ends of the connection. If a packet is not received from the other side within a certain period of time, the link is considered down.

Note that this method is different from TCP keepalive. TCP keepalive send a keepalive packet only when there is a certain amount of inactivity (idle), and expects to receive an ACK for that packet. If TCP keepalive does not receive an ACK, it repeats this a specified number of times and then assumes that the link is down.

The heartbeat\_interval parameter and heartbeat\_timeout parameter are propagated from the conmgr process to the watchdog process, and also apply to the interval between the transmissions of heartbeat packets from the watchdog process. If a watchdog process is connected from both a conmgr process with a heartbeat\_interval parameter of 3 seconds and a conmgr process with a heartbeat\_interval of 5 seconds, it sends heartbeat packets every 3 seconds to the former process and every 5 seconds to the latter process. The unit is seconds. Specify a value equal to or more than 1 second. The default is 10 seconds. You must restart conmgr process for this parameter change to take effect.

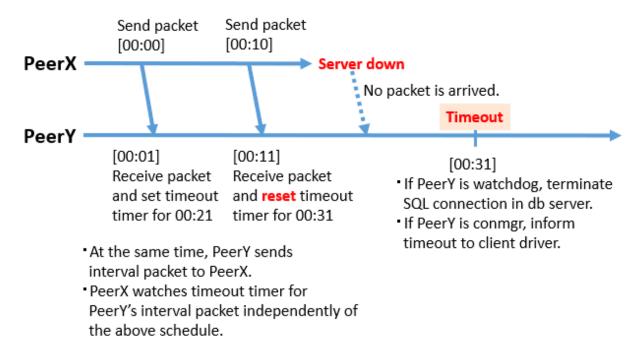
#### heartbeat\_timeout (integer)

If a heartbeat packet for heartbeat monitoring cannot be received for more than the time specified by this parameter, an error is assumed to have occurred and the application is notified of the error.

This parameter should be decide of heartbeat\_interval parameter as the basis. No error is occured when the configuration file is loaded, but is always considered abnormal by heartbeat monitoring if it is at least not greater than heartbeat\_interval parameter. The unit is seconds. Specify a value equal to or more than 1 second. The default is 20 seconds. You must restart conmgr process for this parameter change to take effect.

Refer to the following figure for the relationship between the heartbeat\_interval parameter and heartbeat\_timeout parameter settings and the heartbeat timeout.

#### [Configuration: heartbeat\_interval=10, heartbeat\_timeout=20]



#### heartbeat\_connect\_interval (integer)

Specify the interval between attempts to establish heartbeat monitoring again after detecting an abnormality.

This parameter is useful when only the database server is started, but not the instance. In such a situation, the TCP connection fails immediately, and retries cannot be attempted without an interval. If you specify an excessively long value, you may delay noticing the start of the instance. If a connection attempt fails for a long time, it will attempt the next connection after the time specified by heartbeat\_connect\_interval parameter has elapsed. The unit is seconds. Specify a value equal to or more than 1 second. The default is 1 second. You must restart connegr for this parameter change to take effect.

#### heartbeat\_connect\_timeout (integer)

Specify the connection timeout for establishing heartbeat monitoring.

The connection includes the time it takes to send the TCP connection and the first heartbeat packet to the watchdog process and receive a reply from the watchdog process. This parameter is particularly needed when the other server is down or the network is disconnected. This is because TCP connections are attempted over a long period of time, depending on the operating system configuration, and the connection takes a long time to fail. The unit is seconds. Specify a value equal to or more than 1 second. The default is 10 seconds. You must restart connegr process for this parameter change to take effect.

#### log\_destination (string)

Specify the destination of the message.

You can specify multiple destinations. Use commas to separate multiple entries. "stderr" and "syslog" can be specified. The default is to print only to stderr. You must restart conmgr process for this parameter change to take effect.

#### syslog\_facility (enum)

Specify the syslog facility.

Valid only if log\_destination parameter includes "syslog". LOCAL0, LOCAL1, LOCAL2, LOCAL3, LOCAL4, LOCAL5, LOCAL6, or LOCAL7 can be specified.

The default is "LOCALO". You must restart conmgr process for this parameter change to take effect.

#### syslog\_ident (string)

Specify the program name used to identify the output from the conmgr process.

The default is "conmgr". You must restart conmgr process for this parameter change to take effect.

#### log\_min\_messages (enum)

Specifies the level of messages to output.

It can be DEBUG, INFO, NOTICE, WARNING, ERROR, LOG, FATAL, or PANIC. Messages below the specified level are not output. The default is "WARNING". You must restart conmgr process for this parameter change to take effect.

#### max\_connections (integer)

Specifies the maximum number of simultaneous connections to the conmgr process.

If there are more than this maximum number of client connections, it forces the connection to be closed without sending an error message to the client.

The conmgr process also outputs this fact at level "LOG" to the destination specified by log\_destination. Specify a value equal to or more than 0.

If 0 is specified, there is no limit. The default is 0. You must restart conmgr process for this parameter change to take effect.



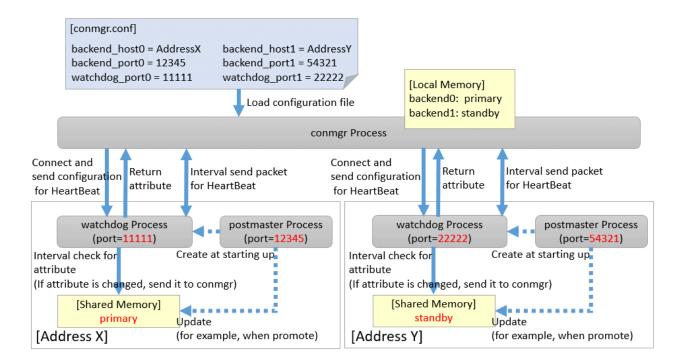
The maximum number of file descriptors that can be opened simultaneously (You can check it with -n of the ulimit command.) imposed on a conmgr process by the OS user limit should be greater than the value derived from the following equation:. Otherwise, the conmgr process will abort if the user limit is violated.

......

9 + Number of database instances specified in conmgr.conf x 2 + max\_connections specified in conmgr.conf

#### Overview of connections definitions

The following figure shows the relationship between the IP address or host name and the port number set in conmgr.conf and the processes.



## 2.2 Setting Up the Server Side

On the server side, configure settings for the watchdog process.

## 2.2.1 Configuring postgresql.conf

Describes the postgresql.conf parameters that must be set when using the Connection Manager.

#### Parameters to Set

#### max\_connections

An existing PostgreSQL parameter. Add 2 to the value already set.

Connection to the instance is maintained from the time the instance is started to do the following:

- The watchdog process checks the state of the instance.
- The terminator process forces the client to terminate the SQL connection.

#### shared\_preload\_libraries

An existing PostgreSQL parameter. Add a watchdog.

The watchdog process and terminator process start when you add watchdog and restart the instance.

#### watchdog.port (integer)

Specify the port number on which the watchdog process accepts connections for heartbeat monitoring from the conmgr process.

The value must be greater than or equal to 1 and less than or equal to 65535. The default is 27545. The instance must be restarted for this parameter change to take effect.

#### watchdog.check\_status\_interval (integer)

Specify the interval between checking the attributes of a instance.

watchdog process immediately notifies the conmgr process if the attribute changes.

The unit is milliseconds. Specify a value equal to or more than 1 millisecond. The default is 1000 milliseconds. The instance must be restarted for this parameter change to take effect.

watchdog.max\_hb\_connections (integer)

Specify the maximum number of conmgr processes that connect to watchdog process.

The default is the value specified in max\_connections of postgresql.conf.

There is no upper limit, but about 200 bytes of memory are consumed for 1 connection when PostgreSQL is started.



Normally you do not need to consider, but if you have a heartbeat connection with a very large number of conmgr processes, it may violate on the maximum number of file descriptors (You can check it with -n of the ulimit command.) of the OS user limit. This is because the socket for the heartbeat connection consumes the file descriptor. Set the maximum number of file descriptors of the OS user limit to a value larger than the value calculated below from the max\_files\_per\_process parameter value and watchdog.max\_hb\_connections parameter value in postgresql.conf.

max\_files\_per\_process + watchdog.max\_hb\_connections x 2

## 2.2.2 Introducing the watchdog extension

Execute the CREATE EXTENSION statement with watchdog.

Example)

postgres=# CREATE EXTENSION watchdog;
CREATE EXTENSION

This allows you to see the pgx\_stat\_watchdog view for information about the watchdog process.

## 2.3 Removing Setup

No work is required on the client side.

On the server side, drop watchdog extension by DROP EXTENSION statement and remove it from shared\_preload\_libraries.

Example)

postgres=# DROP EXTENSION watchdog; DROP EXTENSION

## Chapter 3 Using from an Application

Describes how to use the Connection Manager from an application.

#### 3.1 Connection Method

When connecting to the instance using ConnectionManager, specify the following values in the connection parameters of the application. Application connection parameters are parameters that specify the database IP address, host name, port number, etc., which are originally specified when connecting to the database from the application. For example, when using libpq, specify "localhost" for the host parameter and specify the port number on which the connegr process listens for the port parameter.

Connection parameters not shown here are used directly by the instance in the second stage of the connection, connecting to the instance (connecting to an instance without the Connection Manager), and the conmgr process does not check or use it.

#### Connection destination address

Specify "localhost". Unix domain sockets are not allowed.

It is possible to connect to a remote conmgr process, but it should not be used for other purposes expect such as testing. This is because there is no mechanism between the application and the conmgr process to detect the remote server down or the network link down, making the Connection Manager meaningless.

#### Port number

Specify the value specified for the port parameter in conmgr.conf.

#### Connection destination instance attributes

Follow the "Target server" in the application connection switch feature. Refer to "Taget server" in "Connection Information for the Application Connection Switch Feature" in the "Application Development Guide" for information on the target server in the application connection switch feature.

### 3.2 How to Detect Instance Errors

Only special if you are using libpq's asynchronous communication method. For additional discovery methods, refer to "Errors when an Application Connection Switch Occurs and Corresponding Actions" of the for each client driver in the "Application Development Guide". If the conmgr process goes down while accessing it, or if the conmgr process tries to establish a SQL connection while it is down, the same error is returned as if the instance went down.

## 3.3 How to Use in libpq

libpq provides very detailed communication control. Therefore, to detect a heartbeat error through the conmgr process, you may need to modify the existing application logic.



See

Refer to "libpq - C Library" in the PostgreSQL Documentation on functions described below.

## 3.3.1 Specifying a Connection Destination

The host or hostaddr parameter of the connection string should specify a connection destination for only one conmgr process.



The host or hostaddr parameter of the connection string can be a mixture of multiple conmgr process destinations and the postmaster destination. The first connection from an application process is attempted in the listed order, and the second and subsequent connections

are attempted in the order in which the previous connection was made. However, if a connection to the conmgr process is confirmed, subsequent connections are always attempted in the listed order.

For example, if conmgr1, conmgr2 are specified, and if conmgr1 does not initially know a server with the attributes specified in the target\_session\_attrs parameter, it will query conmgr2 for a destination. In the second and subsequent connections, since there was a connection to the conmgr1 process, the connection destinations are queried in the order of the listed conmgr1 and conmgr2.

Also, for example, if postmaster1, conmgr1 is specified, it will try to connect directly to the database instance initially named postmaster1. If this fails, it queries conmgr1 for a destination. On subsequent connections, since there was a connection to the conmgr1 process, it queries the listed postmaster1, then conmgr1, for connections.

## 3.3.2 Using the Asynchronous Connection Method

An asynchronous connection method is to use a function like PQconnectStart() instead of a function like PQconnectdb (). PQconnectStart() returns without synchronizing the completion of the connection to the database. The user application must then monitor the sockets returned by PQsocket() to be readable or writable, for example by using the poll() system call, according to the values required by the return value of PQconnectPoll().

With the Connection Manager, the socket returned by PQsocket() may change after a call to PQconnectPoll(), so be sure to reacquire the socket that you want to give to the poll() system call using PQsocket(). This behavior is similar to simply specifying multiple hosts in the connection string without using the Connection Manager.

## 3.3.3 Using an Asynchronous Communication Method

An asynchronous communication method is one in which the application returns control without waiting for a response from the database, and PQsetnonblocking() is used to asynchronize completion of transmission or completion of receipt of all results. Instead of using a function like PQexec(), use a function like PQsendQuery(). In this method, the user application monitors the socket that connects to the database returned by PQsocket(), for example, by using the poll() system call.

For example, if the link to the database goes down, simply monitoring the socket returned by PQsocket with the poll() system call will not detect it.

However, it is possible to detect the reception of database anomaly detection packets sent from the conmgr process, for example, by monitoring the reception of data on the socket (POLLIN) connecting to the conmgr process returned by PQcmSocket(). Once a reception is detected, the user application need not directly manipulate the packet. By calling something like PQgetResult() or PQcosumeInput() according to the existing application logic, it behaves as if the connection were disconnected. Refer to "Errors when an Application Connection Switch Occurs and Corresponding Actions" in the Application Development Guide on SQLSTATE returned, etc. If you are not using the Connection Manager, PQcmSocket() returns -1.

## 3.3.4 Behavior of PQhost() or PQhostaddr() or PQport()

PQhost(), PQhostaddr() or PQport() typically return a host parameter or hostaddr parameter or port parameter specified in the connection string by the user application. However, if you specify a connection destination for the conmgr process, the destination for the conmgr process you specify will be changed to the database connection destination information provided by the conmgr process before the connection is completed. This behavior is similar to simply specifying multiple hosts in the connection string without using the Connection Manager.

## 3.3.5 Behavior of PQstatus()

If you are using an asynchronous connection method, you can monitor the intermediate state of the connection to the database with PQstatus(). If you are using the Connection Manager, the enum value returned by PQstatus() is appended with the following:

```
CONNECTION_AWAITING_CMRESPONSE
/ * Waiting for a response from the conmgr process * /
```

## 3.3.6 PQcmSocket()

You can get a socket that leads to the conmgr process. It returns a value equal to or more than 0 for a valid socket, or -1 if you are not connected to the conmgr process.

int PQcmSocket(const PGconn \*conn);

## Appendix A System Views

## A.1 pgx\_stat\_watchdog

A row in this view corresponds to conmgr process, which is connected to watchdog process. Additional columns may be added in future versions versions.

Column	Туре	Description
conmgr_addr	inet	IP address of conmgr process.
conmgr_port	integer	The conmgr (ephemeral) port number that conmgr process is using to communicate with watchdog process. This is not the port number to be set in conmgr.conf.
heartbeat_interval	integer	The interval at which heartbeat packets are sent to and from this conmgr process. The unit is seconds.
heartbeat_timeout	integer	The timeout value for the heartbeat to and from this conmgr process. The unit is seconds.

## Index

	[B]
backend_host*	4
backend_hostaddr*	5
backend_port*	5
	[C]
	4
conmgr process	2
	(I I)
II 4	[H]
	1
	6
	6
	5
heartbeat_timeout	5
	[L]
log destination	6
	7
log_mm_messages	/
	[M]
max connections (integer)	7
	8
	[P]
pgx_stat_watchdog	13
	4
postgresql.conf	8
PQcmSocket()	12
	[S]
_	8
	6
syslog_ident	7
	(T)
	[T]
	2
Transparent connection support	ort feature1
	[W]
watchdog.check status interv	/al8
	ns9
	8
	2
	5



# FUJITSU Enterprise Postgres 12



## Reference

Glossary >

Reference Guide >

Message Guide >

FJQSS User Guide >



# FUJITSU Enterprise Postgres 12



## Reference Guide

## Preface

#### Purpose of this document

This document is a command reference, and explains FUJITSU Enterprise Postgres commands and options with features expanded on from PostgreSQL.

#### Intended readers

This document is aimed at people who manage and operate FUJITSU Enterprise Postgres. Readers of this document are also assumed to have general knowledge of:

- PostgreSQL
- SQL
- Linux
- W
- Windows

#### Structure of this document

This document is structured as follows:

#### Chapter 1 Command List and Specification Format

Lists commands and describes their specification format.

#### Chapter 2 Client Commands

Explains options not listed in "PostgreSQL Client Applications" in the PostgreSQL Documentation.

#### Chapter 3 Server Commands

Explains commands and options not listed in "PostgreSQL Server Applications" in the PostgreSQL Documentation.

#### Chapter 4 Mirroring Controller Commands

Explains the Mirroring Controller commands.

#### Chapter 5 Connection Manager Commands

Explains the Connection Manager commands.

#### How to read this document

Examples in this document are predominantly for UNIX/Linux.

For Windows, replace values (such as paths in the examples) as appropriate.

#### **Export restrictions**

If this document is to be exported or provided overseas, confirm legal requirements for the Foreign Exchange and Foreign Trade Act as well as other laws and regulations, including U.S. Export Administration Regulations, and follow the required procedures.

#### Issue date and version

Edition 1.0: July 2020

#### Copyright

Copyright 2015-2020 FUJITSU LIMITED

## Contents

Chapter 1 Command List and Specification Format	1
1.1 Command List	1
1.1.1 Client Commands	1
1.1.2 Server Commands	1
1.1.3 Mirroring Controller Commands	1
1.1.4 Connection Manager Commands	
1.2 Command Specification Format	
Chapter 2 Client Commands	
2.1 pg_dumpall	
2.2 pgx_loader	3
Chapter 3 Server Commands	6
3.1 pg_ctl	
3.2 pgx_dmpall	6
3.3 pgx_keystore	8
3.4 pgx_rcvall	
3.5 postgres	
Chapter 4 Mirroring Controller Commands	13
4.1 mc_arb	
4.2 mc_ctl	
Chapter 5 Connection Manager Commands	20
5.1 cm_ctl	

## Chapter 1 Command List and Specification Format

This chapter lists commands and describes their specification format.

#### 1.1 Command List

This chapter lists commands and options not explained in "PostgreSQL Client Applications" or in "PostgreSQL Server Applications" in the PostgreSQL Documentation.

#### 1.1.1 Client Commands

The commands below have options not explained in "PostgreSQL Client Applications" in the PostgreSQL Documentation.

Command	Functional overview
pg_dumpall	Extract a PostgreSQL database cluster into a script file

The command below is not explained in "Client Applications" in the PostgreSQL Documentation.

Command	Functional overview
pgx_loader	Loads data from an external file into a PostgreSQL table.

### 1.1.2 Server Commands

The commands below have options not explained in "PostgreSQL Server Applications" in the PostgreSQL Documentation.

Command	Functional overview
pg_ctl	Initialize, start, stop, or control a PostgreSQL server
postgres	PostgreSQL database server

The commands below are not explained in "PostgreSQL Server Applications" in the PostgreSQL Documentation.

Command	Functional overview
pgx_dmpall	Backs up the data directory, tablespaces, and configuration files.
pgx_keystore	Manages keystore
pgx_rcvall	Recovers the data directory, tablespaces, and configuration files.

## **1.1.3 Mirroring Controller Commands**

Mirroring Controller has the following commands:

Command	Functional overview
mc_arb	L Start, stop and display the status of the Mirroring Controller arbitration process.
	W Start, stop and display the status of the Mirroring Controller arbitration process, or register and unregister as a Windows service.

Command	Functional overview
mc_ctl	L Start and stop Mirroring Controller, switch/disconnect the server, or display the server status.
	W Start and stop Mirroring Controller, switch/disconnect the server, display the server status, or register and unregister as a Windows service.

## 1.1.4 Connection Manager Commands

Connection Manager has the following commands:

Command	Functional overview
cm_ctl	Start, stop or display the status of the conmgr process

## 1.2 Command Specification Format

The table below shows the command specification format.

Item	Explanation
[]	Indicates optional element.
	Indicates that the item can be specified repeatedly.

## **Chapter 2 Client Commands**

This chapter explains options not listed in "PostgreSQL Client Applications" in the PostgreSQL Documentation.

## 2.1 pg\_dumpall

#### Name

pg\_dumpall -- Extract a PostgreSQL database cluster into a script file

#### **Synopsis**

```
pg_dumpall [connectionOption...] [option...]
```

#### **Options**

--no-tablespace-encryption

Do not output commands to encrypt tablespaces. Running the generated SQL script will restore the originally encrypted data without being encrypted.

#### See

Refer to "pg\_dumpall" in the PostgreSQL Documentation for details.

## 2.2 pgx\_loader

#### Name

pgx\_loader --Loads data from a file into a PostgreSQL table.

#### Overview

```
pgx_loader load -c command [options...]
pgx_loader recovery -t table
```

#### **Description**

The pgx\_loader command loads data from an external file into PostgreSQL tables, and commits or rolls back transactions prepared during data load.

In load mode, data is loaded at high speed by executing the COPY FROM command specified in *command* at a certain degree of parallelism. If load is completed successfully, the message below is output to the standard output.

```
Note: count is the number of rows loaded.
```

In recovery mode, commit or rollback of transactions prepared during data load is performed.

#### **Options**

-a

--echo-sql

Display the executed command in the standard output.

- -c command
- --copy-command=command

Specify the COPY FROM command to be executed. If STDIN is specified for the FROM clause, data will be loaded from the standard input. In this case, SQL superuser privileges (or having one of the roles pg\_read\_server\_files or

pg\_execute\_server\_program) are not required, because local user access privileges will be used for external files and external programs, instead of database server access privileges.

"binary" cannot be specified for the FORMAT option of the COPY FROM command specified in this option. Therefore, input files in binary format cannot be specified.

The FREEZE option cannot be specified for the COPY FROM command specified in this option.



#### See

Refer to "COPY" in the PostgreSQL Documentation for information on the COPY FROM command.

#### -i number-of-jobs

#### --jobs=number-of-jobs

Specify the number of background workers (parallel workers) that the COPY COMMAND should use to simultaneously perform data conversion, table creation, and index creation. This option can dramatically reduce the time for loading a large amount of data on instances that runs on multiple processor machines.

The optimal value depends on the server, client, and network configurations. The number of CPU cores and disk configuration also affect the optimal value. The number of CPU cores of the server is recommended as the initial value to try. Naturally, if a value that is too large is used, performance degradation will occur due to thrashing and context switching.

Specify a value from 2 to 128. The default is 2.

#### -t table

#### --table=table

Complete the prepared transactions only for the specified table.

-?

#### --help

Show how to use pgx\_loader command line arguments, and exit.

The command line options below control the database connection parameters.

#### -d connstr

#### --dbname=connstr

Specify the database name to connect to.

If this option is not specified, the PGDATABASE environment variable will be used. If the environment variable is not set, your operating-system user name will be used.

#### -h host

#### --host=host

Specify the host name of the machine the database server runs on. If the specified value starts with a slash, it will be used as the directory for a Unix domain socket.

If this option is not specified, the PGHOST environment variable will be used. If the environment variable is not set, it will be considered a Unix domain socket connection.

#### -p port

#### --port=port

Specify the TCP port to be used by the server to monitor the connection, or extension of the local Unix domain socket file.

If this option is not specified, the PGPORT environment variable will be use. If the environment variables is not set, 27500 will be used.

#### -U username

#### --username=username

User name for connection to the database.

-w

#### --no-password

Never prompt for the password. If the server requires password authentication but other means (such as a .pgpass file) are not available, the connection attempt will fail. This option can be useful in batch jobs, scripts, and so on, where no user is present to enter a password.

-W

#### --password

Force pgx\_loader to prompt for the password before connecting to the database. This option is never essential, since pgx\_loader will automatically prompt for the password if the server demands password authentication. However, pgx\_loader will waste a connection attempt finding out if the server requires a password. In some cases it is worth specifying this option to avoid the extra connection attempt.

#### **Diagnostics**

load mode

0: Normal exit

Other than 0: Abnormal exit

#### recovery mode

- 0: There are no prepared transactions that must be completed
- 3: A prepared transaction was committed
- 4: A prepared transaction was rolled back

Other than the above: Abnormal exit



The order of the table rows loaded by pgx\_loader may not match the order of the file lines. This is because the file lines will have been inserted into the table in parallel, by multiple parallel workers.

#### **Example**

The example below loads the file /path/to/data.csv (2000 records) into table tbl using a degree of parallelism of 3.

```
\ pgx\_loader\ load\ -j\ 3\ -c\ "COPY\ tbl\ FROM\ '/path/to/data.csv'\ WITH\ CSV''\ LOAD\ 2000
```

The example below reads the file /path/to/data.csv (2000 records) from the standard input and loads into table tbl using a degree of parallelism of 3.

```
\ pgx\_loader\ load\ -j\ 3\ -c\ "COPY\ tbl\ FROM\ STDIN\ WITH\ CSV"\ <\ /path/to/data.csv\ LOAD\ 2000
```

The example below completes the transactions prepared for table tbl.

```
$ pgx_loader recovery -t tbl
```

## Chapter 3 Server Commands

This chapter explains commands and options not listed in "PostgreSQL Server Applications" in the PostgreSQL Documentation.

## 3.1 pg\_ctl

#### Name

pg\_ctl -- Initialize, start, stop, or control a PostgreSQL server

#### **Synopsis**

#### **Options**

--keystore-passphrase

Prompt for the passphrase to open the keystore.

#### See

Refer to "pg\_ctl" in the PostgreSQL Documentation for details.

## 3.2 pgx\_dmpall

#### Name

pgx\_dmpall - Backs up the data directory, tablespaces, and configuration files.

#### **Synopsis**

```
pgx_dmpall [option...]
```

#### **Description**

The pgx\_dmpall command backs up the data directory, tablespaces, and configuration files. The backup data is stored in the directory specified by backup\_destination parameter of postgresql.conf. The pgx\_dmpall command also deletes archived Write Ahead Logs (WAL) that are no longer necessary for recovery when the backup completes successfully.

#### **Options**

-c

This option only backs up configuration files. The configuration files are as follows:

- postgresql.conf (postgresql.conf)
- File for host-based authentication (pg\_hba.conf)
- Configuration file for ident authentication (pg\_ident.conf)

If an external reference, such as 'include' in postgresql.conf, is set, the reference destination files are also backed up.

#### -C fast|spread

#### --checkpoint=fast|spread

Sets checkpoint mode to fast or spread (default).

If fast is specified, the checkpoint processing at the start of backup becomes quick, but the impact on performance of running applications gets larger due to intense I/O. In spread mode, the impact on applications is smaller but the backup takes longer, because the checkpoint is performed slowly.

#### -D datadir

Specify the data directory. If this option is omitted, the value of the environment variable PGDATA is used.

#### -f configFile

Specify the postgresql.conf configuration file. This option is set if the data directory and the configuration file set in the 'data\_directory' parameter of the postgresql.conf file are running in separate directories.

#### -P tablespacesListFile

#### --tablespaces-list-file=tablespacesListFile

Specify the full path of the file containing the names of the tablespaces to be backed up using the copy command, using less than 1024 bytes.

The file format is described below:

```
tablespaceName<newline>
tablespaceName<newline>
```

Tablespaces not listed in the specified file are backed up to the backup storage directory. If this option is not specified, all tablespaces are backed up using the copy command.

This option can be specified if the -Y option has been specified, and it is used to limit the tablespaces backed up using the copy command.

#### -U username

#### --username=username

Specify the user name of the database superuser. This defaults to the name of the effective user running pgx\_dmpall.

#### -Y copyCommandFile

#### --copy-command=copyCommandFile

Specify the full path of the file of the copy command for backup, using less than 1024 bytes. This option cannot be specified together with the -c option.

#### -w

#### --no-password

Never issue a password prompt. If the server requires password authentication and a password is not available by other means such as a .pgpass file, the connection attempt will fail. This option can be useful in batch jobs and scripts where no user is present to enter a password.

#### -W

#### --password

Force pgx\_dmpall to prompt for a password before connecting to a database.

This option is never essential, since pgx\_dmpall will automatically prompt for a password if the server demands password authentication. However, pgx\_dmpall will waste a connection attempt finding out that the server wants a password. In some cases it is worth typing -W to avoid the extra connection attempt.

#### --maintenance-db=dbname

Specifies the name of the database to connect to. If not specified, the postgres database will be used; if that does not exist, template1 will be used.

Any database can be specified as long as it can be connected to.

#### --exclude-copy-cluster

Excludes a database cluster from backup via the copy command. This option can be specified if the -Y option has been specified. If this option is not specified, the database cluster will be backed up using the copy command.

#### **Environment**

#### **PGDATA**

Specify the data directory. You can overwrite using the -D option.

#### **Diagnostics**

0: Normal end

Other than 0: Abnormal end

#### **Notes**

This command can only be executed when the database server is running.

Execute this command as a PostgreSQL user account.

Do not update or delete files in the backup storage directory. Otherwise, you may not be able to recover the database.

Do not store other files in the backup storage directory.

This command uses one database connection. To establish a connection, this command uses the IPv4 loopback address 127.0.0.1 on Windows, and the UNIX domain socket on other operating systems. Therefore, permit these connections in pg\_hba.conf.

This command cannot be executed on the standby server.

#### **Example**

In the following example, the data directory, tablespaces, and configuration files are backed up. At this time, stored WALs are no longer necessary because the backups are destroyed.

```
$ pgx_dmpall
```

#### Related item

pgx\_rcvall

## 3.3 pgx\_keystore

#### Name

```
pgx_keystore -- Manages keystore
```

#### **Synopsis**

```
pgx_keystore [option...] keystore_location
```

#### **Description**

pgx\_keystore enables auto-open of a keystore.

#### **Options**

-a

--enable-auto-open

Enables auto-open of a keystore. This allows the keystore to open automatically without entering the passphrase when the database server starts.

When auto-open is enabled, an obfuscated copy keystore.aks is created in the same directory where the keystore file keystore.ks is stored. To disable auto-open, delete keystore.aks.

#### -P passphrase

--passphrase=passphrase

Specify the passphrase to open the keystore. If this option is omitted, the prompt to enter the passphrase is displayed.

keystore\_location

Specify the absolute or relative path of the keystore file.

#### **Diagnostics**

0: Normal exit

Other than 0: Abnormal exit

#### **Notes**

This command can be executed whether the database server is running or stopped. It does not connect to the database server.

#### **Example**

Enables auto-open of a keystore.

```
$ pgx_keystore -a /key/store/location/keystore.ks
```

## 3.4 pgx\_rcvall

#### Name

pgx\_rcvall - Recovers the data directory, tablespaces, and configuration files.

#### **Synopsis**

```
pgx_rcvall [option...]
```

#### **Description**

The pgx\_rcvall command recovers the data directory, tablespaces, and configuration files using the data that was backed up with pgx\_dmpall command and archived Write-Ahead-Log (WAL). If none of the options that indicate the recovery point is specified, all archived WAL are applied and the data will be recovered to the latest point.

#### **Options**

-B backupdir

Specify the backup storage directory. If the data directory is damaged, this option cannot be omitted.

-D datadir

Specify the data directory. If this option is omitted, the value of the environment variable PGDATA is used.

-e targetTime

Specify this option to recover the data as of the specified date and time.

targetTime

Specify the time at which the data is recovered. The format is as follows:

```
"YYYY-MM-DD HH:MM:SS"
```

-1

This option displays a list of the backup data information in the backup storage directory that was obtained using the pgx\_dmpall command. If the pgx\_dmpall command was executed using the copy command (-Y option) for backup, the resources backed up using the copy command will also be listed. This cannot be specified together with -p, -e or -n option.

#### -n restorePoint

Specify this option to recover the data to the specified restore point. Restore points are created with SQL function pg\_create\_restore\_point. If multiple restore points with the same names were created, the first one after the backup was taken is used for recovery. If the specified restore point does not exist, all archived WAL are applied. This cannot be specified together with -e or -p option.

-p

Specify this option to recover the data as of the time when the last backup completed. This cannot be specified together with -e or -n option.

#### -s connectionString

Construct a standby server from the backup created by pgx\_dmpall. Specify as an argument a connection string for connecting to the primary server. This is the same as primary\_conninfo parameter in recovery.conf. This option can only be specified together with -D and -B.

-X

Specify this option if you do not want to include transactions committed at the time specified in the -e option as part of the recovery.

#### -Y copyCommandFile

#### --copy-command=copyCommandFile

Specify the full path of the file of the copy command for recovery, using less than 1024 bytes. This option cannot be specified together with the -l option.

#### --keystore-passphrase

Prompt for the passphrase to open the keystore.

#### --view-results-of-copying

Output the backup information file that was written by the copy command executed via the pgx\_dmpall command. This option cannot be specified together with the -l, -p, -e, -n, or -Y option.

#### **Environment**

#### **PGDATA**

Specify the data directory. You can overwrite using the -D option.

#### **PGPORT**

Specify the port number for connecting to the database.

#### **PGUSER**

Specify the user name of the database superuser. This defaults to the name of the effective user running pgx\_dmpall.

#### **Diagnosis**

0: Normal exit

Other than 0: Abnormal exit

#### **Backup data information**

#### Date

Date the backup data was created using the pgx\_dmpall command.

#### Dir

This is the name of the directory in the backup storage directory that is used to store the backup data.

Directory naming format: Time format (YYYY-MM-DD\_HH-MM-SS)

#### Status

This is the status of the pgx\_dmpall command backup data.

COMPLETE: Complete INCOMPLETE: Incomplete

#### Resources backed up by the copy command

List of resources that were backed up by the copy command executed via the pgx\_dmpall command.

If there are resources that were backed up by the copy command, then database clusters ('pg\_data') or tablespace names will be listed, delimited by header and halfwidth comma.

#### **Notes**

This command can only be executed when the database server is stopped, except when it is executed with -l option.

Execute this command as a PostgreSQL user account.

Use backup data that was taken from the recovery target data directory.

Before executing this command, disconnect all application database connections. Additionally, do not connect to the database during recovery.

The configuration files are restored from those files that were taken by the last pgx\_dmpall (including -c option).

This command connects to the database to determine whether the recovery has completed. So ensure that you set the port number with PGPORT environment variable in the environment where multiple instances exist.

Match the OS timezone setting when running pgx\_dmpall/pgx\_rcvall to the timezone specified by timezone parameter in postgresql.conf.

Otherwise, data might be recovered to an unexpected time when -e or -p is specified.

If you recover to a past point, a new timeline (history of database updates) begins at that point. That recovery point is the latest point in the new timeline when the recovery is completed. If you subsequently recover to the latest point, the database updates in the new timeline will be replayed.

Valid restore points are the ones that were created in the timeline where the backup had been taken. That means that if you recover to a past point, those restore points created thereafter are unavailable. Therefore, take a backup when you have restored the past data desired.

If the pgx\_dmpall command was executed using the copy command (-Y option) for backup, it is necessary to execute this command using the copy command (-Y option) for recovery. However, because the list of resources (database cluster or tablespace) that were backed up using the copy command is recorded in the backup directory, there is no need to specify the target resources when executing this command. The -I option can be used to check the target resources for which a backup is retrieved using the copy command.

#### **Example**

In the following example, the data directory, tablespaces, and configuration files are recovered.

```
$ pgx_rcvall -B /home/pgsql/Backupdir
```

In the following example, the data directory and tablespaces are recovered at 10:00:00 on 01-05-2017. The configuration files are recovered at the point at which the last of the data is obtained.

```
$ pgx_rcvall -B /home/pgsql/Backupdir -e "2017-05-01 10:00:00"
```

In the following example, the data directory and tablespaces are recovered up to the time of restore point "before\_match\_20170510\_1". The configuration files are restored from the latest backup.

```
$ pgx_rcvall -B /home/pgsql/Backupdir -n before_match_20170510_1
```

In the following example, the obtained backup data information in the backup storage directory is displayed in a list.

```
$ pgx_rcvall -1
```

#### Related item

pgx\_dmpall

## 3.5 postgres

#### Name

```
postgres -- PostgreSQL database server
```

#### **Synopsis**

```
postgres [option...]
```

#### **Options**

-K

Prompt for the passphrase to open the keystore. This option works in single-user mode only, so you must also specify the --single option, as shown below:

```
postgres --single -K
```

#### See

Refer to "postgres" in the PostgreSQL Documentation for details.

# Chapter 4 Mirroring Controller Commands

This chapter explains the Mirroring Controller commands.

### 4.1 mc\_arb

#### Name



mc\_arb - Start, stop, and display the status of the Mirroring Controller arbitration process



mc\_arb - Start, stop, and display the status of the Mirroring Controller arbitration process, or register and unregister it as a Windows service

#### Overview

```
mc_arb start [-M mcdir] [-w| -W]
mc_arb stop [-M mcdir] [-e]
mc_arb status [-M mcdir]
```



```
mc_arb register [-M mcdir] [-w| -W] -P password [-S a[uto] | d[emand]] mc_arb unregister [-M mcdir]
```

#### Description



mc\_arb starts, stops, and displays the status of the Mirroring Controller arbitration process.

W

mc\_arb starts, stops, and displays the status of the Mirroring Controller arbitration process, or registers and unregisters it as a Windows service.

The start mode starts the Mirroring Controller arbitration process.

The stop mode stops the Mirroring Controller arbitration process.

The status mode displays the connection status of the Mirroring Controller arbitration process with the Mirroring Controller processes of the primary server and standby server.

W

The register mode registers the Mirroring Controller arbitration process as a Windows service. The -w and -W options can be used to start/stop the Mirroring Controller arbitration process through the Windows service, and when it executes the mc\_arb command, the specified option will take effect.

The unregister mode unregisters the Mirroring Controller arbitration process as a Windows service.

If the Mirroring Controller arbitration process has not been started on the server executing the command, stop mode and status mode will terminate with an error.

Additionally, if Mirroring Controller is forcibly stopped on the database server, it may take a few moments until the status mode displays the status of the server connection as offline.



This command can be executed by any user.



Execute this command as a user with administrator privileges (operating system user that belongs to the "Administrators" group).

#### **Options**

-е

Specify this option to forcibly stop the Mirroring Controller arbitration process on the active server.

Specify this option to stop the Mirroring Controller arbitration process but keep Mirroring Controller running (to stop both, first stop Mirroring Controller of the database server, and then the Mirroring Controller arbitration process). This option can also be specified to halt the arbitration process (by stopping the Mirroring Controller arbitration process) when the fencing command called by the arbitration process becomes unresponsive.

#### -M mcdir

Specify the Mirroring Controller arbitration process management directory.

ASCII characters can be specified in the directory path.

If this option is omitted, the value of the environment variable ARBCONTROLDIR is used.

-P password

For the register mode, specify the password for the user who executed the command.

-S a[uto] | d[emand]

Specify the start type for the Windows service to be registered. You can choose auto or demand as the start type by specifying the entire word or just its first letter. The default is auto.

-W

Waits for operations to finish.

This option is the default of start mode.

This option is the default of start mode and register mode.

-W

Does not wait for operations to finish.

#### **Environment**

#### **ARBCONTROLDIR**

Specify the Mirroring Controller arbitration process management directory.

ASCII characters can be specified in the directory path.

You can specify the -M option to override this value.

#### **Diagnostics**

0: Normal end

Other than 0: Abnormal end

#### **Notes**

w

If the Mirroring Controller arbitration process is forcibly stopped or communication between the command and the Mirroring Controller arbitration process is interrupted while the Mirroring Controller arbitration process is being stopped, a message that the command is being executed may be output and stopping may terminate in error, even though no other instances of the mc\_arb command are being executed.

To solve this issue, ensure that other instances of the mc\_arb command are not being executed before forcibly stopping the Mirroring Controller arbitration process.

Use the start mode and stop mode to start and stop the Windows service. Before being able to use these modes, use the register mode to register the Mirroring Controller arbitration process as a Windows service.

The mc\_arb command must be executed from "Administrator: Command Prompt". Right-click [Command Prompt], and then select [Run as administrator] from the context menu to display the [Administrator: Command Prompt] window.

#### **Example**

Start the Mirroring Controller arbitration process.

\$ mc\_arb start -M /mcarb\_dir/arbiter1

W

Start the Mirroring Controller arbitration process.

```
> mc_arb start -M D:\mcarb_dir\arbiter1
```

Display details of mc\_arb status

```
        server_id
        host
        status

        (1)
        (2)
        (3)
```

- (1) Server identifier
- (2) Host name or IP address
- (3) Server connection status online : Connected offline : Disconnected

### 4.2 mc\_ctl

#### Name



mc\_ctl - Start and stop Mirroring Controller, switch/disconnect the server, or display the server status.



mc\_ctl - Start and stop Mirroring Controller, switch/disconnect the server, display the server status, or register and unregister as a Windows service.

#### Overview

```
mc_ctl start [-M mcdir] [-w| -W] [-f| -F] [--mc-only] [--async-connect-arbiter] [--local-server server_id]
mc_ctl stop [-M mcdir] [[-a] [--mc-only]| [-e][--local-server server_id]]
mc_ctl status [-M mcdir] [--arbiter] [--local-server server_id]
mc_ctl switch [-M mcdir] [--force [--no-fencing]] [--local-server server_id]
mc_ctl detach [-M mcdir] [--no-fencing] [--local-server server_id]
mc_ctl enable-failover [-M mcdir] [--local-server server_id]
mc_ctl disable-failover [-M mcdir] [--local-server server_id]
```



#### **Description**



mc\_ctl starts and stops Mirroring Controller, switches/disconnects the server, or displays the server status.



mc\_ctl starts and stops Mirroring Controller, switches/disconnects the server, displays the server status, or registers and unregisters as a Windows service.

The start mode starts Mirroring Controller. If the --mc-only option is omitted, the command starts a database instance.

The stop mode stops Mirroring Controller. If the --mc-only option is omitted, the database instance is stopped. If --mc-only option is not specified, database instance is also stopped. When executes on standby server without --mc-only, standby server will be detached from primary server.

The status mode displays the status of the servers, database instance processes, and disks monitored by Mirroring Controller. Additionally, if the --arbiter option is specified, Mirroring Controller arbitration process connection status is displayed.

The switch mode switches the primary server. When the server is switched, the database instance on the primary server stops, and the database instance on the standby server is upgraded to primary server and begins degrading operation. This mode can be executed on the primary server and standby server in an environment where the Mirroring Controller process can communicate with the Mirroring Controller process on the other server.

The detach mode forcibly disconnects the standby server. This mode is used to forcibly disconnect the other server when stopping of Mirroring Controller cannot be performed using stop mode (which requires login to the standby server). It can only be executed on the primary server.

The enable-failover mode enables automatic switching and disconnection.

The disable-failover mode disables automatic switching and disconnection.

The register mode registers Mirroring Controller as a Windows service. The -w, -W, -f, -F, --mc-only, --async-connectarbiter, and --local-server options are used when Mirroring Controller is started and stopped as a Windows service. If the mc\_ctl command is used to start and stop Mirroring Controller, the option specified in the command will be valid.

The unregister mode unregisters Mirroring Controller as a Windows service.

If Mirroring Controller has not been started on the server that executes the command, commands for any mode other than the start mode, and status mode.

Execute this command as an instance administrator.

If Mirroring Controller has not been started on the server that executes the command, commands for any mode other than the start mode, status mode, register mode, and unregister mode terminate with an error.

Execute this command as a user with the "Administrator" privilege (operating system user ID that belongs to the Administrator group).

Until you start Mirroring Controller of standby server after starting Mirroring Controller of the primary server, operation can be started with only the primary server. Standby server is incorporated when you start the Mirroring Controller of standby server, and you should be able to operate in the multiplexing configuration.

#### **Options**

-a

w

L

w

Specify this option to stop Mirroring Controller on all servers.

-е

Specify this option to forcibly stop Mirroring Controller on the active server.

-f

Specify this option to enable automatic switching and disconnection of Mirroring Controller immediately after startup.

This option is the default of start mode.

This option is the default of start mode and register mode.

-F

Specify this option to disable automatic switching and disconnection immediately after startup of Mirroring Controller.

#### --async-connect-arbiter

Specify this option to connect the Mirroring Controller start process asynchronously to the Mirroring Controller arbitration process. This option can be specified to forcibly start Mirroring Controller if the Mirroring Controller arbitration process is not started.

Specify this option if using the Mirroring Controller arbitration server.

#### --arbiter

Specify this option to display the connection status of the Mirroring Controller arbitration process. This option can be specified if using status mode.

Specify this option if using the Mirroring Controller arbitration server.

#### --local-server server\_id

If you run a simulation build of the primary and standby servers in a single server (for system testing, for example), specify this option to identify the server to be operated.

For *server\_id*, specify the server identifier specified in the network.conf file. ASCII characters other than single-byte space can be specified in the server identifier. The operations will be executed as if the user has logged in to *server\_id*.

#### --mc-only

Specify this option to start and stop only Mirroring Controller processes. At the start mode, this option can be specified only while the database instance is running. If this option is omitted, the database instance is simultaneously started and stopped.

#### -M mcdir

Specify the Mirroring Controller management directory.

ASCII characters can be specified in the directory path.

If this option is omitted, the value of the environment variable MCCONTROLDIR is used.

#### --force

Switching with this option specified can only be specified on the standby server. This option is used to perform switching forcibly after performing fencing on the primary server if communication with the Mirroring Controller process of the other server is not possible (due to a network issue between database servers or unresponsive server, for example), thus preventing normal switching. However, if the --no-fencing option is specified, fencing will not be performed on the primary server.

#### --no-fencing

When switching or disconnection is executed with the --force option specified, fencing of the server to be degraded is circumvented. Therefore, it is necessary for the user to isolate the server to be degraded from the cluster system in advance.

#### -P password

W

For the register mode, specify the password for the user who executed the command.

#### -S a[uto] | d[emand]

Specify the start type for the Windows service to be registered. You can choose auto or demand as the start type by specifying the entire word or just its first letter. The default is auto.

-w

Waits for operations to finish.

This option is the default of start mode.

This option is the default of start mode and register mode.

#### -W

Does not wait for operations to finish.

#### **Environment variable**

#### **MCCONTROLDIR**

Specifies the Mirroring Controller management directory.

ASCII characters can be specified in the directory path.

You can specify the -M option to override this value.

#### **Diagnostics**

0: Normal end

Other: Abnormal end

#### **Notes**

The message under execution might be output though the mc\_ctl command is not being executed and, besides, it terminate abnormally when the server is downed while processing execution of this command, an automatic switch, and an automatic separation, and the communication between a primary server and the standby server is cut off. Besides, restart Mirroring Controller to solve this problem after confirming that the mc\_ctl command is not in progress. Afterwards, execute a necessary operation.

If a time-out error occurs when the mc\_ctl command is in progress, the messages may be different from the processes. Take the actions described in the "Action" section of the message.

Automatic switching and disconnection by the enable-failover mode, the disable-failover mode, and the start mode can be enabled/disabled only while Mirroring Controller is running. Therefore, if you want to enable/disable automatic switching and disconnection on starting, specify the -f option or -F option each time you start Mirroring Controller.

Use the start mode and stop mode to start and stop the Windows service. To do this, use the register mode to register Mirroring Controller as a Windows service in advance.

In case of postgresql.conf has any incorrect parameter when this command is executed, this command may be abnormally terminated. If this is the case, re-execute it again after correct the parameter in postgresql.conf.

If the arbitration server of Mirroring Controller is used, connection with the arbitration server will be performed on startup even if startup using start mode is executed with the -F option specified, or if executed with --local-server specified.

The mc\_ctl command must be executed from "Administrator: Command Prompt". Right-click [Command Prompt], and then select [Run as administrator] from the menu to display the [Administrator: Command Prompt] window.

For operation using an arbitration server, the amount of time spent attempting to connect with the arbitration server is calculated, so the Mirroring Controller startup time may take longer.

#### **Example**

To start Mirroring Controller:

```
$ mc_ctl start -M /mcdir/inst1
```

To start Mirroring Controller:

```
> mc_ctl start -M D:\mcdir\inst1
```

#### Display details of mc\_ctl status

```
mirroring status
-----(1)
server_id host_role host host_status db_proc_status disk_status
(2) (3) (4) (5) (6) (7)
```







(1) Multiplexing status

switched : Switched (displayed when switching has finished and the degrading

operations status has been enabled)

not-switchable : Not switchable (displayed when a server is disconnected and

switching is not possible)

unknown : Unknown (\*1)

failover-disabled : Failover is disabled

(2) Server identifier

(3) Server role

primary : Primary
standby : Standby

none(inactivated primary): No role

(primary is stopped or being defined as primary)

none(inactivated standby): No role

(standby is stopped or being defined as primary)

- (4) Host name or IP address
- (5) Live/dead state of the server
   normal : Normal operation

abnormal : Abnormal unknown : Unknown (\*1)

(6) DBMS process status

normal : Normal

abnormal (abnormal process name (\*2)) : Abnormal unknown : Unknown (\*1)

(7) Disk status

normal : Normal abnormal (abnormal disk type (\*3)) : Abnormal unknown : Unknown (\*1)

- \*1: Displayed when Mirroring Controller is stop state, the management network is abnormal, or Mirroring Controller has failed or is unresponsive.
- \*2: The names of the DBMS processes in which the abnormality was detected are output. The name has the following meaning: However, if multiple DBMS process issues are detected, only the DBMS for which the first issue was detected is displayed.
  - -postmaster: Process (postmaster) that accepts application connections
- -wal\_sender or wal\_receiver: Process (WAL sender or WAL receiver) that sends and receives transaction  $\log s$
- \*3: The types of disks where the abnormality was detected are output separated by a comma. The type has the following meaning:

-data: Data storage disk

-tran\_log: Transaction log storage disk

-tablespace: Tablespace storage disk

#### Display details of mc\_ctl status (with the --arbiter option specified)

arbiter\_id host status
(1) (2) (3)

- (1) Arbitration server identifier
- (2) Host name or IP address
- (3) Mirroring Controller arbitration process connection status

online : Connected

offline : Disconnected (\*1)

\*1: When Mirroring Controller is stopped, connections to the Mirroring Controller arbitration process cannot be established, so it will be displayed as "offline".

# Chapter 5 Connection Manager Commands

This chapter explains the Connection Manager commands.

### 5.1 cm ctl

#### Name

cm\_ctl - Start, stop or display the status of the conmgr process

#### **Synopsis**

```
cm_ctl start [-D diretory] [-W] [--complete] [-t seconds]

cm_ctl stop [-D diretory] [-W] [-m {smart | fast | immediate}] [-t seconds]

cm_ctl status [-D directory] [-t seconds] [-i {all | instance | application}]
```

#### **Description**

The start mode starts the commgr process. The command returns at least after the heartbeat monitoring connection is completed with the primary server's instance.

When --complete is specified, wait until all instances configured in conmgr.conf have completed their heartbeat monitoring connections. You can set a timeout for these waits. The default of timeout is 60 seconds. Can be changed using the -t option. If it times out, it simply gives up waiting and the conmgr process remains up.

The stop mode sends a signal to the conmgr process to shut down and wait until the process disappeares.

The default of wait time is 60 seconds. Can be changed using the -t option. If it times out, it simply gives up waiting. There are three shutdown methods, "smart", "fast", and "immediate", specify with the -m option. The "smart" waits until all applications using the conmgr process run out of SQL connections before shutting down. The "fast" forces all applications using the conmgr process to disconnect from the conmgr process before shutting it down. As a result, the SQL connection for the application receives an error. The "immediate" terminates the conmgr process immediately. If nothing is specified, it stops in smart mode.

The status mode, if the commgr process exists, queries the conmgr process for instance and application information known to the conmgr process, and display them to standard output along with the state of conmgr itself.

The -i option allows you to specify what information to query. The "instance" queries information about the instance; The "application" queries information about the application; The "all" queries information about both. conmgr's own information is always displayed. The default time to wait for a query to return is 60 seconds. Can be changed using the -t option.

#### **Options**

--complete

Wait until all instances configured in conmgr.conf have completed their heartbeat monitoring connections. If the When used with the -W option, the -W option takes precedence.

-D

-- directory=directory

Specify the directory where conmgr.conf is located. If omitted, it refers to the directory specified by the environment variable CMDATA. You cannot omit both.

-i {all | instance | application}

Specify the information to display the status.

-m

--mode={smart | fast | immediate}

Specify the mode of shutdown. The default is smart.

- -t seconds
- --timeout=seconds

Specify how long to wait for the operation to complete. The unit is seconds. The default is 60 seconds.

-W

--no-wait

In start mode, cm\_ctl command returns immediately after forking conmgr process. In stop mode, the cm\_ctl command returns without waiting for the process to disappeare.

#### **Diagnostics**

start mode or stop mode

- 0: Normal end
- 2: Timeout occurred
- 3: Unable to access the specified directory

Other than the above: None of the above

status mode

- 0: Normal end
- 2: Timeout occurred
- 3: Unable to access the specified directory
- 4: conmgr process does not exist

Other than the above: None of the above

#### **Privileges**

The conmgr process cannot be started by the administrator(e.g. superuser(root) on Linux).

#### **Example**

Display details of start mode

The block of information that can be specified with the -i option is used as a unit. There is one blank line between the blocks and no blank lines within the blocks. It includes one or more spaces between columns and between data.

```
$ cm_ctl status -i all
conmgr_status:
status pid
(1)
      (2)
ready 3456
instance_information:
addr
        port database_attr
(3)
            (4) (5)
192.0.2.100 27500 standby
192.0.2.110 27500 primary
192.0.2.120 27500 standby
192.0.2.130
            27500 unknown
application_infomation:
addr
        port pid connected_time
(6)
           (7) (8) (9)
127.0.0.1 12345 5678 2020-01-15 02:03:04
```

- (1) Status of the conmgr process
  - starting : The process is starting its startup sequence but is not ready to accept connections from clients.
  - ready : Ready to accept connections from clients.
  - stopping: It has received an end instruction and is starting the stop sequence.
  - inactive : The conmgr process does not exist.
- (3) IP address of the instance
- (4) Port number the postmaster listens on
- (5) Status of the instance (primary | standby | unknown)
- (6) Connection source IP address for conmgr process
- (7) Connection source (ephemeral) port number for conmgr process
- (8) PID of the connection source
- (9) Date and time conmgr process connection

The ISO 8601 compliant date is followed by a blank, followed by the ISO 8601 compliant second precision time.

This representation is a PostgreSQL string representation of type timestamp.



# FUJITSU Enterprise Postgres 12



Message Guide



# FUJITSU Enterprise Postgres 12



FJQSS User Guide (Information Collection Tool)

# Preface

This manual describes information to use the outline and the tool of FJQSS (Information Collection Tool).

#### Purpose of this Manual

The purpose of this manual is to be able to gather the collected information by using FJQSS (Information Collection Tool) when trouble occurs.

#### Reader of this Manual

This manual is targeted at the user who gathers the collected information when the trouble occurs.

You can obtain information on the method of gathering the function and the collected information of FJQSS (Information Collection Tool) obtain by referring to this manual.

#### Structure of this Manual

The structure of this manual is as follows:

Chapter 1 outline

Chapter 2 FJQSS (Information Collection Tool)

Appendix A Messages and error codes

#### **Notations**

The following notations are used in this manual;

Symbols used in this manual

The following symbols are used in this manual.

Example:[Enter]key

Symbols used in command

This subsection describes the symbols used in the examples of commands.

Example:

fjqss\_collect [-s]

Meaning of Symbols

Symbol	Meaning
[]	The item enclosed with this sign is shown and omissible is shown.

#### Abbreviations

The following abbreviations are used in this manual:

Operating Systems

Full Name	Abbreviations
Oracle Solaris 9	Solaris 9
Oracle Solaris 10	Solaris 10
Oracle Solaris 11	Solaris 11
Oracle Solaris	Solaris
Red Hat Enterprise Linux 6 (for x86)	RHEL6(x86)
Red Hat Enterprise Linux 6 (for Intel64)	RHEL6(Intel64)
Red Hat Enterprise Linux 5 (for x86)	RHEL5(x86)
Red Hat Enterprise Linux 5 (for Intel64)	RHEL5(Intel64)

#### **Export Management Regulations**

Our documentation may include special technology based on Foreign Exchange and Foreign Trade Control Law.

In case you export the relevant documentation or to provide any overseas resident with the relevant documentation, permission based on the above law is necessary.

#### Trademarks and Screenshots

Oracle and Java are registered trademarks of Oracle Corporation, its subsidiaries and its affiliated companies in the United States and other countries.

The names of companies and products, etc. in the sentences inside can be trademarks or registered trademarks of each company. Date of Publication

Date of Publication and Version	Manual Code
November 2012 Version1	J2S2-1625-01ENZ2(00)
April 2013 Version1.1.	J2S2-1625-02ENZ2(00)
October 2013 Version 1.2	J2S2-1625-03ENZ2(00)
July 2014 Version 1.3	J2UL-1958-01EN2Z(00)

#### Copyright Notice

Copyright 2014 FUJITSU LIMITED

# Contents

Chantan 4 Outline	4
Chapter 1 Outline	1 1
Chapter 2 Collecting the information	2
2.1 Collection method.	2
2.1.1 Collecting with showing menu	
2.1.2 Collecting without showing menu.	
2.1.3 Getting the product identification name	5
2.2 Checking the collected information.	<i>6</i>
2.3 Changing a path of the output directory	<i>6</i>
Appendix A Messages and error codes	ç
A.1 Error Messages	
A 2 Error Codes	 10
A Z ETTOT COORS	1 (

# Chapter 1 Outline

This chapter explains the outline of FJQSS (Information Collection Tool).

# 1.1 FJQSS (Information Collection Tool)

This appendix explains FJQSS (Information Collection Tool) that archives the information for investigation (middleware logs and operation system logs) with an easy operation. Executing FJQSS right after a problem occurs makes a rapid investigating of the causes.

# Chapter 2 Collecting the information

This chapter explains the method of gathering the survey material, the confirm method of the survey material, and notes.

### 2.1 Collection method

This section explains the operating procedure of FJQSS (Information Collection Tool).

Collecting information can do by using the command. You can operate the Collection by the method that shows the product selection menu, or the method that does not show such the menu.

- 2.1.1 Collecting with showing menu
- 2.1.2 Collecting without showing menu

### 2.1.1 Collecting with showing menu

This subsection explains the operating procedure of FJQSS (Information Collection Tool) with showing the product selection menu.

#### Authority to execute

User Authority of the root user is necessary.

#### **Prior confirmation**

Set the output directory of the collected information if necessary. The output directory of default is the /var/tmp directory on Solaris, or the /tmp directory on Linux. See [2.3 Changing a path of the output directory] for details of changing the output directory.

#### **Synopsis**

If you execute the command specifying the output directory by the parameter:

```
fjqss_collect output_directory
```

If you execute the command specifying the output directory by the environment variable or the environmental configuration file:

fjqss\_collect

#### Path to this command

opt/FJSVqstl

#### Notes in executing the command

- Do not put fjqss\_collect or its copy into any other directory. They work only if they are in the path shown above in [Path to this command].
- Execute this command immediately after the problem occurs.
- Multiple execution of this command is not available. If you do that, the following error message appears:

FJQSS is already started. Please restart FJQSS after the executing one ends.

#### Operating procedure

This part explains the operating procedure of FJQSS (Information Collection Tool).

1. Log on to the server, open the console window and execute the information collection command in the console window.

Hereafter, it explains by the example when outputting the information to the default output directory:

```
> cd /opt/FJSVqstl
```

./fjqss\_collect

- 2. The product selection menu appears. It shows the list of the products whose information can be collected.
- 3. Put the number of the product of which you want to collect the information:

```
FJQSS Product Selection

1.MW Product B (Server Side)
2.MW Product B (Client Side)

Please input a number of the product of which you want to collect the information.>>1
```

4. Check the indicated name(s) of the middleware (s) whose investigation information are to be collected, then press the [Y] key to start.

Collection starts and the following indicator appear in the command prompt:

The time and the disk space required for the collection depend on the condition of the target system.

5. Check the contents of the output directory.

After the FJQSS (Information Collection Tool) has completed the collection, the following prompt appears. Verify that the information have been collected in the directory whose name is indicated in the prompt:

This is an example on Linux.

The gathered investigation informations are archived to /var/tmp/··· directory on Solaris, or /tmp/··· directory in Linux.

The following files are created in the output directory (ezcollectYYYYMMDDHHMMSS: YYYYMMDDHHMMSS indicates the time (year, month, day, hours, minutes and seconds) that the collection started). See [2.2 Checking the collected information] for the output directory:

- resultYYYYMMDDHHMMSS.tar.gz

The collected information.

- result.txt

The result log is a text file to indicate if the collection of information has succeeded.

This file indicates the names and the results of collection (OK/NG) of each material.

- ez.log

The activity log of the FJQSS (Information Collection Tool).

#### **Aborting FJQSS (Information Collection Tool)**

This subsection explains how to abort running FJQSS (Information Collection Tool).

Although it is recommended that you wait until the collection ends, you can press [CTRL] + [C] key combination to abort the execution if it is necessary. The following prompt appears if [CTRL] + [C] is pressed and the execution is aborted.

The output directory that has been made in the collecting might remain when the executing is aborted. There is no influence in the operation of FJQSS because the output directory will be newly made when FJQSS is executed again.

Moreover, there is no influence in the operation of FJQSS even if the remaining output directory was deleted.



Do not close the command prompt before the collection ends. if you do so, the temporary directories might remain in the work directory. Please delete the following directories if they remain in the work directory.

ezcollectYYYYMMDDHHMMSS

ezworkYYYYMMDDHHMMSS

See 2.3 Changing a path of the output directory] for details of the work directory.

### 2.1.2 Collecting without showing menu

This subsection explains the operating procedure of FJQSS (Information Collection Tool) without showing the product selection menu.

You can use the procedure in the shell script.

#### Authority to execute

User Authority of the root user is necessary.

#### **Prior confirmation**

Verify the product identification name of the product of which you collect the information by the product information showing command. See 2.1.3 Getting the product identification name] for details.

Set the output directory of the collected information if necessary. The output directory of default is the /var/tmp directory on Solaris, or the /tmp directory on Linux. See [2.3 Changing a path of the output directory] for details of changing the output directory.

#### **Synopsis**

The specification of product identification name by "-pr product identification name>" is necessary if you specify the silent mode option by "-s".

If you execute the command specifying the output directory by the parameter:

```
fjqss_collect -pr collect -pr collect identification name [-s] output_directory
```

If you execute the command specifying the output directory by the environment variable or the environmental configuration file:

```
fjqss_collect -pr collect -pr collect identification name
```

#### Description

The fjqss\_collect command collects the information required for the investigation.

#### Option

-pr

This option specifies the product identification name of the product of which you collect the information.

-S

The command runs without asking to the operator. (It runs in silent mode.)

#### Path to this command

```
/opt/FJSVqstl
```

#### Notes in executing the command

- Do not put fjqss\_collect or its copy into any other directory. They work only if they are in the path shown above in [Path to this command].
- Execute this command immediately after the problem occurs.
- Multiple execution of this command is not available. If you do that, the following error message appears:

```
FJQSS is already started. Please restart FJQSS after the executing one ends.
```

### 2.1.3 Getting the product identification name

This subsection explains how to get the product identification name which you specify when collecting the information without showing Product Selection Menu.

The product identification name can be got by executing the product information showing command.

#### Authority to execute

User Authority of the root user is necessary.

#### **Synopsis**

```
fjqss_mwpinfolist
```

#### Description

This command shows the list of the information of products of which the information collection is available.

Each product's information is shown in one line.

The product's Information is shown in order of "No." (number), "KEY" (product identification name), "DispName" (product name) and "Path" (destination directory of the information)

Example of the result of executing the product information showing command:

```
MW Product List

No. KEY DispName Path

1. PRODUCTA MW PRODUCT A /opt/FJSVPRDA/server/fjqss_tool1

2. PRODUCTB MW PRODUCT B /opt/FJSVPRDB/server/fjqss_tool1

3. PRODUCTC MW PRODUCT C /opt/FJSVPRDC/server/fjqss_tool1
```

If you collect the information of the product No.1, specify the product identification name "PRODUCTA" as the parameter

to collect the information of the "MW PRODUCT A" by the procedure of 2.1.2 Collecting without showing menu].

#### Path to this command

/opt/FJSVqstl

## 2.2 Checking the collected information

The collected information are created in the output directory (ezcollectYYYYMMDDHHMMSS) as file "resultYYYYMMDDHHMMSS.tar.gz".

The output directory is created at the output specified for an environmental configuration file destination at the output destination. If no destination is specified, then the output directory is created at the /var/tmp directory on Solaris, or the /tmp directory on LInux.

You can specify a path of the output directory if you execute the information collecting command in the console window. See [2.3 Changing a path of the output directory], for details of the environmental configuration file and the output destination.

Verify that the following file and directory exist:

- Result log (result.txt)

The result log is a text file to indicate if the collection of information has succeeded.

This file indicates the names and the results of collection (OK/NG) of each material.

- Collected information (resultYYYYMMDDHHMMSS.tar.gz)



Collected information (resultYYYYMMDDHHMMSS.tar.gz) is compressed and archived by gzip and tar. Collected information contains the following files.

- resultdir

This directory contains the collected files.

- result.txt

The result log is a text file to indicate if the collection of information has succeeded.

This file indicates the names and the results of collection (OK/NG) of each material.

- ez.log

The activity log of the FJQSS (Information Collection Tool).

- convert.txt

If the file path name of information collecting is too long and information collecting can't get in former file name, file path name is converted and gathered. Former file path name and Converted file path name are output to convert.txt.convert.txt is output in the form of the following.

Output form

<Former file path name>,<Converted file path name>

convert.txt is output when there was only information collecting for which the conversion of the file path name was necessary.

## 2.3 Changing a path of the output directory

You can change a path of the output directory by two methods below.

- Specifying a path of the output directory in the option of the information collecting command.
- Specifying a path of the output directory by setting the environment variable.
- Specifying a path of the output directory by updating the environmental configuration file.

The method of specifying in the option of the command explains in 2.1.1 Collecting with showing menu and 2.1.2 Collecting without showing menu]. Hereafter, the method of specifying by the environment variable and by the environmental configuration file will be explained.

\_

#### Specifying a path of the output directory by setting the environment variable

You can specify the directory where the output directory will be created by setting the environment variable EZ\_OUTPUTDIR before executing this command.

#### **Description format:**

EZ\_OUTPUTDIR="the directory path where the output directory will be created" export EZ\_OUTPUTDIR

#### **Example:**

If "/etc/work" is specified to %EZ\_OUTPUTDIR%:

EZ\_OUTPUTDIR=/etc/work export EZ\_OUTPUTDIR

Then the directory below will be created as the output directory:

/etc/work/ezcollectYYYYMMDDHHMMSS

#### Specifying a path of the output directory by updating the environmental configuration file

Environmental configuration file (FJQSSconf.txt) is a file that defines the operating environment of FJQSS.

It exists in following directory.

<Directory where the product is installed>/fjqss\_tooln (n=1,2,..,9)

The items that can be specified for environmental configuration file are shown below.

Item	Outline	Omission
OutputDir	The directory path where the output directory will be created (*3).	Possible (*2)
WorkDir	The directory path where the work directory (*1) will be created by FJQSS execution (*3).	Possible(*2)

- (\*1)The work directory is a directory that temporarily stores information necessary for the execution of FJQSS. When FJQSS ends, stored information is automatically deleted.
- (\*2) If the specification is wrong or omitted, the /var/tmp directory (on Solaris), or the /tmp directory (on Linux) is used.
- (\*3) You can specify the directory by the relative path from an absolute path or environmental configuration file.

The format of environmental configuration file is shown below.

#### **Description format:**

OutputDir="the directory path where the output directory will be created" WorkDir=" work directory"

#### Example:

If "/etc/work" is specified to OutputDir, "/etc/tmp" is specified to WorkDIr.

OutputDir=/etc/work WorkDir=/etc/tmp

Then the directory below will be created as the output directory:

/etc/work/ezcollectYYYYMMDDHHMMSS

And the directory below will be created as the work directory:

/etc/tmp/ezworkYYYMMDDHHMMSS



If the plural specifications of the output directory is done, then the specification becomes effective in the following order by priority:

- 1. Specification in the option of the information collecting command.
- 2. Specification by setting the environment variable.
- 3. Specification by updating the environmental configuration file.

# Appendix A Messages and error codes

# A.1 Error Messages

The message which FJQSS outputs can have an error code. See [A.2 Error Codes]. for detail of error codes.

Look for	What to do	
FJQSS is already started.	Check if FJQSS is running. Multiple execution of it is not	
Please restart FJQSS after the executing one ends.	available.	
The file used with FJQSS is not found.	Please contact the Fujitsu engineer, and pass on the	
Please contact a Fujitsu technological member.	PRODUCT_IDENTIFICATION_NAME, PATH_NAME and ERROR_CODE in the message. About detail	
(product= PRODUCT_IDETIFICATION_NAME, errcode= ERROR_CODE)	meanings of the ERROR_CODE, see [A.2 Error Codes].	
Failed in the copy of materials used with FJQSS to the work directory.		
Please contact a Fujitsu technological member.		
(PATH= PATH_NAME, errcode= ERROR_CODE)		
Failed in the access of the file used with FJQSS.		
Please contact a Fujitsu technological member.		
(errcode= ERROR_CODE)		
The file used with FJQSS is not found or broken.		
Please contact a Fujitsu technological member.		
(PATH= PATH_NAME)		
Failed in the access of the directory used with FJQSS.		
Please contact a Fujitsu technological member.		
(PATH= PATH_NAME)		
Failed in the current directory change to the work directory.		
Please contact a Fujitsu technological member.		
(PATH= PATH_NAME)		
Failed to execute FJQSS.Please contact Fujitsu technical staff.		
There is a wrong option in the command line.	Execute FJQSS again after removing the cause according	
Please restart FJQSS with the correct option.	to the left messages. About Environmental configuration file, see [2.3 Changing a path of the output directory].	
Not super-user.		
Please restart FJQSS by super-user		
A memory is insufficient.		
Please stop unnecessary applications and restart FJQSS.		
The product specified by the -pr option is not found.		
Please restart FJQSS with the correct option.		
(product= PRODUCT_IDETIFICATION_NAME)		
Japanese characters are included in the directory name at the 'OutputDir' of environmental file. Japanese characters cannot be used.		
Please correct environmental file and restart FJQSS.		
(environmental file=PATH_NAME)		

Look for	What to do
Japanese characters are included in the directory name at the 'WorkDir' of environmental file. Japanese characters cannot be used.	
Please correct environmental file and restart FJQSS.	
(environmental file=PATH_NAME)	
Failed in making the output directory.	
Please confirm the following.	
-There must be an authority to make the folder.	
-Passing in xx bytes or less must be specified.	
-The length of each directory name must be 255 bytes or less.	
Failed in making the work directory.	
Please confirm the following.	
-There must be an authority to make the folder.	
-Passing in xx bytes or less must be specified.	
-The length of each directory name must be 255 bytes or less.	
Output directory is too long at the 'OutputDir' of environmental file.Please adjust the output directory to xx bytes or less and restart FJQSS	
(environmental file= PATH_NAME)	
Work directory is too long at the 'WorkDir' of environmental file.	
Please adjust the output directory to xx bytes or less and restart FJQSS	
(environmental file= PATH_NAME)	
The output directory is too long specified in the option.	
Please adjust the output directory to %s bytes or less and restart FJQSS.	
Japanese characters are included in the output directory specified in the option. Japanese characters cannot be used.	
Please correct the output directory and restart FJQSS.	
The output directory is too long specified with the environment variable:EZ_OUTPUTDIR.	
Please adjust environment variable:EZ_OUTPUTDIR to xx byte or less.	
Failed to create log.(PATH=PATH_NAME)	
Take action indicated the following items:	
- Create the output folder if it doesn't exist.	
- Set the role to the output folder.	
Products are not installed. Please make sure that the products that are installed.	
The is no installed product of which FJQSS can collect the information. FJQSS ends.	
Stopped FJQSS process.	No user action is necessary.

# A.2 Error Codes

Error Code	Cause
0	No error.

Error Code	Cause
1	Command operand error
2	Insufficient memory
3	Failed to execute without administrator permissions
5	Though "-s" option exists, "-pr" option does not exist
6	FJQSS common component has been already done
10	Middleware products' list file does not exist
11	Middleware product specified with -pr option does not exist
13	The built-in directory of the target middleware product of collecting information does not exist
15	Japanese characters are included in the directory name at the output directory of environmental file
16	Japanese characters are included in the directory name at the work directory of environmental file
17	Too long path name is specified in the directory name at the output directory of environmental file
18	Too long path name is specified in the directory name at the work directory of environmental file
19	Failed to create the output directory
20	Failed to create the work directory
21	Failed to access the directory
22	Failed to copy the common materials and built-in materials of the product directory to the work directory
23	Failed to show the information of the product to collect its information materials
24	Failed to change the current directory to the work directory
25	Too long path name is specified to the parameter that specifies the output directory
26	Japanese characters are included in the parameter that specifies the output directory
27	Products of which collecting information is available are not installed
28	Failed to execute FJQSS