

# FUJITSU Enterprise Postgres 14

## Operation Guide

Windows/Linux



# FUJITSU Enterprise Postgres 14

## 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 Global Meta Cache feature.

### [Chapter 13 Local Meta Cache Limit](#)

Describes how to use Local Meta Cache Limit feature.

## [Chapter 14 Backup/Recovery Using the Copy Command](#)

Describes backup and recovery using the copy command created by the user.

## [Chapter 15 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: January 2022
---------------------------

## **Copyright**

Copyright 2015-2022 FUJITSU LIMITED

# Contents

---

Chapter 1 Operating FUJITSU Enterprise Postgres.....	1
1.1 Operating Methods.....	1
1.2 Starting WebAdmin.....	2
1.2.1 Logging in to WebAdmin.....	2
1.3 Starting pgAdmin.....	3
1.3.1 Starting pgAdmin.....	4
1.3.2 Adding an Instance.....	4
1.3.3 Connecting/Disconnecting an Instance.....	6
1.4 Operations Using Commands.....	6
1.5 Operating Environment of FUJITSU Enterprise Postgres.....	7
1.5.1 Operating Environment.....	7
1.5.2 File Composition.....	9
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.....	12
2.1 Starting and Stopping an Instance.....	12
2.1.1 Using WebAdmin.....	12
2.1.2 Using Server Commands.....	14
2.2 Creating a Database.....	15
2.2.1 Using pgAdmin.....	15
2.2.2 Using Client Commands.....	16
Chapter 3 Backing Up the Database.....	17
3.1 Periodic Backup.....	18
3.2 Backup Methods.....	18
3.2.1 Using WebAdmin.....	18
3.2.2 Using Server Commands.....	19
Chapter 4 Configuring Secure Communication Using Secure Sockets Layer.....	22
4.1 Configuring Communication Data Encryption.....	22
4.1.1 Issuing a Certificate.....	23
4.1.2 Deploying a Server Certificate File and a Server Private Key File.....	23
4.1.3 Distributing a CA Certificate File to the Client.....	23
4.1.4 Configuring the Operating Environment for the Database Server.....	23
4.1.5 Configuring the Operating Environment for the Client.....	23
4.1.6 Performing Database Multiplexing.....	24
Chapter 5 Protecting Storage Data Using Transparent Data Encryption.....	25
5.1 Protecting Data Using Encryption.....	25
5.2 Setting the Master Encryption Key.....	26
5.3 Opening the Keystore.....	27
5.4 Encrypting a Tablespace.....	27
5.5 Checking an Encrypted Tablespace.....	28
5.6 Managing the Keystore.....	29
5.6.1 Changing the Master Encryption Key.....	29
5.6.2 Changing the Keystore Passphrase.....	29
5.6.3 Enabling Automatic Opening of the Keystore.....	29
5.6.4 Backing Up and Recovering the Keystore.....	30
5.7 Backing Up and Restoring/Recovering the Database.....	32
5.8 Importing and Exporting the Database.....	34
5.9 Encrypting Existing Data.....	34
5.10 Operations in Cluster Systems.....	34
5.10.1 HA Clusters that do not Use Database Multiplexing.....	35
5.10.2 Database Multiplexing Mode.....	35

5.11 Security-Related Notes.....	36
5.12 Tips for Installing Built Applications.....	37
<b>Chapter 6 Data Masking.....</b>	<b>38</b>
6.1 Masking Policy.....	38
6.1.1 Masking Target.....	39
6.1.2 Masking Type.....	39
6.1.3 Masking Condition.....	39
6.1.4 Masking Format.....	40
6.2 Usage Method.....	42
6.2.1 Creating a Masking Policy.....	43
6.2.2 Changing a Masking Policy.....	44
6.2.3 Confirming a Masking Policy.....	44
6.2.4 Enabling and Disabling a Masking Policy.....	45
6.2.5 Deleting a Masking Policy.....	46
6.3 Data Types for Masking.....	46
6.4 Security Notes.....	47
<b>Chapter 7 Periodic Operations.....</b>	<b>48</b>
7.1 Configuring and Monitoring the Log.....	48
7.2 Monitoring Disk Usage and Securing Free Space.....	48
7.2.1 Monitoring Disk Usage.....	48
7.2.2 Securing Free Disk Space.....	48
7.3 Automatically Closing Connections.....	49
7.4 Monitoring the Connection State of an Application.....	49
7.4.1 Using the View (pg_stat_activity).....	50
7.4.2 Using pgAdmin.....	51
7.5 Reorganizing Indexes.....	51
7.6 Monitoring Database Activity.....	52
7.6.1 Information that can be Collected.....	53
7.6.2 Collection Configuration.....	54
7.6.3 Information Reset.....	55
<b>Chapter 8 Streaming Replication Using WebAdmin.....</b>	<b>56</b>
8.1 Creating a Standby Instance.....	56
8.2 Promoting a Standby Instance.....	58
8.3 Converting an Asynchronous Replication to Synchronous.....	58
8.4 Converting a Synchronous Replication to Asynchronous.....	59
8.5 Joining a Replication Cluster.....	59
<b>Chapter 9 Installing and Operating the In-memory Feature.....</b>	<b>61</b>
9.1 Installing Vertical Clustered Index (VCI).....	61
9.1.1 Evaluating whether to Install VCI.....	61
9.1.2 Estimating Resources.....	62
9.1.3 Setting up.....	62
9.1.3.1 Setting Parameters.....	62
9.1.3.2 Installing the Extensions.....	63
9.1.3.3 Creating a VCI.....	63
9.1.3.4 Confirming that the VCI has been Created.....	64
9.1.4 Data that can Use VCI.....	64
9.1.4.1 Relation Types.....	64
9.1.4.2 Data Types.....	65
9.2 Operating VCI.....	66
9.2.1 Commands that cannot be Used for VCI.....	67
9.2.2 Data Preload Feature.....	68
<b>Chapter 10 Parallel Query.....</b>	<b>69</b>
10.1 CPU Load Calculation.....	69

10.2 Increase of Workers during Runtime.....	69
<b>Chapter 11 High-Speed Data Load.....</b>	<b>70</b>
11.1 Installing High-Speed Data Load.....	70
11.1.1 Deciding whether to Install.....	70
11.1.2 Estimating Resources.....	70
11.1.3 Setup.....	71
11.1.3.1 Setting Parameters.....	71
11.1.3.2 Installing the Extension.....	72
11.2 Using High-Speed Data Load.....	72
11.2.1 Loading Data.....	72
11.2.2 Checking Progress.....	73
11.2.3 Recovering from a Data Load that Ended Abnormally.....	74
11.3 Removing High-Speed Data Load.....	75
11.3.1 Removing the Extension.....	75
<b>Chapter 12 Global Meta Cache.....</b>	<b>77</b>
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
<b>Chapter 13 Local Meta Cache Limit.....</b>	<b>79</b>
13.1 Usage.....	79
13.1.1 Deciding Whether to Enable the Local Meta Cache Limit Feature.....	79
13.1.2 How to Set Parameters for the Local Meta Cache Limit Feature.....	79
13.1.3 Cache Removal when Local Meta Cache Limit is Enabled.....	79
13.1.4 Performance Impact and Parameter Tuning of the Local Meta Cache Limit Feature.....	80
<b>Chapter 14 Backup/Recovery Using the Copy Command.....</b>	<b>82</b>
14.1 Configuration of the Copy Command.....	82
14.2 Backup Using the Copy Command.....	86
14.3 Recovery Using the Copy Command.....	87
14.4 Copy Command Interface.....	88
14.4.1 Copy Command for Backup.....	88
14.4.2 Copy Command for Recovery.....	90
<b>Chapter 15 Actions when an Error Occurs.....</b>	<b>92</b>
15.1 Recovering from Disk Failure (Hardware).....	93
15.1.1 Using WebAdmin.....	93
15.1.2 Using Server Command.....	94
15.2 Recovering from Data Corruption.....	98
15.2.1 Using WebAdmin.....	99
15.2.2 Using the pgx_rcvall Command.....	99
15.3 Recovering from an Incorrect User Operation.....	100
15.3.1 Using WebAdmin.....	100
15.3.2 Using the pgx_rcvall Command.....	101
15.4 Actions in Response to an Application Error.....	102
15.4.1 When using the view (pg_stat_activity).....	103
15.4.2 Using the ps Command.....	103
15.4.3 Using pgAdmin.....	104
15.5 Actions in Response to an Access Error.....	105
15.6 Actions in Response to Insufficient Space on the Data Storage Destination.....	105
15.6.1 Using a Tablespace.....	106

15.6.2 Replacing the Disk with a Larger Capacity Disk.....	106
15.6.2.1 Using WebAdmin.....	106
15.6.2.2 Using Server Commands.....	107
15.7 Actions in Response to Insufficient Space on the Backup Data Storage Destination.....	108
15.7.1 Temporarily Saving Backup Data.....	108
15.7.1.1 Using WebAdmin.....	109
15.7.1.2 Using Server Commands.....	109
15.7.2 Replacing the Disk with a Larger Capacity Disk.....	112
15.7.2.1 Using WebAdmin.....	112
15.7.2.2 Using Server Commands.....	113
15.8 Actions in Response to Insufficient Space on the Transaction Log Storage Destination.....	116
15.8.1 Replacing the Disk with a Larger Capacity Disk.....	117
15.8.1.1 Using WebAdmin.....	117
15.8.1.2 Using Server Commands.....	118
15.9 Errors in More Than One Storage Disk.....	119
15.10 Actions in Response to Instance Startup Failure.....	119
15.10.1 Errors in the Configuration File.....	119
15.10.2 Errors Caused by Power Failure or Mounting Issues.....	120
15.10.3 Other Errors.....	120
15.10.3.1 Using WebAdmin.....	120
15.10.3.2 Using Server Commands.....	120
15.11 Actions in Response to Failure to Stop an Instance.....	121
15.11.1 Using WebAdmin.....	121
15.11.2 Using Server Commands.....	121
15.11.2.1 Stopping the Instance Using the Fast Mode.....	121
15.11.2.2 Stopping the Instance Using the Immediate Mode.....	121
15.11.2.3 Forcibly Stopping the Server Process.....	121
15.12 Actions in Response to Failure to Create a Streaming Replication Standby Instance.....	122
15.13 Actions in Response to Error in a Distributed Transaction.....	122
15.14 I/O Errors Other than Disk Failure.....	124
15.14.1 Network Error with an External Disk.....	124
15.14.2 Errors Caused by Power Failure or Mounting Issues.....	124
15.15 Anomaly Detection and Resolution.....	124
15.15.1 Port Number and Backup Storage Path Anomalies.....	124
15.15.2 Mirroring Controller Anomalies.....	125
<b>Appendix A Parameters.....</b>	<b>126</b>
<b>Appendix B System Administration Functions.....</b>	<b>132</b>
B.1 WAL Mirroring Control Functions.....	132
B.2 Transparent Data Encryption Control Functions.....	132
B.3 Data Masking Control Functions.....	133
B.3.1 pgx_alter_confidential_policy.....	133
B.3.2 pgx_create_confidential_policy.....	139
B.3.3 pgx_drop_confidential_policy.....	142
B.3.4 pgx_enable_confidential_policy.....	143
B.3.5 pgx_update_confidential_values.....	144
B.4 VCI Data Load Control Function.....	145
B.5 High-Speed Data Load Control Functions.....	146
<b>Appendix C System Views.....</b>	<b>147</b>
C.1 pgx_tablespaces.....	147
C.2 pgx_stat_lwlock.....	147
C.3 pgx_stat_latch.....	147
C.4 pgx_stat_walwriter.....	148
C.5 pgx_stat_sql.....	148
C.6 pgx_stat_gmc.....	149
C.7 pgx_stat_progress_loader.....	149



Appendix D Tables Used by Data Masking.....	150
D.1 pgx_confidential_columns.....	150
D.2 pgx_confidential_policies.....	150
D.3 pgx_confidential_values.....	151
Appendix E Tables Used by High-Speed Data Load.....	152
E.1 pgx_loader_state.....	152
Appendix F Starting and Stopping the Web Server Feature of WebAdmin.....	153
F.1 Starting the Web Server Feature of WebAdmin.....	153
F.2 Stopping the Web Server Feature of WebAdmin.....	153
Appendix G WebAdmin Wallet.....	155
G.1 Creating a Credential.....	155
G.2 Using a Credential.....	156
Appendix H WebAdmin Disallow User Inputs Containing Hazardous Characters.....	157
Appendix I Copy Command Samples that Use the Advanced Copy Feature of the ETERNUS Disk Array.....	158
Appendix J Collecting Failure Investigation Data.....	160
Index.....	161

# 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.



Note

- .....
- 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.	

\*1: 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.



### Example

For a server with IP address "192.0.2.0" and port number "27515"

```
http://192.0.2.0:27515/
```

Display the startup windows. From this window you can log in to WebAdmin or access the product documentation.



### Note

- 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



### Point

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 (64bit) xSPz" (where "x" is the product version, and "z" is the product level (x SPz)).

Windows(R) 8.1

From the [Apps] view, start [pgAdmin 4 (64bit) (x SPz)].

Other than above

Click [Start] >> [All apps] >> [FUJITSU Enterprise Postgres Client(64bit) x SPz] and start [pgAdmin 4 (64bit) (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.

The screenshot shows the 'Create - Server' dialog box with the following fields and controls:

- Name:** An empty text input field with a red warning triangle icon to its left.
- Server group:** A dropdown menu showing 'Servers' with a list icon and a downward arrow.
- Background:** A checkbox with an 'X' icon, currently unchecked.
- Foreground:** A checkbox with an 'X' icon, currently unchecked.
- Connect now?:** A toggle switch that is currently turned on (blue).
- Comments:** A large empty text area.

A red error message at the bottom of the dialog reads: **'Name' cannot be empty.** The bottom bar contains three buttons: 'Close', 'Reset', and 'Save'.

[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.



## Note

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.



## Note

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.



## See

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.



\*1: To distribute the I/O load, place the transaction log on a different disk from the data storage destination.



Table 1.1 OS resources

Type	Role
Shared memory	Used when a database process exchanges information with an external process.
Semaphore	

Table 1.2 FUJITSU Enterprise Postgres client resources

Type	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

Type	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.

Type	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
Program files	Y	Multiple	Note that "<x>" indicates the product version. /opt/fsepv<x>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

\*1: Set manually when using the applicable feature.



### Note

- Do not place files for use with FUJITSU Enterprise Postgres in a directory mounted over the network except when creating a database space in a storage device on a network.  
Examples include NFS (Network File System) and CIFS (Common Internet File System).  
This is because the database might hang if the network fails.
- 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.



See

.....  
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;
```



Note

- .....
- 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 or later using `pg_upgrade`, additional steps must be performed because of the incompatibility of the VCI extension between FUJITSU Enterprise Postgres 12 or later 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](#)



### Point

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.



### Note


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.	

\*1: 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 .
3. In the [Instances] tab, click , and select the Immediate mode to stop the instance.

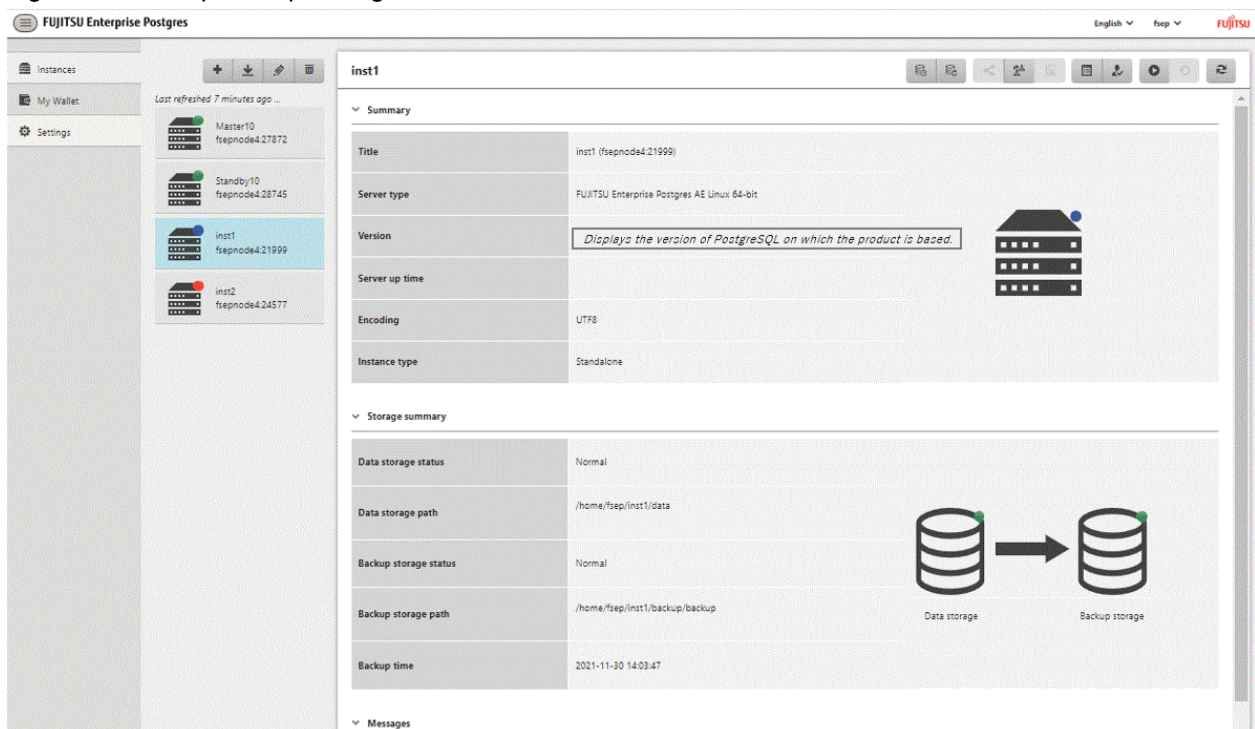
## Checking the operating status of an instance

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.
	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



The screenshot displays the Fujitsu Enterprise Postgres WebAdmin interface. On the left, a sidebar shows the 'Instances' tab with a list of instances: Master10 (fsepnode4.27872), Standby10 (fsepnode4.28745), inst1 (fsepnode4.21999), and inst2 (fsepnode4.24577). The 'inst1' instance is highlighted with a blue background and a green status indicator. The main panel shows the details for 'inst1', including a 'Summary' section with fields for Title, Server type, Version, Server up time, Encoding, and Instance type. Below this is a 'Storage summary' section with fields for Data storage status, Data storage path, Backup storage status, Backup storage path, and Backup time. A diagram illustrates the flow from 'Data storage' to 'Backup storage'.

## Note

- 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.

## See

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.

## Example

```
> pg_ctl start -D /database/inst1
```

## Note

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.

### Example

```
> 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.

### Example

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/inst1
pg_ctl: no server running.
```

### See

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 (14.0)
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.



**See**

.....  
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.



## Information

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 14 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:

$$\text{backupTime} = \text{dataStorageDestinationUsage} / \text{diskWritePerformance} \times 1.5$$

- *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.



## Note

- 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.

---

 **Information**

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.

 **Note**

- 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.

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	<i>'installationDirectory/bin/pgx_walcopy.cmd "%p" "backupDataStorageDestinationDirectory/archived_wal/%f"'</i>	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.



#### Example

```
> pgx_dmpall -D /database/inst1
```



#### Note

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_rvfall` command to check the backup status.

Specify the following values in the `pgx_rvfall` 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 -l -D /database/inst1
Date                Status           Dir
2020-05-01 13:30:40 COMPLETE        /backup/inst1/2020-05-01_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.

### 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.

### Example

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.

```
postgres=# SELECT pg_create_restore_point('batch_20200503_1');
LOG:  restore point "batch_20200503_1" created at 0/20000E8
STATEMENT:  select pg_create_restore_point('batch_20200503_1');
pg_create_restore_point
-----
0/20000E8
(1 row)
```

Refer to "15.3.2 Using the pgx\_rcvall Command" for information on using a restore point to recover the database.

### Note

- 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

---



See

.....  
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

---



See

.....  
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.



See

.....  
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.

### Note

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

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
>
```



### Point

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

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`.



## Point

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.



## Note

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, spcncalgo  
FROM pg_tablespace ts, pgx_tablespaces tsx  
WHERE ts.oid = tsx.spctablespace;
```

## Example

```
postgres=# SELECT spcname, spcencalgo FROM pg_tablespace ts, pgx_tablespaces tsx WHERE ts.oid =
tsx.spcnamespace;
   spcname   | spcencalgo
-----+-----
 pg_default  | none
 pg_global   | none
 secure_tablespace | AES256
(3 rows)
```

## See

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
>
```

## See

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.

## Example

```
# 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.

## Note

- 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

## Point

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, 2020.

```
> pgx_dmpall -D /database/inst1
> cp -p /key/store/location/keystore.ks /keybackup/keystore_20200501.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, 2020.

```
> psql -c "SELECT pgx_set_master_key('passphrase')" postgres
> cp -p /key/store/location/keystore.ks /keybackup/keystore_20200505.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.

## Example

- Restore the keystore containing the latest master encryption key as of May 5, 2020.

```
> cp -p /keybackup/keystore_20200505.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 2020.

```
> cp -p /keybackup/keystore_20200501.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.

## See

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.



### Example

---

#### - Back up the database and the keystore on May 1, 2020.

```
> pgx_dmpall -D /database/inst1
> cp -p /key/store/location/keystore.ks /keybackup/keystore_20200501.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, 2020.

```
> cp -p /keybackup/keystore_20200501.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.



See

---

- 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



### Example

.....

```
# 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;
```



### See

.....

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.



See

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.



Point

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.



See

.....  
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.

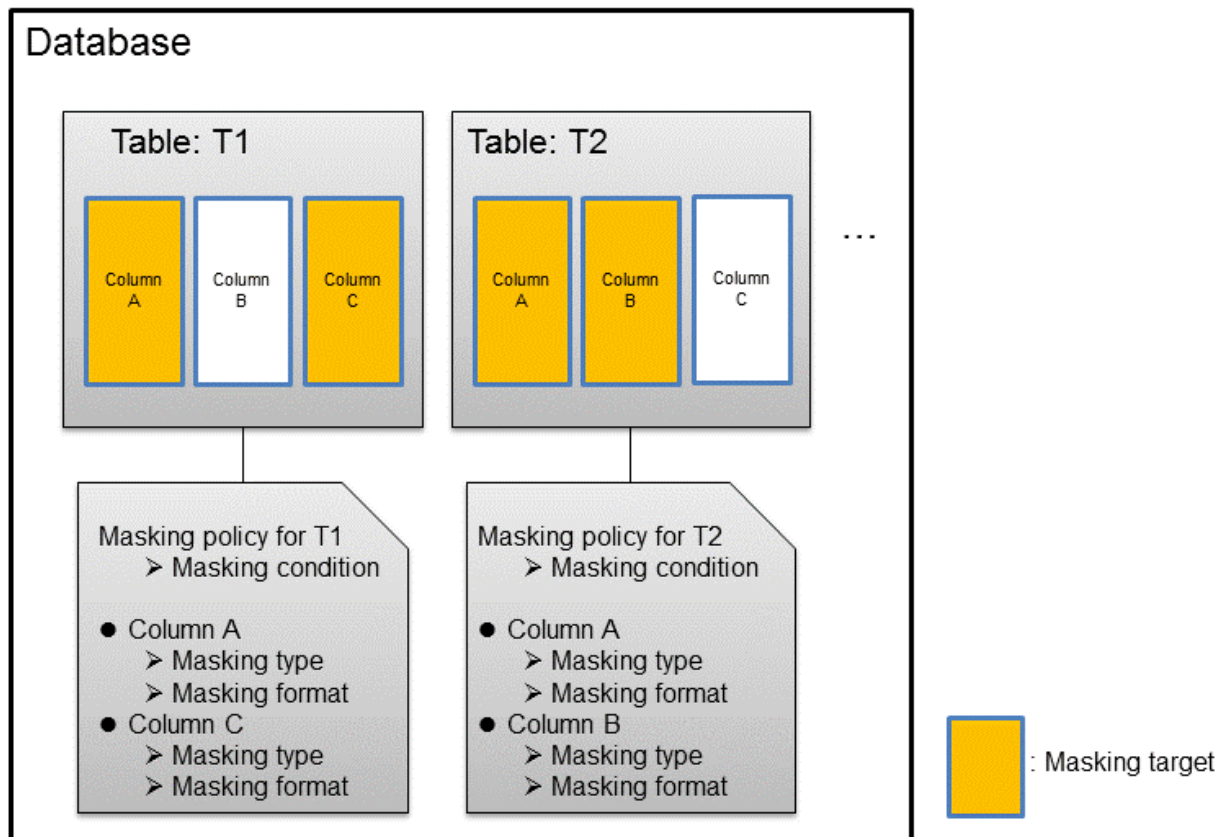
## Note

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



## Note

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.

 **Information**

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.



 **See**

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	<p><i>'replacementCharacter, startPosition, endPosition'</i></p> <ul style="list-style-type: none"> <li>- <i>replacementCharacter</i>: Specify the number to display. Specify a value from 0 to 9.</li> <li>- <i>startPosition</i>: Specify the start position of masking. Specify a positive integer.</li> <li>- <i>endPosition</i>: Specify the end position of masking. Specify a positive integer that is greater than <i>startPosition</i>.</li> </ul> <p> <b>Example</b></p> <p>Specify as below to change the values from the 1st to 5th digits to 9.</p> <p>function_parameters := '9, 1, 5'</p> <p>In this example, if the original data is "123456789", it will be changed to "999996789".</p>
Character type	<p><i>'inputFormat, outputFormat, replacementCharacter, startPosition, endPosition'</i></p> <ul style="list-style-type: none"> <li>- <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.</li> <li>- <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.</li> <li>- <i>replacementCharacter</i>: Specify any single character. If you want to output a single quotation mark, specify two of them consecutively.</li> <li>- <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.</li> </ul>

Category	Method of specifying function_parameters
	<p>- <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>.</p> <p> <b>Example</b></p> <p>Specify as below to mask a telephone number other than the first three digits using *.</p> <p>function_parameters := 'VVVFVVVVFVVVV, VVV-VVVV-VVVV, *, 4, 11'</p> <p>In this example, if the original data is "012-3156-7890", it will be changed to "012-****-*****".</p>
Date/timestamp type	<p>'MDYHMS'</p> <ul style="list-style-type: none"> <li>- 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.</li> <li>- 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.</li> <li>- 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.</li> <li>- 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.</li> <li>- 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.</li> <li>- 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.</li> </ul> <p> <b>Example</b></p> <p>Specify as below to mask hour, minute, and second and display 00:00:00.</p> <p>function_parameters := 'MDYh0m0s0'</p> <p>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".</p>

 **See**

- 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.

 **Example**

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.



### Note



- 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 *n*th 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
```



### Note



You must always prepend "pgx\_datamasking" to the "shared\_preload\_libraries" parameter.



### Information



- 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



### Note

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: 'l=1'

```
postgres=# select pgx_create_confidential_policy(table_name := 't1', policy_name := 'p1',
expression := 'l=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.

```
postgres=# select * from t1;
 c1 |      c2
----+-----
  0 | 012-3456-7890
  0 | 012-3456-7891
  0 | 012-3456-7892
(3 row)
```



### See

- Refer to "[B.3.2 pgx\\_create\\_confidential\\_policy](#)" for information on the `pgx_create_confidential_policy` system management function.



```

public      | t1      | p1      | c2      | PARTIAL      | VVVFVVVVFVVVV, VVV-VVVV-
VVVV, *, 4, 11 |
(2 row)

```

2. Confirm information about the masking policy content  
Refer to `pgx_confidential_policies` to confirm the masking policy content.

```

postgres=# select * from pgx_confidential_policies;
 schema_name | table_name | policy_name | expression | enable | policy_description
-----+-----+-----+-----+-----+-----
public      | t1      | p1      | 1=1      | t      |
(1 row)

```



- 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.

```

postgres=# select * from t1;
 c1 |      c2
-----+-----
  1 | 012-3456-7890
  2 | 012-3456-7891
  3 | 012-3456-7892
(3 row)

```

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)

```

 See

- 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
(1 row)

```

2. Confirm the displayed data

Confirm that the original data is displayed by deleting the masking policy.

```

postgres=# select * from t1;
 c1 |      c2
-----+-----
  1 | 012-3456-7890
  2 | 012-3456-7891
  3 | 012-3456-7892
(3 row)

```

 See

- Refer to "B.3.3 [pgx\\_drop\\_confidential\\_policy](#)" for information on the `pgx_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	Data type	Masking type		
		Full masking	Partial masking	Regular expression masking
Character type	character varying( <i>n</i> )	Y	Y	Y
	varchar( <i>n</i> )	Y	Y	Y
	character( <i>n</i> )	Y	Y	Y
	char( <i>n</i> )	Y	Y	Y
Date/timestamp type	date	Y	Y	N
	timestamp	Y	Y	N

### Note

Even if the data type can be masking, if the data is a special value (NaN, Infinity, -Infinity), it is not.

## 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.



See

- 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.



Information

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.

### Note

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.

### See

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.



Example

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 ]-----+-----
datid          | 13003
datname        | db01
pid            | 4638
leader_pid     |
usesysid      | 10
username       | fsep
application_name | apl01
client_addr    | 192.33.44.15
client_hostname |
client_port    | 27500
backend_start  | 2020-02-24 09:09:21.730641+09
xact_start     | 2020-02-24 09:09:23.858727+09
query_start    | 2020-02-24 09:09:23.858727+09
state_change   | 2020-02-24 09:09:23.858834+09
wait_event_type | Client
wait_event     | ClientRead
state          | idle in transaction
backend_xid    |
backend_xmin   |
query_id      |
query         | begin;
backend_type   | client backend
```




See

- 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](#)".

### Note

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.

### See

Refer to "Routine Reindexing" under "Server Administration" in the PostgreSQL Documentation for information on reorganizing indexes by periodically executing the REINDEX command.

### Point

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	index_size	table_size
public	pgbench_accounts_pkey	0.16	2208 KB	13 MB
pg_catalog	pg_depend_depender_index	0.6	224 KB	368 KB
pg_catalog	pg_depend_reference_index	0.58	216 KB	368 KB
...				



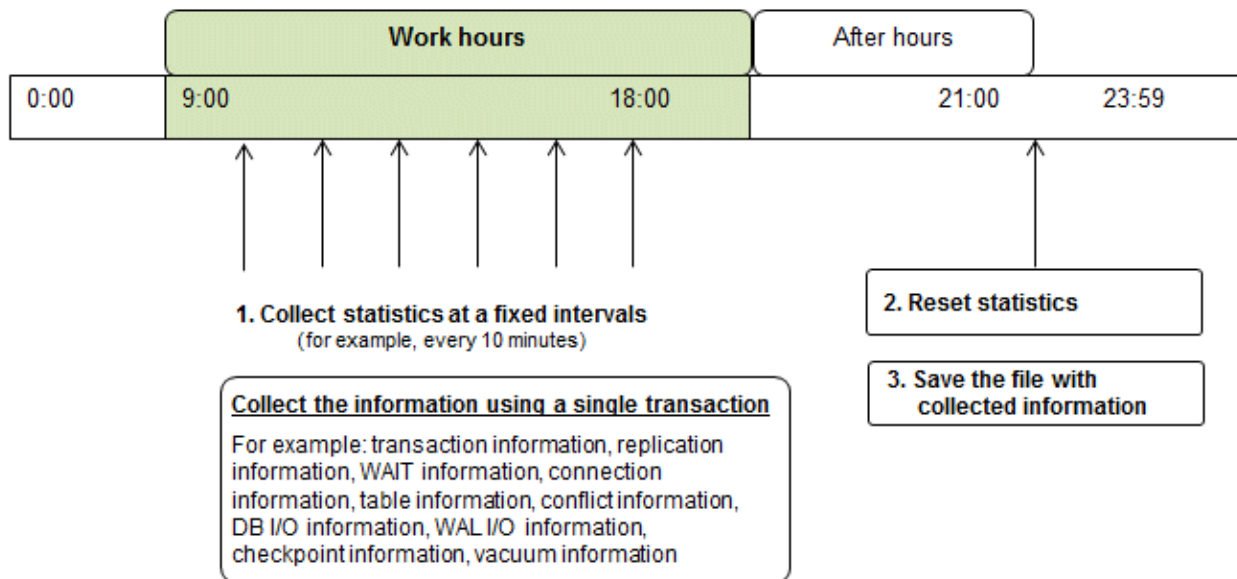
Refer to "Notes on Application Compatibility" in the Application Development Guide for information on maintaining application compatibility.

## 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.

### Note

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



See

Refer to "Monitoring Database Activity" under "Server Administration" in the PostgreSQL Documentation for information on 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.
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



See

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	

View name	Parameter
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	<p>Reset some cluster-wide statistics counters to zero, depending on the argument (requires superuser privileges).</p> <p>Calling pg_stat_reset_shared('lwlock') will zero all counters shown in pgx_stat_lwlock.</p> <p>Similarly, in the following cases, all values of the pertinent statistics counter are reset:</p> <ul style="list-style-type: none"> <li>- If pg_stat_reset_shared('latch') is called: All values displayed in pgx_stat_latch</li> <li>- If pg_stat_reset_shared('walwriter') is called: All values displayed in pgx_stat_walwriter</li> <li>- If pg_stat_reset_shared('sql') is called: All values displayed in pgx_stat_sql</li> <li>- If pg_stat_reset_shared('gmc') is called: All values except size column in pgx_stat_gmc</li> </ul>



See

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.

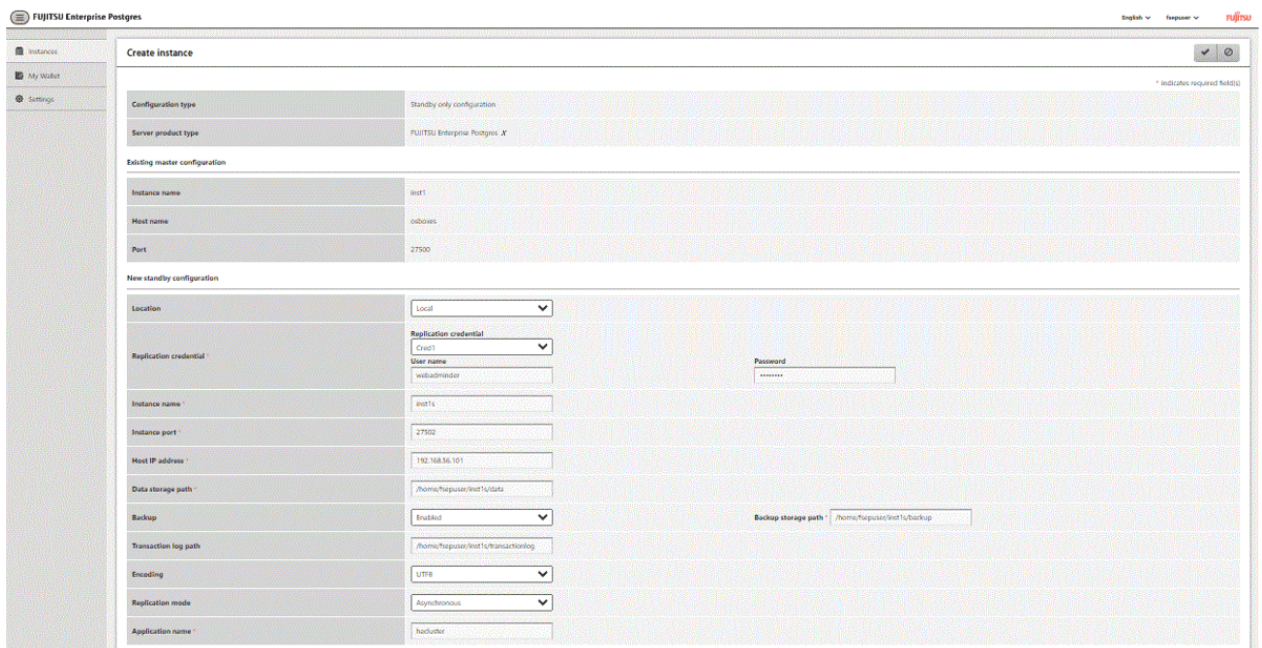
## Point

- 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:



## Note

- 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 .
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.

## Note

- 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  to open the [Join replication cluster] dialog box.

For FUJITSU Enterprise Postgres 13 and earlier instances, select [Restart later] or [Restart now], and then click [Yes] to set up the standby instance.  
For FUJITSU Enterprise Postgres 14, click [Yes] or [Restart now], and then click [Yes] to set up the standby instance.
8. Upon successful completion, the confirmation dialog box will be displayed.
9. 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.

### Information

The operations on VCI can continue even if the memory configured for VCI is insufficient by using VCI data on the disk.

### See

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

### Example

```
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
```

### Note

An error occurs if you use VCI to start instances when procfs is not mounted. Ensure that procfs is mounted before starting instances.

### See

- 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;
```

### Note

- 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);
```

## Note

- 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.

#### Example

```
db01=# SELECT indexdef FROM pg_indexes WHERE indexdef LIKE '%vci%';
          indexdef
-----
CREATE INDEX idx_vci ON table01 USING vci (col01, col02)
(1 row)
```

## 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( <i>n</i> )	Y
	char( <i>n</i> )	Y
	nchar	Y
	nvarchar( <i>n</i> )	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( <i>n</i> )	Y
	bit varying( <i>n</i> )	Y
Text search	tsvector	N
	tsquery	N
UUID	uuid	Y
XML	xml	N
JSON	json	N
	jsonb	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 (Basic type, enumerated type, composite type, and range type)	-	N

## 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 if a VCI is specified.
	SET TABLESPACE	
	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 CLASS	RENAME TO	The subcommands on the left cannot be performed if a VCI is specified.
	OWNER TO	
	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 <i>name</i> [ OWNED BY <i>roleName</i> ] SET TABLESPACE <i>newTablespace</i>	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 ] <i>colName</i> [ RESTRICT   CASCADE ]	A column 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.
	ALTER [ COLUMN ] <i>colName</i> [ SET DATA ] TYPE <i>dataType</i> [ COLLATE <i>collation</i> ] [ USING <i>expr</i> ]	
	CLUSTER ON <i>indexName</i>	A VCI cannot be specified.
	REPLICA IDENTITY {DEFAULT   USING INDEX <i>indexName</i>   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 if a VCI is specified.  This operation is not supported in VCI.
	CONCURRENTLY	
	[ ASC   DESC ]	
	[ 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 the --index 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');
```



See

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.

## Note

- 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`.

#### See

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.

#### Example

```
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.

### Note

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.

### See

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
<code>max_prepared_transactions</code>	Add the number of transactions that can be prepared by parallel 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.	Mandatory
<code>max_worker_processes</code>	Number of parallel workers to perform data load.	Mandatory
<code>max_parallel_workers</code>	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

### Example

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
```

 **Note**

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.  The default value is recommended.	Mandatory

 **See**

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.

 **Example**

The example below installs the extension on the "postgres" database.

```
postgres=# CREATE EXTENSION pgx_loader;
CREATE EXTENSION
```

 **Note**

- 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.

 **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
```

## Point

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.

## Note

The data inserted using this feature is dumped as a COPY command by the `pg_dump` command and the `pg_dumpall` command.

## See

- 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 Checking Progress

If you are performing a data load with a large external file as input, you can verify that the process is continuing by getting progress information during the load. Progress information can be obtained from the `pgx_stat_progress_loader` view. This view displays the sum of the progress information of the back-end process and the number of parallel worker processes. Search the `pgx_stat_progress_loader` view, for example, with a SELECT statement, to locate the appropriate row. After running the `pgx_loader` command, look in the `pg_stat_activity` view and locate a row in the `pg_stat_activity` view with the PID obtained.

## Example

1. See the `pg_stat_activity` view. (9311 for back-end processes, 9312, 9313, 9314 for worker processes)

```
postgres=# select pid, application_name, backend_type from pg_stat_activity
 pid | application_name | backend_type
-----+-----+-----
 6216 |                  | autovacuum launcher
 6218 |                  | logical replication launcher
 6271 | psql             | client backend
 9311 | pgx_loader       | client backend
 9312 |                  | parallel loader for PID 9311
 9313 |                  | parallel loader for PID 9311
 9314 |                  | parallel loader for PID 9311
 6214 |                  | background writer
 6213 |                  | checkpointer
 6215 |                  | walwriter
```

2. Check the information in the `pgx_stat_progress_loader` view.

```
postgres=# SELECT * FROM pgx_stat_progress_loader
 pid | datid | datname | relid | command | type | bytes_processed | bytes_total | tuples_processed | tuples_excluded
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----
 9311 | 222   | testdb  | 333   | COPY FROM | FILE | 192000          | 450000      | 189000          |
 |      |         |       |         |      | 3000            |              |                  |
```

Refer to "C.7 [pgx\\_stat\\_progress\\_loader](#)" for information on the `pgx_stat_progress_loader` view.

## Note

When you run the `pgx_loader` command, the PostgreSQL `pg_stat_progress_copy` view prints the progress of the back-end process and the number of parallel worker processes on each line. The backend process progress information `tuples_processed`, `tuples_excluded` is 0. Also, `bytes_processed` and `bytes_total` for worker processes are 0.

```
postgres=# SELECT * FROM pg_stat_progress_copy
pid | datid | datname | relid | command | type | bytes_processed | bytes_total | tuples_processed | tuples_excluded
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----
9311 | 222   | testdb  | 333   | COPY FROM | FILE | 192000          | 450000      | 0                | 0
|
9312 | 222   | testdb  | 333   | COPY FROM | FILE | 0               | 0           | 63000            | 0
|
9313 | 222   | testdb  | 333   | COPY FROM | FILE | 0               | 0           | 63000            | 0
|
9314 | 222   | testdb  | 333   | COPY FROM | FILE | 0               | 0           | 63000            | 0
|
```

Refer to "`pg_stat_progress_copy` View" in the PostgreSQL Documentation for information on the `pg_stat_progress_copy` view.

## 11.2.3 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".

### Example

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.

```
postgres=# SELECT gid, state FROM pgx_loader.pgx_loader_state WHERE
postgres=# role_oid IN (SELECT oid FROM pg_roles WHERE rolname = 'myrole') AND
postgres=# relation_oid IN (SELECT relid FROM pg_stat_all_tables WHERE
postgres=# relname = 'tbl');
   gid          | state
-----+-----
pgx_loader:9590 | rollback
pgx_loader:9591 | commit
pgx_loader:9589 | rollback
(3 rows)
```

2. Refer to the `pg_prepared_xacts` system view.

Check if the in-doubt transactions retrieved above exist.

## Example

The example below checks if in-doubt transactions with the global transaction identifiers `pgx_loader:9589` and `pgx_loader:9590` exist.

```
postgres=# SELECT gid FROM pg_prepared_xacts WHERE gid IN ('pgx_loader:9589', 'pgx_loader:9590');
gid
-----
pgx_loader:9590
pgx_loader:9589
(2 rows)
```

## See

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.

## Example

The example below completes the in-doubt transactions prepared for table `tbl`.

```
$ pgx_loader recovery -t tbl
```

## Point

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 "[15.13 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.

## Example

The example below removes the extension on the "postgres" database.

```
postgres=# DROP EXTENSION pgx_loader;
DROP EXTENSION
```

 Note

- 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 increase the time it takes to execute a single SQL to reference a meta cache on shared memory, 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 Cache. 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_metacache` 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
<code>pgx_global_metacache</code>	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.



## Example

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	GlobalCatcache	Waiting to find, add, and remove meta caches in the GMC area.
IPC	GMCSweep	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.



## Note

If GMCSweep 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 Local Meta Cache Limit

Local Meta Cache Limit feature limits the size of a Local Meta Cache by removing it if it has not been accessed for a long time.



This feature is available only in the Advanced Edition.

## 13.1 Usage

Describes how to use the Local Meta Cache Limit feature.

### 13.1.1 Deciding Whether to Enable the Local Meta Cache Limit Feature

Refer to “[Appendix A Parameters](#)”, after estimating the total amount of memory to be used as the catalog cache and relation cache, when the total amount of memory exceeds the amount of installed memory or occupies a large amount of installed memory, consider using this feature.

This feature adds the action of discarding the meta cache that has been held permanently. If you attempt to refer to a destroyed meta cache again, the meta cache is recreated, so using this feature will result in poor performance compared to not using it.

Therefore, read the following to understand how to discard a meta cache.

- [13.1.3 Cache Removal when Local Meta Cache Limit is Enabled](#)
- [13.1.4 Performance Impact and Parameter Tuning of the Local Meta Cache Limit Feature](#)
- [Parameters for the Local Meta Cache Limit feature](#)

How to set the upper limit with these considerations is described in detail in the estimation formula in “[Appendix A Parameters](#)”.

### 13.1.2 How to Set Parameters for the Local Meta Cache Limit Feature

To enable the Local Meta Cache Limit feature, set the `pgx_catalog_cache_max_size` and `pgx_relation_cache_max_size` parameters.

Parameter Name	Description
<code>pgx_catalog_cache_max_size</code>	Specify the maximum amount of memory that the backend process should use as the catalog cache. You can enable catalog cache removal by setting it to 8 KB or more. When it is set to 0 (default value), the catalog cache removal is disabled.
<code>pgx_relation_cache_max_size</code>	Specify the maximum amount of memory that the backend process should use as the relation cache. You can enable relation cache removal by setting it to 8 KB or more. When it is set to 0 (default value), the relation cache removal is disabled.



#### Example

Here is an example `postgresql.conf` configuration:

```
pgx_catalog_cache_max_size = 1MB
pgx_relation_cache_max_size = 1MB
```

### 13.1.3 Cache Removal when Local Meta Cache Limit is Enabled

When this feature is enabled, the caching strategy is to keep the cache as long as possible within the specified upper limit. If holding a new cache exceeds the limit, consider locality of reference and remove the cache from the one with the longest unreferenced time.

However, because the cache used by active transactions cannot be removed, if a transaction uses a large number of caches, the cache may be held above the limit. In this case, remove the all caches at the end of the transaction. This is necessary to free up memory.



In PostgreSQL, in order to acquire memory at high speed, a memory block of a certain size is acquired from the OS, and a small memory is cut out from the block and used. The memory for the metacache is cut out in the same way. Therefore, it is possible to return the memory block to the OS by destroying all the meta caches scattered throughout the memory block. When this happens, the next SQL execution will be slowed down due to the re-creation of the meta cache. Therefore, upper limit of feature should be set to a value larger than the size of the meta cache used by at least one transaction.

When the size of the meta cache exceeds the upper limit, the following message is output:

```
WARNING: could not reduce Cat/RelCacheMemoryContext size to AA kilobytes, reduced to BB kilobytes
HINT: consider increasing the configuration parameter pgx_catalog/relation_cache_max_size
```

(**AA**: Upper limit, **BB**: Amount of memory actually used)

CatCacheMemoryContext and RelCacheMemoryContext are memory areas for storing the catalog cache and relation cache, respectively. If this message is output, consider increasing the upper limit.

If the memory consumption by the backend process exceeds the allowable value by increasing the upper limit, reconsider the SQL to be executed, such as reducing the number of tables accessed in one transaction, or add memory adjust to the amount of memory used.

### 13.1.4 Performance Impact and Parameter Tuning of the Local Meta Cache Limit Feature

By observing how much meta cache regeneration is taking place, you can determine if the low upper limit is the cause of the failure to achieve the desired performance.

From the message below, calculate the cache hit ratio as follows:

```
Cache hit ratio = Number of cache hits ÷ Number of times the cache was searched
```

If the cache hit ratio is 80% or higher, this feature will not be the main factor that impedes performance. If not, raise the upper limit and see if performance can reach the goal. In doing so, first try to shift the focus of allocation to the relations cache. This is because when executing SQL, the relation cache generated based on the catalog cache is mainly referenced, so it is advantageous to leave a large amount of relation cache.

```
Catalog cache:catalog cache hit stats: search XX, hits YY
Relation cache:relation cache hit stats: search XX, hits YY
```

(**XX**: Number of times the cache was searched, **YY**: Number of cache hits)

This message is printed when the transaction ends. However, if you output the message frequently, the performance will be degraded by itself, so you can adjust the output interval with the following parameters.

Parameter Name	Description
pgx_cache_hit_log_interval	<p>When the transaction ends, if the time set in this parameter has elapsed since the previous message was output, the message is output.</p> <p>If set to 0, a message will be output each time the transaction ends. Setting -1 disables the output. The default value is 10min.</p> <p>Even if pgx_catalog_cache_max_size and pgx_relation_cache_max_size are disabled, the message output of the corresponding cache will be invalid.</p> <p>Immediately after connecting to the server, a small transaction occurs before the request from the user application, such as for user authentication. Since it is meaningless to know the hit ratio for these, a message is output at the end of the transaction that started after the time set in this parameter has elapsed after connecting to the server.</p> <p>For the same reason, setting a small value such as 0 may result in a message being printed at the end of such a small transaction.</p> <p>You can check which transaction the message corresponds to from the information output at the beginning. This information depends on the setting of the parameter log_line_prefix.</p>



## Example

Here is an example postgresql.conf configuration:

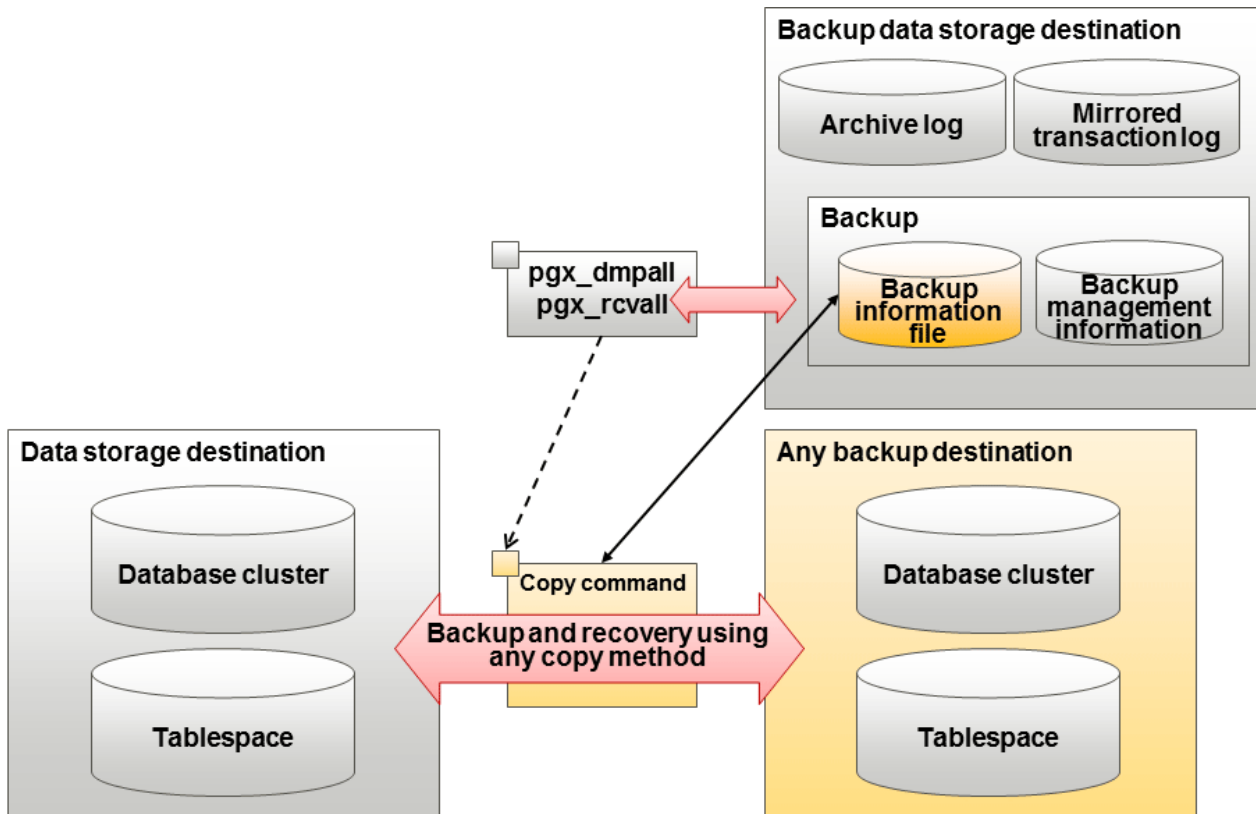
```
pgx_cache_hit_log_interval= 30min
```

# Chapter 14 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.



## Point

- 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.

## Note


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.

## 14.1 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.




#### Note

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.



### Information

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_rcvall` 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.

## Note

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.

## Point

By referring to the mode assigned to the copy command as an argument, backup and recovery can be implemented using a single copy command.

## Example

### Using a bash script

```
case $1 in
  prepare)
    processingRequiredForPrepareMode
    ;;
  backup)
    processingRequiredForBackupMode
    ;;
  finalize)
    processingRequiredForFinalizeMode
    ;;
  restore)
    processingRequiredForRestoreMode
    ;;
esac
```

## 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 and SLES 15.

## 14.2 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 "[14.4 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.



#### Example

```
$ pgx_dmpall -D /database/inst1 -Y '/database/command/backup.sh'  
--exclude-copy-cluster -P '/database/command/tablespace_list.txt'
```



#### Point

- 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.



#### See

- 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.

### Example

```
$ pgx_rcvall -l -D /database/inst1
Date                Status   Dir                Resources backed up by the copy command
2020-05-01 13:30:40 COMPLETE /backup/inst1/2020-05-01_13-30-40 pg_data,dbspace,indexspace
```

## 14.3 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.

### Example

```
$ 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 "14.2 Backup Using the Copy Command".

The example below recover only some tablespaces, but not the database cluster, using the copy command.

### Example

```
$ pgx_rcvall -D /database/inst1 -B /backup/inst1 -Y '/database/command/recovery.sh'
```

### Point

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.

### See

Refer to "pgx\_rcvall" in the Reference for information on the command.



## 14.4 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.

### 14.4.1 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.

#### Point

.....  
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

### Example

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 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 <code>pgx_dmpall</code> command continues processing.
Other than 0	Abnormal end The <code>pgx_dmpall</code> 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.

### Note

- 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.

## 14.4.2 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.

## Note

- 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 15 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



## Note

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 "[15.1 Recovering from Disk Failure \(Hardware\)](#)", and take actions accordingly.

- If the disk is not defective

Refer to "[15.14 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.



## See

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

$$\text{Recovery time} = (\text{usageByTheDataStorageDestination} + \text{usageByTheTransactionLogStorageDestination}) / \text{diskWritePerformance} \times 1.5$$

- *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

$$\text{Recovery time} = \text{usageByTheBackupDataStorageDestination} / \text{diskWritePerformance} \times 1.5$$

- *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

## 15.1 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:

- [15.1.1 Using WebAdmin](#)
- [15.1.2 Using Server Command](#)



### Point

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.

### 15.1.1 Using WebAdmin

---

Recover the database cluster by following the appropriate recovery procedure below for the disk where the failure occurred.



### Note

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.

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 "[15.10.3 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.

## 15.1.2 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 "[15.11 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
```



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 up
```

### 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
(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 "[15.11 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.



#### 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 `pgx_resume_wal_multiplexing`.

## 15.2 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:

- [15.2.1 Using WebAdmin](#)
- [15.2.2 Using the `pgx\_rcvall` Command](#)



#### Note

- 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.

## 15.2.1 Using WebAdmin

---

If using WebAdmin, recover the data to the point immediately prior to data corruption by using the backup data.

Refer to "[15.1.1 Using WebAdmin](#)" for details.

## 15.2.2 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 re-execute 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 "[15.11 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
2020-05-20 10:00:00 COMPLETE          /backup/inst1/2020-05-20_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, 2020 10:00:00" is specified as the recovery time.

```
> pgx_rcvall -D /database/inst1 -B /backup/inst1 -e '2020-05-20 10:00:00'
```

## Note

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 up
```

### 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.

### 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.

## 15.3 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:

- [15.3.1 Using WebAdmin](#)
- [15.3.2 Using the `pgx\_rcvall` Command](#)

## Note

- 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.

### 15.3.1 Using WebAdmin

---

You can use WebAdmin to recover data to a backup point.

## Note

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.

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.

## 15.3.2 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 "[15.11 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
2020-05-01 10:00:00 COMPLETE        /backup/inst1/2020-05-01_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\_20200503\_1".

```
> pgx_rcvall -D /database/inst1 -B /backup/inst1 -n batch_20200503_1
```

### Note

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 up
```

#### 6. Start the instance

Start the instance.

Refer to "[2.1.2 Using Server Commands](#)" for information on how to start an instance.

#### 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.

### See

Refer to "pgx\_rcvall" in the Reference for information on the `pgx_rcvall` command.

## 15.4 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 "15.4.1 When using the view (pg\_stat\_activity)")
- ps command (refer to "15.4.2 Using the ps Command")
- pgAdmin (refer to "15.4.3 Using pgAdmin")

Use the system management function (pg\_terminate\_backend) to disconnect connections.

## 15.4.1 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 (14.0)
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.

```
select pid,username,application_name,client_hostname,pg_terminate_backend(pid) from
pg_stat_activity where state='idle in transaction' and current_timestamp > cast(query_start +
interval '60 minutes' as timestamp);
-[ RECORD 1 ]-----+-----
pid                | 4684
username           | fsepuser
application_name   | apl1
client_addr        | 192.11.11.1
pg_terminate_backend | t
```



### See

- 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.

## 15.4.2 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      1 ?          /opt/fsepv<x>server64/bin/postgres -D /disk01/data
fsepuser 20518 20517 ?          \_ postgres: logger
fsepuser 20520 20517 ?          \_ postgres: checkpointer
fsepuser 20521 20517 ?          \_ postgres: background writer
fsepuser 20522 20517 ?          \_ postgres: walwriter
fsepuser 20523 20517 ?          \_ postgres: autovacuum launcher
```



```
fsepuser 20524 20517 ?      \_ postgres: archiver
fsepuser 20525 20517 ?      \_ postgres: stats collector
fsepuser 18673 20517 ?      \_ postgres: fsepuser postgres 192.168.100.1(49448) idle
fsepuser 16643 20517 ?      \_ postgres: fsepuser db01 192.168.100.11(49449) UPDATE waiting
fsepuser 16644 20517 ?      \_ postgres: fsepuser db01 192.168.100.12(49450) idle in transaction
```

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
(1 row)
```



See

- 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.


### 15.4.3 Using pgAdmin


---

If using pgAdmin, follow the procedure below to close connections.

1. In the [Browser] pane, click the relevant database server.

- 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  to close the session.



## 15.5 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



See

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.

## 15.6 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:

- [15.6.1 Using a Tablespace](#)

- [15.6.2 Replacing the Disk with a Larger Capacity Disk](#)

## 15.6.1 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.



See

Refer to "SQL Commands" under "Reference" in the PostgreSQL Documentation for information on the CREATE TABLESPACE command and ALTER TABLE command.

## 15.6.2 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:

- [15.6.2.1 Using WebAdmin](#)
- [15.6.2.2 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.



Note

- 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.

### 15.6.2.1 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 ["15.1.1 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.

## 15.6.2.2 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 ["15.11 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
```



## Note

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 up
```



## 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.

### 10. Resume applications

Resume applications that are using the database.

### 11. Restore files

Restore the files backed up in step 1.

## 15.7 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:

- [15.7.1 Temporarily Saving Backup Data](#)
- [15.7.2 Replacing the Disk with a Larger Capacity Disk](#)

### 15.7.1 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:

- [15.7.1.1 Using WebAdmin](#)
- [15.7.1.2 Using Server Commands](#)

## 15.7.1.1 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
```

## 15.7.1.2 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.

```
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 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/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 `/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 "[15.11 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.

## 15.7.2 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:

- [15.7.2.1 Using WebAdmin](#)
- [15.7.2.2 Using Server Commands](#)

## Note

Before replacing the disk, stop applications that are using the database.

### 15.7.2.1 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 "[15.1.1 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
```

## 15.7.2.2 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  
(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)
```

### 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 "[15.11 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
```



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`.

## 15.8 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.

## 15.8.1 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:

- [15.8.1.1 Using WebAdmin](#)
- [15.8.1.2 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.



### Note

- 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.

### 15.8.1.1 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 "[15.1.1 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.

# 15.8.1.2 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 "[15.11 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
```

#### Note

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 up
```



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.

10. Resume applications

Resume applications that are using the database.

11. Restore files

Restore the files backed up in step 1.

## 15.9 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



Refer to "Setup" in the Installation and Setup Guide for Server for information on how to create an instance and build the runtime environment.

## 15.10 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.

### 15.10.1 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



See

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

## 15.10.2 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 "[15.14.2 Errors Caused by Power Failure or Mounting Issues](#)", and take actions accordingly.

## 15.10.3 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:

- [15.10.3.1 Using WebAdmin](#)
- [15.10.3.2 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.

### 15.10.3.1 Using WebAdmin

Follow the procedure below to perform recovery.

1. 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 "[15.2.1 Using WebAdmin](#)" for details.

### 15.10.3.2 Using Server Commands

Follow the procedure below to recover the database.

1. 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 "[15.2.2 Using the pgx\\_rcvall Command](#)" for details.

## 15.11 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:

- [15.11.1 Using WebAdmin](#)
- [15.11.2 Using Server Commands](#)

### 15.11.1 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.

### 15.11.2 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.

#### 15.11.2.1 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 "[15.11.2.2 Stopping the Instance Using the Immediate Mode](#)" or "[15.11.2.3 Forcibly Stopping the Server Process](#)".



#### Example

```
> pg_ctl stop -D /database/inst1 -m fast
```

#### 15.11.2.2 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 "[15.11.2.3 Forcibly Stopping the Server Process](#)".



#### Example

```
> pg_ctl stop -D /database/inst1 -m immediate
```

#### 15.11.2.3 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
fsepuser 20520 20517 ?          \_ postgres: checkpointer
fsepuser 20521 20517 ?          \_ postgres: background writer
fsepuser 20522 20517 ?          \_ postgres: walwriter
fsepuser 20523 20517 ?          \_ postgres: autovacuum launcher
fsepuser 20524 20517 ?          \_ postgres: archiver
fsepuser 20525 20517 ?          \_ postgres: stats collector
```

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
```

## 15.12 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.



See

Refer to "Deleting Instances" in the Installation and Setup Guide for details on how to delete an instance.

## 15.13 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

```
postgres=# select * from pg_prepared_xacts;
 transaction | gid          | prepared | owner  | database
-----+-----+-----+-----+-----
 2103 | 374cc221-f6dc-4b73-9d62-d4fec9b430cd | 2020-05-06 16:28:48.471+08 | postgres |
postgres (1 row)
```

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
```

---

## 15.14 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.

- [15.14.1 Network Error with an External Disk](#)
- [15.14.2 Errors Caused by Power Failure or Mounting Issues](#)

---

### 15.14.1 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.

---

### 15.14.2 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.

---

## 15.15 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.

---

### 15.15.1 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.



### Note

Critical errors encountered during anomaly resolution will be displayed, however, rollback of the instance to its previous state is not supported.

## 15.15.2 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.

## Information

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.

### - 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".

## Example

```
search_path = '$user', public, oracle, pg_catalog'
```

## Information

- 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.

- 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 feature

- 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.

```
Size of GMC area
= Max(10MB,
      (All user table x 0.4 KB
       + All user Indexes x 0.3 KB
       + All user columns x 0.8 KB) x 1.5 (*1) )
```

\*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.

#### Parameters for the Local Meta Cache Limit feature

- pgx\_catalog\_cache\_max\_size(numerical value)

Specifies the maximum amount of memory that the backend process should use as the catalog cache.

You can enable catalog cache deletion by setting it to 8 KB or more.

A setting of 0 disables the catalog cache removal. The default is 0.

If no units are specified, they are treated as KB.

- Minimum value: 8KB
- Maximum value: Maximum value that can be expressed as a 4-byte signed integer

When calculating the parameter settings, the factors that determine the cache size are calculated as the number of tables, the number of indexes, and the number of columns. What is kept as a catalog cache or relation cache also includes objects such as databases, roles, or procedures, but these are small compared to the above factors and do not need to be factored into them. It also includes a calculation method for pgx\_relation\_cache\_max\_size because the given memory is distributed between the catalog cache and the relation cache.



#### Note

The calculation method here assumes that all backends have similar access and that the transaction also has access to a similar number of resources. If you have a small number of singular backends or transactions, consider excluding them as errors.

1. Determine how much memory a backend process can use. Decide by subtracting the memory size required by the entire system such as the database cache from the installed memory and dividing the rest by the number of connections.
2. For best performance, use the following formula to calculate the total memory size of the catalog cache when the backend holds the catalog cache for all resources accessed during its lifetime.  
The amount of memory varies depending on whether Global Meta Cache is enabled or disabled. Enabling Global Meta Cache reduces the amount of memory required because most of the cache is located on shared memory.

When Global Meta Cache is enabled:

```
(Number of tables to access + Number of indexes to access + Number of columns to access)
× 0.1KB × 1.5 (*1)
```

When Global Meta Cache is disabled:

```
{ Number of tables to access × 0.5KB(pg_class tuple size)
```

```
+ Number of indexes to access × 0.5KB(pg_index tuple size)
+ Number of columns to access × 1.0KB(pg_statistic tuple size)} × 1.5 (*1)
```

\*1) Safety Factor (1.5)

The system catalog contains columns with variable-length types. For example, the tuple size in pg\_class is a constant value multiplied by the number of tables, while relname in pg\_class is variable length data.

It is not practical to calculate every definition in detail, so we added 50% to the above formula.

3. In the same way as in 2., calculate the relation cache using the following formula.

```
(1.4KB × Number of tables to access + 2.4KB × Number of indexes to access) × 1.5 (*1)
```

\*1) Safety Factor (1.5)

The relation cache is structured to facilitate the use of table and index definitions, holds pointers to various objects, and is sized to include them. It is variable length because the type of object allocated by the table definition and its size change. Since it is not realistic to calculate for all definitions, 50% is added.

4. If the value of 1. the value of 2. + the value of 3., the backend process can keep all caches to the extent allowed, so there is no need to limit the caches. If you want to cap for safety, set the value of 2. to pgx\_catalog\_cache\_max\_size and the value of 3. to pgx\_relation\_cache\_max\_size.
5. If the value of 1. < the value of 2. + the value of 3. then you need to limit the cache. However, this parameter does not limit the size of the cache used by a transaction. Therefore, take the following steps.
6. Calculate the catalog cache used by a transaction using the formula in 2.
7. Calculate the relation cache used by a transaction using the formula in 3.
8. If the value of 1. < the value of 6. + the value of 7., then the value of 1. needs to be increased. In other words, in some cases, it may be necessary to increase the installed memory or reduce the number of connections.
9. If the value of 1. the value of 6. + the value of 7., the condition of 1. can be satisfied by limiting the cache with this parameter. Divide the value of 1. by the ratio of 2. and 3. and set it as a parameter. Set the value distributed to 2. to pgx\_catalog\_cache\_max\_size and the value distributed to 3. to pgx\_relation\_cache\_max\_size.
10. The value calculated in 9. is a provisional value. If you cannot meet your target performance, first try to shift the focus of allocation to the relation cache. This is because when executing SQL, the relation cache generated based on the catalog cache is mainly referenced, so it is advantageous to leave a large amount of relation cache. If the performance is still not satisfied, adjust the parameters by referring to "[13.1.4 Performance Impact and Parameter Tuning of the Local Meta Cache Limit Feature](#)".

## Note

Be careful when partitioning the table.

The cached definition changes depending on whether the parent table is specified in the SQL statement or the child table is specified. In particular, note that if you specify a parent table, the definitions of all child tables are cached. This is because when you specify a parent table in an SQL statement, you need to know the definitions of all the child tables in order to determine which child table will contain the desired data. Note that the column information of the parent table is not cached.

When specifying the parent table:

```
Number of tables to access = Number of parent tables to access + Number of defined child tables
Number of columns = Number of defined columns × number of defined child tables
```

When specifying the child table directly:

```
Number of tables to access = Number of child tables actually accessed
Number of columns = Number of defined columns × number of child tables actually accessed
```

Example)

Suppose the parent table T (1 index, 3 columns) is split from child tables T1 to T5 (1 index, 3 columns, respectively). If the parent table T is specified in SQL, when the child tables that contain the data to be queried are limited to T1 and T2, and when accessing the data using the indexes defined by T1 and T2, calculate as follows.

```
Number of tables = 1(parent table) + 5(child table) = 6
Number of indexes = 2 (index to access)
Number of columns = 3 (number of columns) x 5 (child table) = 15
```

If you specify child tables T1 and T2 in SQL and use the indexes defined on T1 and T2 when accessing data, the calculation is as follows.

```
Number of tables = 2(child table)
Number of indexes = 2 (index to access)
Number of columns = 3 (number of columns) x 2 (child table) = 6
```



- `pgx_relation_cache_max_size`(numerical value)

Specifies the maximum amount of memory that the backend process should use as the relation cache.

You can enable catalog cache deletion by setting it to 8 KB or more.

A setting of 0 disables the relation cache removal. The default is 0.

If no units are specified, they are treated as KB.

- Minimum value: 8KB
- Maximum value: Maximum value that can be expressed as a 4-byte signed integer

For the calculation method for parameter setting, refer to the calculation method of `pgx_catalog_cache_max_size`.

- `pgx_cache_hit_log_interval`(numerical value)

Specifies the time interval to output a message indicating the cache reference status for each backend process.

When the transaction ends, if the time set in this parameter has elapsed since the previous message was output, the message is output.

If set to 0, a message will be output each time the transaction ends.

Setting -1 disables the output. The default value is 10min.

If no units are specified, they are treated as ms.

Even if `pgx_catalog_cache_max_size` and `pgx_relation_cache_max_size` are disabled, the message output of the corresponding cache will be invalid.

Immediately after connecting to the server, a small transaction occurs before the request from the user application, such as for user authentication. Since it is meaningless to know the hit rate for these, a message will be output at the end of the transaction that started after the time set in this parameter has elapsed after connecting to the server.

For the same reason, setting a small value such as 0 may result in a message being printed at the end of such a small transaction.

You can check which transaction the message corresponds to from the information output at the beginning.

This information depends on the setting of the parameter `log_line_prefix`.

- Minimum value: 0
- Maximum value: 2147483647ms



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
<code>pgx_pause_wal_multiplexing()</code>	void	Stops WAL multiplexing
<code>pgx_resume_wal_multiplexing()</code>	void	Resumes WAL multiplexing
<code>pgx_is_wal_multiplexing_paused()</code>	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
<code>pgx_open_keystore(<i>passphrase</i>)</code>	void	Opens the keystore
<code>pgx_set_master_key(<i>passphrase</i>)</code>	void	Sets the master encryption key
<code>pgx_set_keystore_passphrase(<i>oldPassphrase</i>, <i>newPassphrase</i>)</code>	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
<a href="#">pgx_alter_confidential_policy</a>	boolean	Changes masking policies
<a href="#">pgx_create_confidential_policy</a>	boolean	Creates masking policies
<a href="#">pgx_drop_confidential_policy</a>	boolean	Deletes masking policies
<a href="#">pgx_enable_confidential_policy</a>	boolean	Enables or disables masking policies
<a href="#">pgx_update_confidential_values</a>	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

```
common_arg:
[schema_name      := 'schemaName', ]


```

- 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

```
pgx_alter_confidential_policy(  
  commonArg,  
  action      := 'DROP_COLUMN',  
  column_name := 'colName'  
)
```

- Change the masking condition

```
pgx_alter_confidential_policy(  
  commonArg,  
  action      := 'MODIFY_EXPRESSION',  
  expression  := 'expr'  
)
```

- Change the content of a masking policy set for a masking target

```
pgx_alter_confidential_policy(  
  commonArg,  
  action      := 'MODIFY_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']
```

- 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_COLUMN'
	column_name	varchar(63)	Masking target name	Mandatory
	function_type	varchar(63)	Masking type - 'FULL': Full masking - 'PARTIAL': Partial masking - 'REGEXP': Regular expression masking	'FULL'
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

- 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 masking - 'PARTIAL': Partial masking - 'REGEXP': Regular expression masking	'FULL'
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

- 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_COLUMN	MODIFY_EXPRESSION	MODIFY_COLUMN			SET_POLICY_DESCRIPTION	SET_COLUMN_DESCRIPTION
	Full masking	Partial masking	Regular expression masking			Full masking	Partial masking	Regular expression masking		
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_COLUMN	MODIFY_EXPRESSION	MODIFY_COLUMN			SET_POLICY_DESCRIPTION	SET_COLUMN_DESCRIPTION
	Full masking	Partial masking	Regular expression masking			Full masking	Partial masking	Regular expression masking		
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

```
postgres=# select pgx_alter_confidential_policy(table_name := 't1', policy_name := 'p1', action :=
'ADD_COLUMN', column_name := 'c2', function_type := 'PARTIAL', function_parameters := 'VVVFVVVFVVVV,
VVV-VVVV-VVVV, *, 4, 11');
pgx_alter_confidential_policy
-----
t
(1 row)
```

### 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
(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

```
postgres=# select pgx_alter_confidential_policy(table_name := 't1', policy_name := 'p1', action :=
'SET_POLICY_DESCRIPTION', policy_description := 'this policy is an example. ');
pgx_alter_confidential_policy
-----
t
(1 row)
```

### 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
<code>ADD_COLUMN</code>	Adds a masking target to a masking policy.
<code>DROP_COLUMN</code>	Deletes a masking target from a masking policy.
<code>MODIFY_EXPRESSION</code>	Changes expression.
<code>MODIFY_COLUMN</code>	Changes the content of a masking policy set for a masking target.
<code>SET_POLICY_DESCRIPTION</code>	Changes <code>policy_description</code> .
<code>SET_COLUMN_DESCRIPTION</code>	Changes <code>column_description</code> .

- 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`



- 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',]
table_name       := 'tableName',
policy_name      := 'policyName',
expression       := 'expr'
[, enable        := 'policyStatus']
[, policy_description := 'policyDesc']
[, column_name    := 'colName'
  [, function_type := 'FULL'] |
  [, function_type := 'PARTIAL', partialOpt] |
  [, function_type := 'REGEXP', regexpOpt]
[, column_description := 'colDesc']
])
```

```
partialOpt:
function_parameters := 'maskingFmt'
```

```
regexpOpt:
regexp_pattern      := 'regexpPattern',
regexp_replacement  := 'regexpReplacementChar',
[, regexp_flags     := 'regexpFlags']
```

### Argument

Details are as follows.

Masking type for which an argument can be specified	Argument	Data type	Description	Default value
All	<code>schema_name</code>	<code>varchar(63)</code>	Schema name of table for which the masking policy is created	'public'
	<code>table_name</code>	<code>varchar(63)</code>	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': Enabled - 'f': Disabled	't'
	policy_description	varchar(1024)	Masking policy description	NULL
	column_name	varchar(63)	Masking target name	NULL
	function_type	varchar(63)	Masking type - 'FULL': Full masking - 'PARTIAL': Partial masking - 'REGEXP': Regular expression masking	'FULL'
	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

```
postgres=# select pgx_create_confidential_policy(table_name := 't1', policy_name := 'p1',
expression := 'l=1');
pgx_create_confidential_policy
-----
t
(1 row)
```

## 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 := 'l=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

```
postgres=# select pgx_create_confidential_policy( table_name := 't1', policy_name := 'p1',
expression := 'l=1', column_name := 'c2', function_type := 'PARTIAL', function_parameters :=
'VVVFVVVFVVVFVVV, VVV-VVVV-VVVV, *, 4, 11');
```

```
pgx_create_confidential_policy
-----
t
(1 row)
```

## 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 := 'l=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`

### Note

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', ]


```

### Argument

Details are as follows.

Argument	Data type	Description	Default value
<code>schema_name</code>	<code>varchar(63)</code>	Schema name of table for which a masking policy is deleted	'public'
<code>table_name</code>	<code>varchar(63)</code>	Name of table for which a masking policy is deleted	Mandatory
<code>policy_name</code>	<code>varchar(63)</code>	Masking policy name	Mandatory

Details about whether arguments can be omitted are as follows.

Argument	Mandatory or optional
<code>schema_name</code>	Y
<code>table_name</code>	N
<code>policy_name</code>	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.



### Note

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.

## B.3.4 `pgx_enable_confidential_policy`

### Description

Enables or disables masking policies

### Format

```
pgx_enable_confidential_policy(  
[schema_name      := 'schemaName', ]  
table_name        := 'tableName',  
policy_name       := 'policyName',  
enable            := 'policyStatus'  
)
```

### Argument

Details are as follows.

Argument	Data type	Description	Default value
<code>schema_name</code>	<code>varchar(63)</code>	Schema name of table for which a masking policy is enabled or disabled	'public'
<code>table_name</code>	<code>varchar(63)</code>	Name of table for which a masking policy is enabled or disabled	Mandatory
<code>policy_name</code>	<code>varchar(63)</code>	Masking policy name	Mandatory
<code>enable</code>	<code>boolean</code>	Masking policy status - 't': Enabled	Mandatory



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.



See

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
<code>pgx_prewarm_vci(vci_index regclass)</code>	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.



See

Refer to "System Views" under "Internals" in the PostgreSQL Documentation for information on other system views.

### C.1 pgx\_tablespaces

The `pgx_tablespaces` view provides information related to the encryption of tablespaces.

Table C.1 `pgx_tablespaces` view

Column	Type	References	Description
<code>spctablespace</code>	<code>oid</code>	<code>pg_tablespace.oid</code>	Tablespace OID
<code>spcencalgo</code>	<code>text</code>		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.2 `pgx_stat_lwlock` view

Column	Type	Description
<code>lwlock_name</code>	<code>name</code>	Name of the lightweight lock
<code>total_waits</code>	<code>bigint</code>	Number of waits caused by the lightweight lock
<code>total_wait_time</code>	<code>double precision</code>	Number of milliseconds spent in waits caused by the lightweight lock
<code>stats_reset</code>	<code>timestamp with timezone</code>	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.3 `pgx_stat_latch` view

Column	Type	Description
<code>latch_name</code>	<code>name</code>	Name of the latch
<code>total_waits</code>	<code>bigint</code>	Number of waits caused a wait
<code>total_wait_time</code>	<code>double precision</code>	Number of milliseconds spent in waits caused by the latch
<code>stats_reset</code>	<code>timestamp with timezone</code>	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.4 `pgx_stat_walwriter` view

Column	Type	Description
<code>dirty_writes</code>	<code>bigint</code>	Number of times old WAL buffers were written to the disk because the WAL buffer was full when WAL records were added
<code>writes</code>	<code>bigint</code>	Number of WAL writes
<code>write_blocks</code>	<code>bigint</code>	Number of WAL write blocks
<code>total_write_time</code>	<code>double precision</code>	Number of milliseconds spent on WAL writing
<code>stats_reset</code>	<code>timestamp with timezone</code>	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.5 `pgx_stat_sql` view

Column	Type	Description
<code>selects</code>	<code>bigint</code>	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 <code>heartbeat_interval</code> of the server definition file (milliseconds).
<code>inserts</code>	<code>bigint</code>	Number of INSERT statements executed
<code>deletes</code>	<code>bigint</code>	Number of DELETE statements executed
<code>updates</code>	<code>bigint</code>	Number of UPDATE statements executed
<code>selects_with_parallelism</code>	<code>bigint</code>	Number of times parallel scan was used in SELECT statements
<code>inserts_with_parallelism</code>	<code>bigint</code>	Not used
<code>deletes_with_parallelism</code>	<code>bigint</code>	Not used
<code>updates_with_parallelism</code>	<code>bigint</code>	Not used
<code>copies_with_parallelism</code>	<code>bigint</code>	Not used
<code>declares</code>	<code>bigint</code>	Number of DECLARE statements executed (number of cursor OPENS)
<code>fetches</code>	<code>bigint</code>	Number of FETCH statements executed
<code>checkpoints</code>	<code>bigint</code>	Number of CHECKPOINT statements executed
<code>clusters</code>	<code>bigint</code>	Number of CLUSTER statements executed
<code>copies</code>	<code>bigint</code>	Number of COPY statements executed
<code>reindexes</code>	<code>bigint</code>	Number of REINDEX statements executed
<code>truncates</code>	<code>bigint</code>	Number of TRUNCATE statements executed
<code>locks</code>	<code>bigint</code>	Number of times a lock occurred
<code>stats_reset</code>	<code>timestamp with timezone</code>	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.6 `pgx_stat_gmc` view

Column	Type	Description
<code>searches</code>	<code>bigint</code>	Number of times the cache table is searched.
<code>hits</code>	<code>bigint</code>	Number of times the cache table is hit.
<code>size</code>	<code>bigint</code>	The current amount of memory (bytes) used in the GMC area.
<code>stats_reset</code>	<code>timestamp with timezone</code>	Last time these statistics were reset.

## C.7 pgx\_stat\_progress\_loader

---

The `pgx_stat_progress_loader` view provides overall progress information for `pgx_loader` command.

The `pgx_stat_progress_loader` view displays the sum of the progress information of the back-end processes and the number of parallel worker processes when `pgx_loader` runs.

Table C.7 `pgx_stat_progress_loader` view

Column	Type	Description
<code>pid</code>	<code>integer</code>	Process ID of the backend.
<code>datid</code>	<code>Oid</code>	Oid of the database to connect to.
<code>datname</code>	<code>name</code>	Name of the database to connect to.
<code>relid</code>	<code>Oid</code>	Oid of the table to load.
<code>command</code>	<code>text</code>	Command executes the load process. (Always "COPY FROM" for <code>pgx_loader</code> )
<code>type</code>	<code>text</code>	Type of data source for the load operation.
<code>bytes_processed</code>	<code>bigint</code>	Size of the data at the end of the load. (Backend and worker totals)
<code>bytes_total</code>	<code>bigint</code>	Size of the data to load. (Backend and worker totals)
<code>tuples_processed</code>	<code>bigint</code>	Number of rows that have completed loading. (Backend and worker totals)
<code>tuples_excluded</code>	<code>bigint</code>	Number of rows skipped during the load process. (Backend and worker totals)

# Appendix D Tables Used by Data Masking

This appendix explains tables used by the data masking feature.

## Note

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	Type	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

```
postgres=# select * from pgx_confidential_columns;
 schema_name | table_name | policy_name | column_name | function_type |
function_parameters          | regexp_pattern | regexp_replacement | regexp_flags |
column_description
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+-----+-----+-----+
public       | t1        | p1         | c1          | FULL         |
|
public       | t1        | p1         | c2          | PARTIAL      | VVVFVVVFVVVV, VVV-VVVV-VVVV,
*, 4, 11 |
(2 row)
```

## D.2 pgx\_confidential\_policies

This table provides information on masking policies.

Column	Type	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

```
postgres=# select * from pgx_confidential_policies;
 schema_name | table_name | policy_name | expression | enable | policy_description
-----+-----+-----+-----+-----+-----
 public      | t1         | p1          | 1=1        | t      |
(1 row)
```

## 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

```
postgres=# select * from pgx_confidential_values;
 number_value | char_value | varchar_value | date_value | ts_value
-----+-----+-----+-----+-----
          0 |           |              | 1970-01-01 | 1970-01-01 00:00:00
(1 row)
```



# 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	Type	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: <ul style="list-style-type: none"><li>- commit: The prepared transaction has been committed.</li><li>- rollback: The prepared transaction is in in-doubt state.</li></ul>
leader_pid	integer	Process ID of the backend process (leader 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).

### Note

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.



See

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.

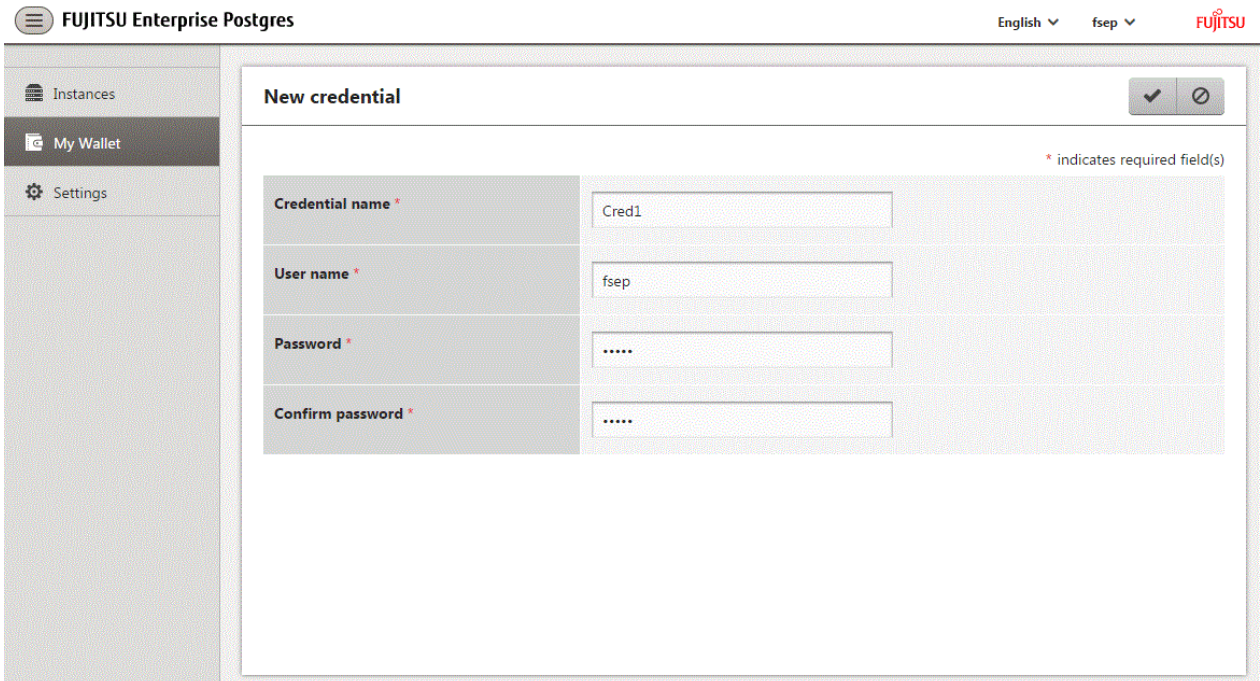
## Note

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.



The screenshot displays the 'New credential' page in the Fujitsu Enterprise Postgres WebAdmin interface. The page title is 'New credential'. On the left, a sidebar menu shows 'Instances', 'My Wallet' (selected), and 'Settings'. The main content area contains a form with the following fields:

Field Label	Value
Credential name *	Cred1
User name *	fsep
Password *	.....
Confirm password *	.....

A note at the top right of the form states: '\* indicates required field(s)'. The top navigation bar includes 'FUJITSU Enterprise Postgres', 'English', 'fsep', and the 'FUJITSU' logo.

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  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.

The screenshot shows the 'Create instance' form in the Fujitsu Enterprise Postgres console. The form is for a standalone configuration. It includes fields for Configuration type (Standalone configuration), Server product type (FUJITSU Enterprise Postgres .X), Location (Remote), Instance name (inst2), Instance port (27502), Host name (Standby/Host), Operating system credential (Cred1), Data storage path (/database/fsep/inst2/data), Backup (Enabled), Backup storage path (/database/fsep/inst2/backup), Transaction log path (/database/fsep/inst2/transactionlog), Encoding (UTF8), WAL file size (8), and Remote WebAdmin port (25515). The 'Operating system credential' dropdown is set to 'Cred1', which has automatically populated the 'User name' field with 'fsep' and the 'Password' field with a masked password.

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)
- ,
- \ (backslash)

# 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_rcvall` 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_acm1.sh.sample
```

- Advanced version

```
/installDir/share/copy_command.esf_acm2.sh.sample
```



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.

#### Note

- The samples cannot be used on SLES 12 and SLES 15.
- 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.

RHEL7,RHEL8 and SLES 12

Use FJQSS (Information Collection Tool) to collect data for initial investigation.

SLES 15

Use the `pgx_fjqssinf` command to collect data for initial investigation.



### See

- Refer to the FJQSS manual for information on how to use FJQSS.
- Refer to the Reference for information on the `pgx_fjqssinf` command.



### Note

- 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.

# Index

	[A]	
Actions in Response to Instance Startup Failure.....	119	Encrypting a Tablespace.....
All user data within the specified tablespace.....	25	Encrypting Existing Data.....
Approximate backup time.....	17	Encryption mechanisms.....
Approximate recovery time.....	92	Errors in More Than One Storage Disk.....
Automatically opening the keystore.....	36	
	[B]	
Backing Up and Recovering the Keystore.....	30	[F]
Backing Up and Restoring/Recovering the Database.....	32	Faster encryption and decryption based on hardware.....
Backup/Recovery Using the Copy Command.....	82	File system level backup and restore.....
Backup and recovery using the <code>pgx_dmpall</code> and <code>pgx_rcvall</code>		
commands.....	32	[H]
backup cycle.....	18	High-Speed Data Load.....
Backup data.....	25	
Backup operation.....	19	[I]
Backup operation (file backup).....	20	If failure occurred in the data storage disk or the transaction log
Backup status.....	19,20	storage disk.....
Backup using the backup information file.....	83	storage disk.....
Backup Using the Copy Command.....	86	If failure occurred on the backup data storage disk.....
<code>backup_destination</code> (string).....	126	If failure occurred on the data storage disk or the transaction log
Building and starting a standby server.....	36	storage directory.....
		Importing and Exporting the Database.....
		Installing and Operating the In-memory Feature.....
	[C]	[K]
Changing a Masking Policy.....	44	<code>keystore_location</code> (string).....
Changing the Keystore Passphrase.....	29	
Changing the Master Encryption Key.....	29	[L]
Changing the master encryption key and the passphrase.....	36	Logging in to WebAdmin.....
Checking an Encrypted Tablespace.....	28	log in.....
Checking backup status.....	87	
Checking the operating status of an instance.....	13,15	[M]
Collecting Failure Investigation Data .....	160	Managing the Keystore.....
Configuration of the Copy Command.....	82	Masking Condition.....
Configuration of the copy command for backup.....	84	Masking Format.....
Configuration of the copy command for recovery.....	85	Masking Policy.....
Confirming a Masking Policy.....	44	Masking Target.....
Continuous archiving and point-in-time recovery.....	33	Masking Type.....
Copy Command for Backup.....	88	Monitoring Database Activity.....
Copy Command for Recovery.....	90	
Copy Command Interface.....	88	[O]
Copy Command Samples that Use the Advanced Copy Feature of		Opening the Keystore.....
the ETERNUS Disk Array.....	158	Operating FUJITSU Enterprise Postgres.....
<code>core_contents</code> (string).....	126	
<code>core_directory</code> (string).....	126	[P]
Creating a Masking Policy.....	43	Parallel Query.....
Cyclic usage of the backup area.....	83	Performing backup.....
		Perform recovery.....
		Periodic Backup.....
		<code>pgx_global_metacache</code> (numerical value).....
		<code>pgx_stat_gmc</code> view.....
		<code>pgx_stat_latch</code> view.....
		<code>pgx_stat_lwlock</code> view.....
		<code>pgx_stat_progress_loader</code> view.....
		<code>pgx_stat_sql</code> view.....
		<code>pgx_stat_walwriter</code> view.....
		<code>pgx_tablespaces</code> .....
		<code>pgx_tablespaces</code> view.....
		Placement and automatic opening of the keystore file.....
		Placing the keystore file.....
	[D]	
Data Masking.....	38	
Data Types for Masking.....	46	
Deleting a Masking Policy.....	46	
Determining the backup area of the latest backup.....	87	
	[E]	
Enabling and Disabling a Masking Policy.....	45	
Enabling Automatic Opening of the Keystore.....	29	

Preparing for backup.....	86
[R]	
Recovery Using the Copy Command.....	87
reserve_buffer_ratio (numerical value).....	127
[S]	
Scope of encryption.....	25
search_path (string).....	126
Security-Related Notes.....	36
Security Notes.....	47
Setting a restore point.....	21
Setting the Master Encryption Key.....	26
Starting and Stopping the Web Server Feature of WebAdmin.....	153
Starting an instance.....	12,14
Starting pgAdmin.....	3,4
Startup URL for WebAdmin.....	3
Stopping an instance.....	12,14
Streaming replication support.....	26
Streaming Replication Using WebAdmin.....	56
Strong encryption algorithms.....	25
System Administration Functions.....	132
System Views.....	147
[T]	
tablespace_encryption_algorithm (string).....	126
Tables Used by Data Masking .....	150
Tips for Installing Built Applications.....	37
track_gmc (string).....	129
track_sql (string).....	127
track_waits (string).....	127
Transparent Data Encryption Control Functions.....	132
Two-layer encryption key and the keystore.....	25
[U]	
User environment.....	3
Using Server Commands.....	14
[V]	
vci.control_max_workers (numerical value).....	127
vci.cost_threshold (numeric).....	127
vci.enable (string).....	127
vci.log_query (string).....	128
vci.maintenance_work_mem (numerical value).....	128
vci.max_local_ros (numerical value).....	128
vci.max_parallel_degree (numerical value).....	128
vci.shared_work_mem (numerical value).....	128
[W]	
WAL and temporary files.....	25
WAL Mirroring Control Functions.....	132
WebAdmin Wallet.....	155

# FUJITSU Enterprise Postgres 14

## 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: February 2022
----------------------------

### **Copyright**

Copyright 2015-2022 FUJITSU LIMITED

# Contents

---

Chapter 1 Overview of Installation.....	1
1.1 Features that can be Installed.....	1
1.2 Installation Types.....	1
1.2.1 New Installation.....	1
1.2.2 Reinstallation.....	1
1.2.3 Multi-Version Installation.....	1
1.3 Installation Procedure.....	1
1.3.1 Installation in Interactive Mode.....	2
1.3.2 Installation in Silent Mode.....	2
1.4 Uninstallation.....	2
Chapter 2 Operating Environment.....	3
2.1 Required Operating System.....	3
2.2 Related Software.....	3
2.3 Excluded Software.....	4
2.4 Required Patches.....	5
2.5 Hardware Environment.....	5
2.6 Disk Space Required for Installation.....	5
2.7 Supported System Environment.....	5
2.7.1 TCP/IP Protocol.....	5
2.7.2 File System.....	5
2.8 PostgreSQL Version Used for FUJITSU Enterprise Postgres.....	6
2.9 Notes on Using Streaming Replication.....	6
Chapter 3 Installation.....	7
3.1 Pre-installation Tasks.....	8
3.2 Installation in Interactive Mode.....	9
3.3 Installation in Silent Mode.....	11
Chapter 4 Setup.....	13
4.1 Operating Method Types and Selection.....	13
4.2 Preparations for Setup.....	14
4.2.1 Creating an Instance Administrator.....	14
4.2.1.1 Security policy settings.....	15
4.2.2 Preparing Directories for Resource Deployment.....	15
4.2.3 Estimating Resources.....	19
4.2.4 Windows Firewall Settings.....	19
4.2.5 Preparing for Output to the Event Log.....	20
4.3 Creating Instances.....	21
4.3.1 Using WebAdmin.....	22
4.3.1.1 Logging in to WebAdmin.....	22
4.3.1.2 Creating an Instance.....	23
4.3.1.3 Changing Instance Settings.....	26
4.3.1.4 Importing Instances.....	28
4.3.2 Using the initdb Command.....	30
4.3.2.1 Creating an Instance.....	30
4.4 Configuring Remote Connections.....	32
4.4.1 When an Instance was Created with WebAdmin.....	33
4.4.2 When an Instance was Created with the initdb Command.....	33
4.4.3 Windows Firewall Settings.....	33
4.5 Other Settings.....	34
4.5.1 Error Log Settings.....	34
4.5.2 Configuring Automatic Start and Stop of an Instance.....	35
4.5.3 Settings when Using the Features Compatible with Oracle Databases.....	36
4.6 Integration with Message-Monitoring Software.....	36
4.7 Setting Up and Removing OSS.....	36

4.7.1 oracle_fdw.....	37
4.7.1.1 Setting Up oracle_fdw.....	37
4.7.1.2 Removing oracle_fdw.....	37
4.7.2 pg_hint_plan.....	38
4.7.2.1 Setting Up pg_hint_plan.....	38
4.7.2.2 Removing pg_hint_plan.....	38
4.7.3 pg_dbms_stats.....	39
4.7.3.1 Setting Up pg_dbms_stats.....	39
4.7.3.2 Removing pg_dbms_stats.....	39
4.8 Deleting Instances.....	40
4.8.1 Using WebAdmin.....	40
4.8.2 Using Server Commands.....	40
<b>Chapter 5 Uninstallation.....</b>	<b>42</b>
5.1 Uninstallation in Interactive Mode.....	42
5.2 Uninstallation in Silent Mode.....	44
<b>Appendix A Recommended WebAdmin Environments.....</b>	<b>48</b>
A.1 Recommended Browser Settings.....	48
A.2 How to Set Up the Pop-up Blocker.....	48
<b>Appendix B Setting Up and Removing WebAdmin.....</b>	<b>49</b>
B.1 Setting Up WebAdmin.....	49
B.1.1 Setting Up WebAdmin.....	49
B.1.2 Starting the Web Server Feature of WebAdmin.....	50
B.1.3 Stopping the Web Server Feature of WebAdmin.....	50
B.2 Removing WebAdmin.....	51
B.3 Using an External Repository for WebAdmin.....	51
B.4 Using the WebAdmin Auto-Refresh Feature.....	53
<b>Appendix C WebAdmin Disallow User Inputs Containing Hazardous Characters.....</b>	<b>54</b>
<b>Appendix D Configuring Parameters.....</b>	<b>55</b>
<b>Appendix E Estimating Database Disk Space Requirements.....</b>	<b>59</b>
E.1 Estimating Table Size Requirements.....	59
E.2 Estimating Index Size Requirements.....	60
E.3 Sizes of Data Types.....	61
E.3.1 Sizes of Fixed-Length Data Types.....	61
E.3.2 Sizes of Variable-Length Data Types.....	62
E.3.3 Sizes of Array Data Types.....	62
E.3.4 Number of Bytes per Character.....	63
E.4 Estimating Transaction Log Space Requirements.....	63
E.5 Estimating Archive Log Space Requirements.....	63
E.6 Estimating Backup Disk Space Requirements.....	63
E.7 Estimating VCI Disk Space Requirements.....	63
<b>Appendix F Estimating Memory Requirements.....</b>	<b>65</b>
F.1 FUJITSU Enterprise Postgres Memory Requirements.....	65
F.2 Database Multiplexing Memory Requirements.....	66
F.3 VCI Memory Requirements.....	66
F.4 High-Speed Data Load Memory Requirements.....	68
F.5 Global Meta Cache Memory Requirements.....	68
<b>Appendix G Quantitative Limits.....</b>	<b>69</b>
<b>Appendix H Determining the Preferred WebAdmin Configuration.....</b>	<b>74</b>
H.1 WebAdmin Configurations.....	74
H.1.1 Single-Server Configuration.....	74



H.1.2 Multiserver Configuration.....	74
H.2 Installing WebAdmin in a Single-Server Configuration.....	75
H.3 Installing WebAdmin in a Multiserver Configuration.....	76
Appendix I Supported contrib Modules and Extensions Provided by External Projects.....	77
Index.....	78

# 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.



#### Note

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.



See

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) 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
- Microsoft(R) Windows Server(R) 2022 Datacenter
- Microsoft(R) Windows Server(R) 2022 Standard
- Microsoft(R) Windows Server(R) 2022 Essentials



Note

- The following components of Windows Server(R) 2016, Windows Server(R) 2019 and Windows Server(R) 2022 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 name	Version	Product name of FUJITSU Enterprise Postgres		Remarks
			AE	SE	
1	Visual Studio	2015 2017 2019	Y	Y	Required when using applications that are developed in Visual Studio.

No.	Software name	Version	Product name of FUJITSU Enterprise Postgres		Remarks
			AE	SE	
2	.NET Framework	4.6.1 or later 4.7/4.7.x 4.8	Y	Y	
3	Perl	5.26	Y	Y	Required when using PL/Perl.
4	Python	3.9.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

### Note

- The following programs are installed during installation of FUJITSU Enterprise Postgres:
  - 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 client that can be connected to the FUJITSU Enterprise Postgres server feature.

Table 2.3 Connectable client

OS	Product name
Windows	FUJITSU Software Enterprise Postgres Client 14 or later
Linux	

### Note

The connection from a client product of a different version to this server function depends on the compatibility of each function included in the client product with PostgreSQL, so some functions may not be available.

The following table lists server assistant that can be connected to the FUJITSU Enterprise Postgres server feature.

Table 2.4 Connectable server assistant

OS	Product name
Windows	- FUJITSU Software Enterprise Postgres Advanced Edition 14
Linux	- FUJITSU Software Enterprise Postgres Standard Edition 14

## 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) 14
- The product to be installed is FUJITSU Software Enterprise Postgres Advanced Edition (64bit) 14

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 2.5 Disk space required for installation

Directory	Required disk space (Unit: MB)
Windows system drive	7 + 266 (*1) + 18 (*2)
Installation destination of the Server	834
Installation destination of the WebAdmin	714
Installation destination of the Client (32bit)	467
Installation destination of the Client (64 bit)	485

\*1: Required for the installation of the Uninstall (middleware) tool.

\*2: Required for the installation of FJQSS.

## 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.



#### Note

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.

## 2.8 PostgreSQL Version Used for FUJITSU Enterprise Postgres

---

FUJITSU Enterprise Postgres is based on PostgreSQL 14.0.

## 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.

### Note

- 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].



### Information

- 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.

## Note

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

## See

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).

## Note

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

```
webAdminInstallFolder\data\feqwa
```

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\feqwa
```

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

## Point

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:\mcdire\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].



### Note

.....

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](#)".

## Information

- 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 multi-version 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.

## Note

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
:
```

### Note

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.

### Information

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 <code>initdb</code> 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 <code>pg_basebackup</code> 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.	

\*1: This can be used together with system log monitoring using operations management middleware (Systemwalker Centric Manager, for example).



See

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)

## Note

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.

### 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.

## Information

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.

### Note

- 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 rollback. This is the same as the WAL (Write Ahead Log).
Archive log	Stores log information for recovery

Resource	Role
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 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)	O
Transaction log storage destination	O (*1) (*2)	O
Corefile output destination	N (*3)	O

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\fssep_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 14 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\fssep_140_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.



## Note

- 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 [Applies to:], 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 [Applies to] in the [Permission entries] list.
8. Click [OK].



### Information

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) 2016:

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.

#### **In cases other than the above:**

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.

## **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 inst1" "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**



See

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.



Note

- 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.
- 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 14 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\fssep_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 screen is displayed. From this window you can log in to WebAdmin or access the product documentation.

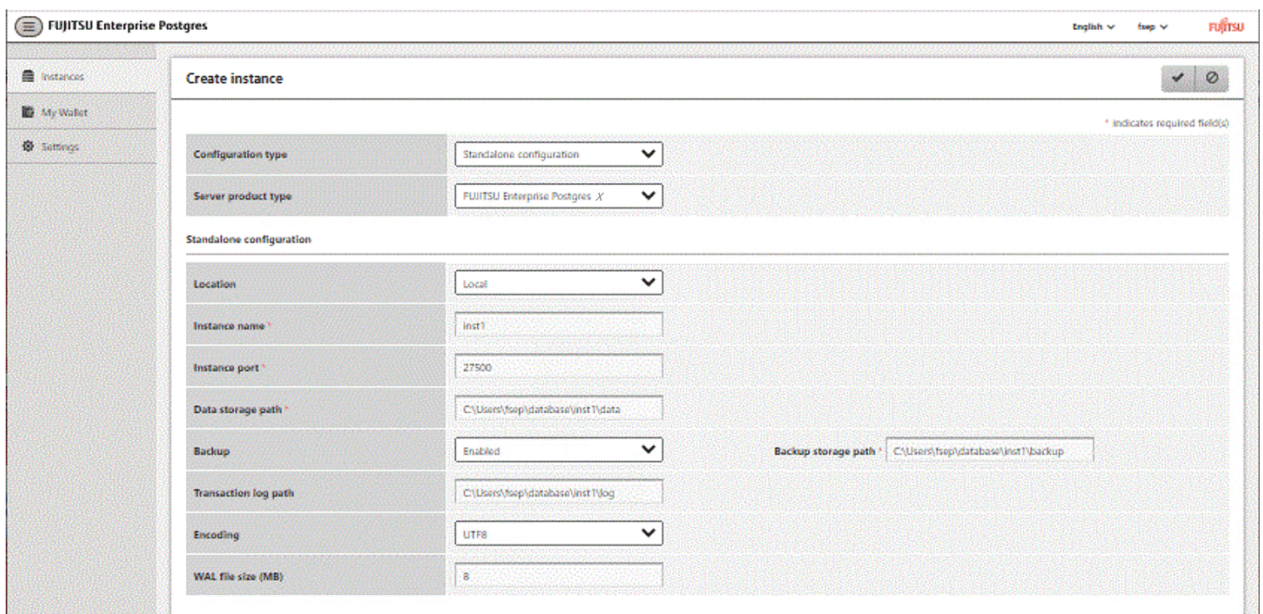
## Logging in to WebAdmin

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.



The screenshot shows the 'FUJITSU Enterprise Postgres' WebAdmin interface. The 'Create instance' form is displayed with the following fields and values:

Field	Value
Configuration type	Standalone configuration
Server product type	FUJITSU Enterprise Postgres X
Standalone configuration	
Location	Local
Instance name *	inst1
Instance port *	27500
Data storage path *	C:\Users\step\database\inst1\data
Backup	Enabled
Backup storage path *	C:\Users\step\database\inst1\backup
Transaction log path	C:\Users\step\database\inst1\log
Encoding	UTF8
WAL file size (MB)	8

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
- FUJITSU Enterprise Postgres 13 Instances
- FUJITSU Enterprise Postgres 14 Instances

The default is "FUJITSU Enterprise Postgres 14".

WebAdmin can create and manage instances compatible with the following, but new functionality in FUJITSU Enterprise Postgres 14 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
  - FUJITSU Enterprise Postgres 12
  - FUJITSU Enterprise Postgres 13
- [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

### Note

- 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

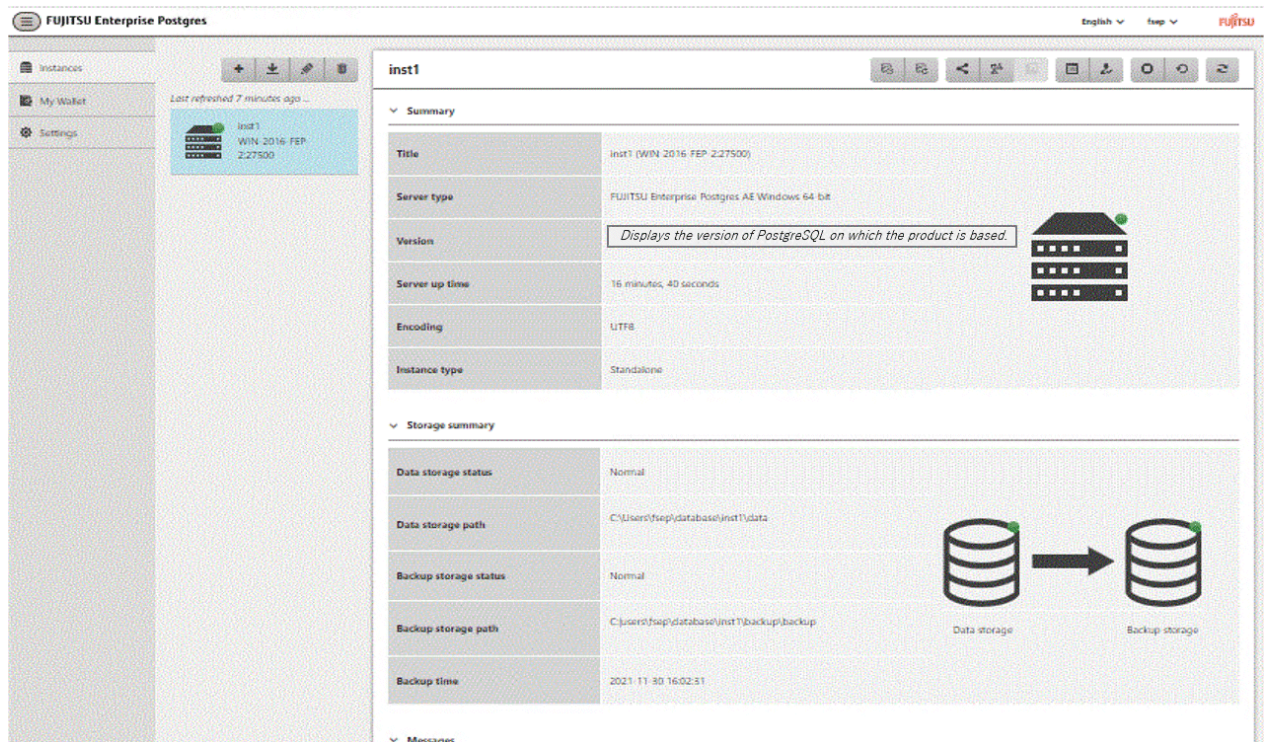
Refer to "[Editing instance information](#)" for details.

- 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  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\feepwa
```

### Note

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](#)

## 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

## Note




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

You can also modify the following files directly, but WebAdmin may not work correctly if the records span multiple lines. Therefore, change the record to a single row.

- pg\_hba.conf
- pg\_ident.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  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 .

## Note

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.



## Note

- 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.

## Note

- 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 following file contains records that span multiple lines, change the record to a single line before importing.
    - pg\_hba.conf
    - pg\_ident.conf
  - 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: 'C:\Users\naomi\AppData\Local\Fujitsu\fsep\_140\_WA\_64\myinst27599\log'

core\_directory: 'C:\Users\naomi\AppData\Local\Fujitsu\fsep\_140\_WA\_64\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.

### Note

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.

### Note

- 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.

### See

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\inst1 --waldir=E:\transaction\inst1 --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

## Note

- 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 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 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.

## Note

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.



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 "inst1" -U fepuser -P ***** -D D:\database\inst1
```



- 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: <i>numberOfSimultaneousConnectionsToInstance</i> + 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.



See

.....  
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) 2016:

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.

**In cases other than the above:**

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.

## 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 style="list-style-type: none"> <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 log directory in the database cluster.	<ul style="list-style-type: none"> <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	<p>Specify information to be added at the start of messages output by an instance. This information is useful for automatic monitoring of messages.</p> <p>You can output the SQLSTATE value, output time, executing host, application name, and user ID.</p> <p>Refer to "What To Log" in the PostgreSQL Documentation for details.</p> <p>Example: log_line_prefix = '%e: %t [%p]: [%l-1] user = %u,db = %d,remote = %r app = %a '</p>	reload option of the pg_ctl mode

### 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.

### Note

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 '
```



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.



Information

- 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 PATH environment variable.
2. Open a command prompt with administrator privileges and run the following command:

```
> xcopy /E "c:\Program Files\Fujitsu\fsepv<x>server64\OSS\oracle_fdw\*"
"c:\Program Files\Fujitsu\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 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 PATH environment variable to the updated path.
- 



#### Note

---

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"
```



#### Information

---

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

1. 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:

```
> xcopy /E "c:\Program Files\Fujitsu\fsepv<x>server64\OSS\pg_hint_plan\*" "c:\Program Files\Fujitsu\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
```



Refer to "Optimizer Hints" in the Application Development Guide for details.

### 4.7.2.2 Removing pg\_hint\_plan



Unsetting pg\_hint\_plan will cause hints registered in the hint\_plan.hints table to be lost. Therefore, it is recommended that pg\_dump back up the hint\_plan.hints table for each database if it is likely that pg\_hint\_plan will be used again later.

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"
```



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

1. 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:

```
> xcopy /E "c:\Program Files\Fujitsu\fsepv<x>server64\OSS\pg_dbms_stats\*" "c:\Program Files\Fujitsu\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.

### 4.7.3.2 Removing pg\_dbms\_stats



Unsetting pg\_dbms\_stats causes statistics managed by pg\_dbms\_stats to be lost. Therefore, it is recommended that you back up each table in the dbms\_stats folder of each database in binary format if you may want to use pg\_dbms\_stats again later.

```
postgres > # COPY <dbms_stats Schema's table name> TO '<Filename>' FORMAT binary;
```

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\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.



## 4.8 Deleting Instances

---

This section explains how to delete an instance.

- [4.8.1 Using WebAdmin](#)
- [4.8.2 Using Server Commands](#)

### Note

- 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 .

### Note

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.

## Information

- 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.

## Note

- 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.

## Information

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\feepwa
```

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 "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 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.

fep\_140\_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.

fep\_140\_WA\_64\_WebAdmin\_Port2

#### WebAdmin automatic start

Select whether or not to start WebAdmin when the machine is started.

## Note

- 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.

## Note

- 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


---

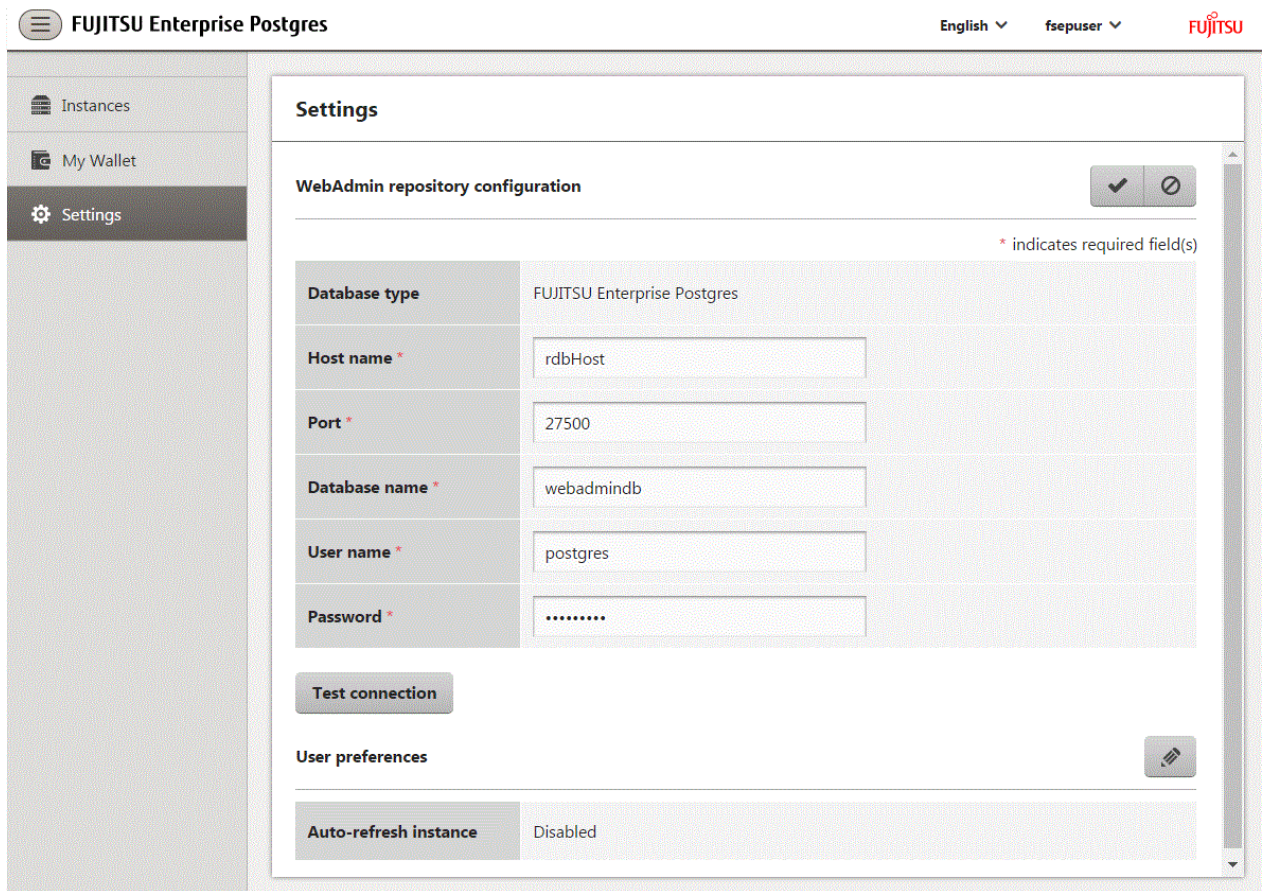
WebAdmin can be configured to use an external database, where it can store the various metadata information it uses. WebAdmin will use this database as a repository to store the information it uses to manage all the created instances. This can be a FUJITSU Enterprise PostgreSQL database or an Open Source 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.




The screenshot shows the 'Settings' page for 'FUJITSU Enterprise Postgres'. The left sidebar contains 'Instances', 'My Wallet', and 'Settings'. The main content area is titled 'Settings' and has a sub-section 'WebAdmin repository configuration' with a checkmark and a pencil icon. Below this, there are several input fields: 'Database type' (FUJITSU Enterprise Postgres), 'Host name \*' (rdbHost), 'Port \*' (27500), 'Database name \*' (webadmindb), 'User name \*' (postgres), and 'Password \*' (masked with dots). A 'Test connection' button is located below these fields. At the bottom, there is a 'User preferences' section with a pencil icon and a field for 'Auto-refresh instance' set to 'Disabled'. A note at the top right of the configuration section states '\* indicates required field(s)'.

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

### Note

- 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.

## Note



- 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.

## Note

- 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)
- 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.



## Note

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
<b>Instance Configuration</b>		
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
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 save size (MB)	wal_keep_size
	Hot standby	hot_standby
	Synchronous standby names	synchronous_standby_names
	WAL receiver timeout (ms)	wal_receiver_timeout
<b>Edit instance</b>		
	Instance name	n/a
	Instance port	port
	Backup storage path	backup_destination

\*1: 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.

### Information

- Calculate the maximum number of connections using the formula below:

$$\text{maximumNumberOfConnections} = \text{maximumNumberOfConnectionsFromApplications} + 3 (*1)$$

\*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`.

$$\text{maximumNumberOfConnections} = \text{maximumNumberOfConnectionsFromApplications} + \text{superuser_reserved_connections}$$

- 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
<code>event_source (*1)</code>	<code>'fsep_version_userName_instanceNamePortNumber'</code>
<code>listen_addresses</code>	<code>*</code>
<code>log_destination</code>	<code>'stderr,eventlog'</code>
<code>logging_collector</code>	<code>on</code>
<code>log_line_prefix</code>	<code>'%e: %t [%p]: [%l-1] user = %u,db = %d,remote = %r app = %a '</code>
<code>log_filename (*2) (*3)</code>	<code>'logfile-%a.log'</code>
<code>log_truncate_on_rotation</code>	<code>on</code>
<code>log_rotation_age</code>	<code>1d</code>

\*1: *PortNumber* is the port number of the database server specified when creating the instance.

\*2: The server logs are split into files based on the day of the week, and are rotated after each week.

\*3: 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)

\*1: 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.

### Note

- 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
  - 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	<p><math>27(*1) + \text{NULL map} + \text{OID} + \text{column data}</math></p> <p>NULL map: <math>\text{Number of columns} / 8 (*2)</math>            OID: 4            Column data: Sum of column lengths</p> <p>*1: Record header section            *2: Round the result up to the next integer.</p> <ul style="list-style-type: none"> <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.                For example, if the calculated length is <math>27 + 1 / 8</math> (rounded up) + 0 = 28 bytes, add 4 to make the length 32 bytes.</li> <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.                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 bigint, c3 int, c4 box)                Estimation: CHAR type 1 byte + boundary adjustment of 7 bytes for BIGINT type 8 bytes + BIGINT type 8 bytes + INT type 4 bytes + boundary adjustment of 12 bytes for BOX type 32 bytes + BOX type 32 bytes = 64 bytes</li> <li>- 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.</li> <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 "<a href="#">Table E.2 Estimation formula when the record length exceeds 2032 bytes</a>" to estimate the table size.</li> </ul>
(2) Page size requirement	<p><math>8192 (*1) \times \text{fillfactor} (*2) - 24 (*3)</math></p> <p>*1: Page length (8192)            *2: Value of the fillfactor specified in the table definitions (if omitted, 100%)            *3: Page header (24)</p> <ul style="list-style-type: none"> <li>- The calculated (2) page size requirement will be rounded down to the nearest integer.</li> </ul>
(3) Number of records per page	<p><math>(2) \text{ Page size requirement} / ((1) \text{ record length} + 4 (*1))</math></p> <p>*1: Pointer length (4)</p> <ul style="list-style-type: none"> <li>- The result will be rounded down to the nearest integer.</li> </ul>

Item	Estimation formula (bytes)
(4) Number of pages required for storing records	Total number of records / (3) number of records per page - 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. - This is the safety factor assumed if vacuuming is performed for garbage collection in tables and indexes.

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. - This is the safety factor assumed if vacuuming is performed for garbage collection in tables and indexes.

## 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 " <a href="#">E.3 Sizes of Data Types</a> " 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 " <a href="#">Table E.4 Estimation formula when the key data length exceeds 512 bytes</a> " 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 fillfactor 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 for storing indexes	Total number of records / (3) number of entries per page - 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. 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. 3) The number of bytes per character depends on the character set (refer to "E.3.4 Number of Bytes per Character" for details).
polygon	Length of size portion + 36 + 16 x number of vertices	
decimal	Length of size portion + 2 + (integer precision / 4 + decimal precision / 4) x 2	
numeric		
bytea	Length of size portion + real data length	
character varying( <i>n</i> ), varchar( <i>n</i> )	Length of size portion + number of characters x number of bytes per character	
character( <i>n</i> ), char( <i>n</i> )	Length of size portion + <i>n</i> 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.

## Note

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
<code>shared_buffers</code>	30 percent of the internal memory of the machine.
<code>max_connections</code>	100
<code>max_prepared_transactions</code>	100

### Local memory amount

$$\begin{aligned} localMemoryAmount = & processStackArea \\ & + memoryUsedInDbSessionsThatUseTempTables \\ & + memoryUsedInDbSessionsThatPerformSortAndHashTableOperations \\ & + memoryUsedInMaintenanceOperations \\ & + baseMemoryUsedInEachProcess \\ & + memoryUsedPreparingForDataAccess \end{aligned}$$

### Process stack area

$$\begin{aligned} processStackArea \\ = & max\_stack\_depth \times (max\_connections + autovacuum\_max\_workers + 9) \end{aligned}$$

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

$$\begin{aligned} memoryUsedInDbSessionsThatUseTempTables \\ = & temp\_buffers \times max\_connections \end{aligned}$$

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

$$\begin{aligned} memoryUsedInDbSessionsThatPerformSortAndHashTableOperations \\ = & work\_mem (*1) \times max\_connections \end{aligned}$$

\*1) For hash table operations, multiply work\_mem by hash\_mem\_multiplier.

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

```
baseMemoryUsedInEachProcess
= baseMemoryUsedInOneProcess x (max_connections + autovacuum_max_workers + 9)
```

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:

```
baseMemoryUsedInOneProcess
= (1.9KB x All user tables + 2.9KB x All user indexes + 1.0KB x All user columns) x 1.5(*1)
```

If you enable the Global Meta Cache feature, use the following formula:

```
baseMemoryUsedInOneProcess
= (All user tables + All user indexes + All user columns) x 1.0KB x 1.5 (*1)
+ (All user tables x 1.4KB + All user indexes x 2.4KB)
```

\*1) Safety Factor (1.5)

There are variable length information. This value takes that into account.

#### Memory used preparing for data access

```
memoryUsedPreparingForDataAccess
= variationAmount x (max_connections + autovacuum_max_workers + 4)
```

```
where variationAmount = shared_buffers / 8KB x 4 bytes
(note that 8KB is the page length, and 4 bytes is the size of page management data)
```

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).

$$memForData = (numOfRowsInTables) \times (numOfBytesPerRow) + (wosSize)$$

Number of bytes per row

$$\begin{aligned} numOfBytesPerRow \\ = (19 + (numOfColsInCreateIndexStatement) / 8 + (numOfBytesPerSingleColValue)) \times 1.1 \end{aligned}$$

Note: Round up the result to the nearest integer.

WOS size

$$wosSize = (numOfWosRows) / 185 \times 8096$$

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

$$\begin{aligned} memForEachProcess \\ = memUsedPerScanning \\ + memUsedForVciMaintenance \\ + memUsedByCreateIndexStatement \end{aligned}$$

Memory used per scanning

- Parallel scan

$$\begin{aligned} memUsedPerScanning \\ = vci.shared\_work\_mem + (numOfParallelWorkers + 1) \times vci.maintenance\_work\_mem \end{aligned}$$

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

$$memUsedPerScanning = vci.max\_local\_ros + vci.maintenance\_work\_mem$$

### Note

- 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

$$memUsedForVciMaintenance = vci.maintenance\_work\_mem \times vci.control\_max\_workers$$

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 0.1KB x max_connections x 1.5 (*1)

```

\*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.

# 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)

\*1: This is the character string byte length when converted by the server character set character code.

\*2: If an identifier that exceeds 63 bytes in length is specified, the excess characters are truncated and it is processed.

Table G.2 Database object

Item	Limit
Number of databases	Less than 4,294,967,296 (*1)

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)

\*1: 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

\*1: Operation might proceed correctly even if operations are performed with a quantity outside the limits.

\*2: This is the character string byte length when converted by the server character set character code.

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)

\*1: 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 expression	2 bytes	From -32,768 to 32,767
		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	

\*1: 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



Table G.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
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)

\*1: Operation might proceed correctly even if operations are performed with a quantity outside the limits.

\*2: The total number of all database objects must be less than 4,294,967,296.

\*3: 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)

\*1: 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



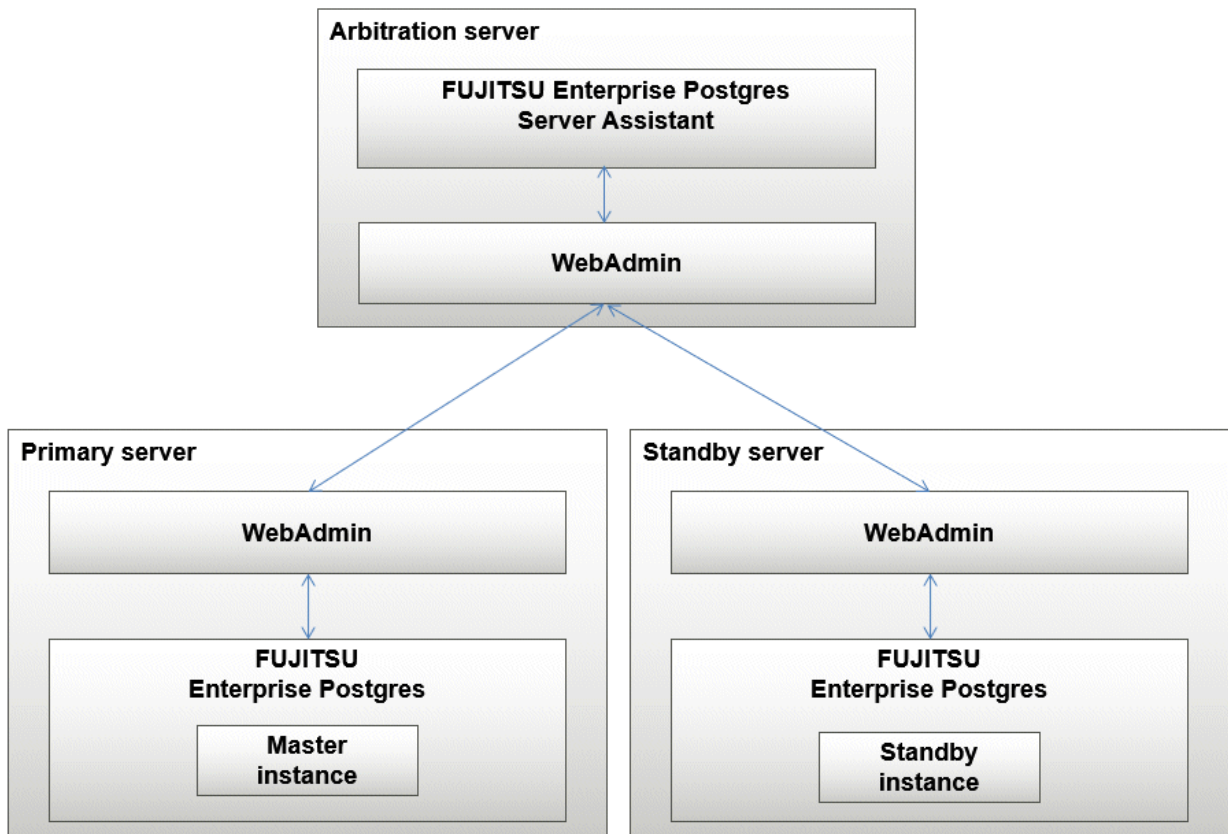
### H.1.2 Multiserver Configuration

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



See

.....  
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



## Information

.....  
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.

# Index

[C]			
Changing client authentication information.....	27	Removing WebAdmin.....	51
Changing instance settings.....	26	Required Operating System.....	3
Check the disk space.....	8	Required Patches.....	5
Check the installed products and determine the installation method.....	8	[S]	
Client Authentication Information settings.....	33	Settings related to connection.....	33
Creating an Instance.....	23,30	Setting Up and Removing WebAdmin.....	49
Creating an Instance Administrator.....	14	Setting Up WebAdmin.....	49
Creating Instances.....	21	Starting the Web Server Feature of WebAdmin.....	50
		Stopping the Web Server Feature of WebAdmin.....	50
[D]		Supported contrib Modules and Extensions Provided by External Projects.....	77
Disk Space Required for Installation.....	5	[T]	
		TCP/IP Protocol.....	5
[E]			
Editing instance information.....	28	[U]	
Excluded Software.....	4	Uninstallation.....	2,42
		Uninstallation in Interactive Mode.....	42
[F]		Uninstallation in Silent Mode.....	44
Firewall.....	19,33	uninstaller.....	47
		Using the initdb Command.....	30
[H]		Using WebAdmin.....	22
Hardware Environment.....	5		
How to Set Up the Pop-up Blocker.....	48	[W]	
		WebAdmin automatic start.....	49
[I]		Web server port number.....	49
Importing Instances.....	28	When an Instance was Created with the initdb Command.....	33
Installation.....	7	When an Instance was Created with WebAdmin.....	33
Installation in Interactive Mode.....	2,9		
Installation in Silent Mode.....	2,11		
Installation Procedure.....	1		
Installation Types.....	1		
Instance configuration.....	27		
[L]			
Logging in to WebAdmin.....	22		
[M]			
Multi-Version Installation.....	1		
[N]			
New Installation.....	1		
[O]			
Operating Environment.....	3		
Operating Method Types and Selection.....	13		
[P]			
Port number to use when Tomcat is stopped.....	49		
postgresql.conf.....	55		
Pre-installation Tasks.....	8		
Preparations for Setup.....	14		
[R]			
Recommended Browser Settings.....	48		
Reinstallation.....	1		
Related Software.....	3		
Remove applied updates.....	8		