

FUJITSU Enterprise Postgres 12 on IBM LinuxONE™



Operation Guide

Linux J2UL-2628-01ZLZ0(00) September 2020

Preface

Purpose of this document

The FUJITSU Enterprise Postgres database system extends the PostgreSQL features and runs on the Linux platform. This document is the FUJITSU Enterprise Postgres Operation Guide.

Intended readers

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

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

- PostgreSQL
- SQL
- Linux

Structure of this document

This document is structured as follows:

Chapter 1 Operating FUJITSU Enterprise Postgres

Describes how to operate FUJITSU Enterprise Postgres.

Chapter 2 Starting an Instance and Creating a Database

Describes how to start a FUJITSU Enterprise Postgres instance, and how to create a database.

Chapter 3 Backing Up the Database

Describes how to back up the database.

Chapter 4 Configuring Secure Communication Using Secure Sockets Layer

Describes communication data encryption between the client and the server.

Chapter 5 Protecting Storage Data Using Transparent Data Encryption

Describes how to encrypt the data to be stored in the database.

Chapter 6 Data Masking

Describes the data masking feature.

Chapter 7 Periodic Operations

Describes the periodic database operations that must be performed on FUJITSU Enterprise Postgres.

Chapter 8 Streaming Replication Using WebAdmin

Describes how to create a streaming replication cluster using WebAdmin.

Chapter 9 Installing and Operating the In-memory Feature

Describes how to install and operate the in-memory feature.

Chapter 10 Parallel Query

Describes the factors taken into consideration by FUJITSU Enterprise Postgres when performing parallel queries.

Chapter 11 High-Speed Data Load

Describes how to install and operate high-speed data load.

Chapter 12 Global Meta Cache

Describes how to use Grobal Meta Cache feature.

Chapter 13 Backup/Recovery Using the Copy Command

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

Chapter 14 Actions when an Error Occurs

Describes how to perform recovery when disk failure or data corruption occurs.

Appendix A Parameters

Describes the FUJITSU Enterprise Postgres parameters.

Appendix B System Administration Functions

Describes the system administration functions of FUJITSU Enterprise Postgres.

Appendix C System Views

Describes how to use the system view in FUJITSU Enterprise Postgres.

Appendix D Tables Used by Data Masking

Describes the tables used by the data masking feature.

Appendix E Tables Used by High-Speed Data Load

Describes the tables used by high-speed data load.

Appendix F Starting and Stopping the Web Server Feature of WebAdmin

Describes how to start and stop WebAdmin (Web server feature).

Appendix G WebAdmin Wallet

Describes how to use the Wallet feature of WebAdmin.

Appendix H WebAdmin Disallow User Inputs Containing Hazardous Characters

Describes characters not allowed in WebAdmin.

Appendix I Collecting Failure Investigation Data

Describes how to collect information for initial investigation.

Appendix J Operation of Transparent data Encryption in File-based Keystores

Describes operation of transparent data Encription in file-based keystores.

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: September 2020

Copyright

Copyright 2019-2020 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 Operations Using Commands | 3 |
| 1.4 Operating Environment of FUJITSU Enterprise Postgres | 4 |
| 1.4.1 Operating Environment | |
| 1.4.2 File Composition | 6 |
| 1.5 Notes on Compatibility of Applications Used for Operations | 7 |
| 1.6 Notes on Upgrading Database Instances | |
| 1.6.1 Additional Steps for upgrading to FUJITSU Enterprise Postgres with Vertical Clustered Index (VCI) Enabled | |
| Chapter 2 Starting an Instance and Creating a Database | ٩ |
| 2.1 Starting and Stopping an Instance | |
| 2.1 Using WebAdmin. | |
| 2.1.2 Using Server Commands | |
| 2.2 Creating a Database | |
| 2.2 Creating a Database | 13 |
| Chapter 3 Backing Up the Database | 14 |
| 3.1 Periodic Backup | |
| 3.2 Backup Methods | |
| 3.2.1 Using WebAdmin | |
| 3.2.2 Using Server Commands | |
| | |
| Chapter 4 Configuring Secure Communication Using Secure Sockets Layer | |
| 4.1 Configuring Communication Data Encryption | 19 |
| 4.1.1 Issuing a Certificate | |
| 4.1.2 Deploying a Server Certificate File and a Server Private Key File | |
| 4.1.3 Distributing a CA Certificate File to the Client | 20 |
| 4.1.4 Configuring the Operating Environment for the Database Server | 20 |
| 4.1.5 Configuring the Operating Environment for the Client | 20 |
| 4.1.6 Performing Database Multiplexing | 21 |
| Chapter 5 Protecting Storage Data Using Transparent Data Encryption | |
| 5.1 Protecting Data Using Encryption | |
| 5.2 Preparing for HSM Collaboration | 23 |
| 5.3 Setting the Master Encryption Key | 23 |
| 5.4 Opening the Keystore. | |
| 5.5 Encrypting a Tablespace | 25 |
| 5.6 Checking an Encrypted Tablespace | |
| 5.7 Managing the Keystore | |
| 5.7.1 Changing the Master Encryption Key | |
| 5.7.2 Changing the HSM master key | |
| 5.7.3 Changing User Pins | |
| 5.7.4 Enabling Automatic Opening of the Keystore | |
| 5.7.5 Backing Up and Recovering the Keystore | |
| 5.8 Backing Up and Restoring/Recovering the Database | |
| 5.9 Importing and Exporting the Database | |
| 5.10 Encrypting Existing Data. | |
| 5.11 Operations in Cluster Systems. | |
| 5.11.1 Database Multiplexing Mode | |
| 5.12 Security-Related Notes. | |
| 5.13 Tips for Installing Built Applications | |
| | - |
| Chapter 6 Data Masking | |
| 6.1 Masking Policy | |

| 6.1.1 Masking Target | |
|---|----|
| 6.1.2 Masking Type | |
| 6.1.3 Masking Condition | |
| 6.1.4 Masking Format | |
| 6.2 Usage Method | |
| 6.2.1 Creating a Masking Policy | |
| 6.2.2 Changing a Masking Policy | |
| 6.2.3 Confirming a Masking Policy | 41 |
| 6.2.4 Enabling and Disabling a Masking Policy | |
| 6.2.5 Deleting a Masking Policy | |
| 6.3 Data Types for Masking | |
| 6.4 Security Notes | |
| Chapter 7 Periodic Operations | |
| 7.1 Configuring and Monitoring the Log | |
| 7.2 Monitoring Disk Usage and Securing Free Space | |
| 7.2.1 Monitoring Disk Usage | |
| 7.2.2 Securing Free Disk Space | |
| 7.3 Automatically Closing Connections | |
| 7.4 Monitoring the Connection State of an Application | |
| 7.4.1 Using the View (pg_stat_activity) | |
| 7.5 Reorganizing Indexes | |
| 7.6 Monitoring Database Activity | |
| 7.6.1 Information that can be Collected | |
| 7.6.2 Collection Configuration | |
| 7.6.3 Information Reset | |
| Chapter 8 Streaming Replication Using WebAdmin | |
| 8.1 Creating a Standby Instance | |
| 8.2 Promoting a Standby Instance | |
| 8.3 Converting an Asynchronous Replication to Synchronous | |
| 8.4 Converting a Synchronous Replication to Asynchronous | |
| 8.5 Joining a Replication Cluster | |
| Chapter 9 Installing and Operating the In-memory Feature | |
| 9.1 Installing Vertical Clustered Index (VCI) | |
| 9.1.1 Evaluating whether to Install VCI | |
| 9.1.2 Estimating Resources | |
| 9.1.3 Setting up | |
| 9.1.3.1 Setting Parameters | |
| 9.1.3.2 Installing the Extensions | |
| 9.1.3.3 Creating a VCI | |
| 9.1.3.4 Confirming that the VCI has been Created | 61 |
| 9.1.4 Data that can Use VCI | |
| 9.1.4.1 Relation Types | |
| 9.1.4.2 Data Types | |
| 9.2 Operating VCI | |
| 9.2.1 Commands that cannot be Used for VCI | |
| 9.2.2 Data Preload Feature | |
| Chapter 10 Parallel Query | |
| 10.1 CPU Load Calculation | |
| 10.2 Increase of Workers during Runtime. | |
| Chapter 11 High-Speed Data Load | 67 |
| 11.1 Installing High-Speed Data Load | |
| 11.1.1 Deciding whether to Install | |
| 11.1.2 Estimating Resources. | |
| 200000 | |

| 11.1.3 Setup | |
|---|----|
| 11.1.3.1 Setting Parameters | |
| 11.1.3.2 Installing the Extension | |
| 11.2 Using High-Speed Data Load | |
| 11.2.1 Loading Data | 69 |
| 11.2.2 Recovering from a Data Load that Ended Abnormally | |
| 11.3 Removing High-Speed Data Load | 71 |
| 11.3.1 Removing the Extension | |
| | |
| Chapter 12 Global Meta Cache | |
| 12.1 Usage | |
| 12.1.1 Deciding Whether to Enable the Global Meta Cache Feature | |
| 12.1.2 Estimating Memory for Global Meta Cache | |
| 12.1.3 How the GMC Memory Area Is Used | |
| 12.1.4 Enabling the Global Meta Cache Feature | |
| 12.1.5 Estimating Resources. | |
| 12.2 Statistics | |
| 12.2.1 System View | |
| Chapter 13 Backup/Recovery Using the Copy Command | 75 |
| 13.1 Configuration of the Copy Command | |
| 13.2 Backup Using the Copy Command | |
| 13.3 Recovery Using the Copy Command | |
| 13.4 Copy Command Interface | |
| 13.4.1 Copy Command for Backup | |
| 13.4.2 Copy Command for Recovery | |
| | |
| Chapter 14 Actions when an Error Occurs | |
| 14.1 Recovering from Disk Failure (Hardware) | |
| 14.1.1 Using WebAdmin | |
| 14.1.2 Using Server Command | 86 |
| 14.2 Recovering from Data Corruption | |
| 14.2.1 Using WebAdmin | |
| 14.2.2 Using the pgx_rcvall Command | |
| 14.3 Recovering from an Incorrect User Operation | |
| 14.3.1 Using WebAdmin | |
| 14.3.2 Using the pgx_rcvall Command | |
| 14.4 Actions in Response to an Application Error | |
| 14.4.1 When using the view (pg_stat_activity) | |
| 14.4.2 Using the ps Command | |
| 14.5 Actions in Response to an Access Error | |
| 14.6 Actions in Response to Insufficient Space on the Data Storage Destination | |
| 14.6.1 Using a Tablespace | |
| 14.6.2 Replacing the Disk with a Larger Capacity Disk | |
| 14.6.2.1 Using WebAdmin | |
| 14.6.2.2 Using Server Commands | |
| 14.7 Actions in Response to Insufficient Space on the Backup Data Storage Destination | |
| 14.7.1 Temporarily Saving Backup Data | |
| 14.7.1.1 Using WebAdmin | |
| 14.7.1.2 Using Server Commands | |
| 14.7.2 Replacing the Disk with a Larger Capacity Disk | |
| 14.7.2.1 Using WebAdmin | |
| 14.7.2.2 Using Server Commands | |
| 14.8 Actions in Response to Insufficient Space on the Transaction Log Storage Destination | |
| 14.8.1 Replacing the Disk with a Larger Capacity Disk | |
| 14.8.1.1 Using WebAdmin. | |
| 14.8.1.2 Using Server Commands | |
| 14.9 Errors in More Than One Storage Disk | |

| 14.10 Actions in Response to Instance Startup Failure | |
|---|-----|
| 14.10.1 Errors in the Configuration File | |
| 14.10.2 Errors Caused by Power Failure or Mounting Issues | |
| 14.10.3 Other Errors | |
| 14.10.3.1 Using WebAdmin | |
| 14.10.3.2 Using Server Commands | |
| 14.11 Actions in Response to Failure to Stop an Instance | |
| 14.11.1 Using WebAdmin | |
| 14.11.2 Using Server Commands | |
| 14.11.2.1 Stopping the Instance Using the Fast Mode | |
| 14.11.2.2 Stopping the Instance Using the Immediate Mode | |
| 14.11.2.3 Forcibly Stopping the Server Process | |
| 14.12 Actions in Response to Failure to Create a Streaming Replication Standby Instance | |
| 14.13 Actions in Response to Error in a Distributed Transaction | |
| 14.14 I/O Errors Other than Disk Failure | |
| 14.14.1 Network Error with an External Disk | |
| 14.14.2 Errors Caused by Power Failure or Mounting Issues | |
| 14.15 Anomaly Detection and Resolution | |
| 14.15.1 Port Number and Backup Storage Path Anomalies | |
| 14.15.2 Mirroring Controller Anomalies | |
| Appendix A Parameters | 118 |
| Appendix B System Administration Functions | |
| B.1 WAL Mirroring Control Functions | |
| B.2 Transparent Data Encryption Control Functions | |
| B.3 Data Masking Control Functions | |
| B.3.1 pgx_alter_confidential_policy | |
| B.3.2 pgx_create_confidential_policy | |
| B.3.3 pgx_drop_confidential_policy | |
| B.3.4 pgx_enable_confidential_policy | |
| B.3.5 pgx_update_confidential_values | |
| B.4 VCI Data Load Control Function | |
| B.5 High-Speed Data Load Control Functions | |
| Appendix C System Views | |
| C.1 pgx_tablespaces | |
| C.2 pgx_stat_lwlock | |
| C.3 pgx_stat_latch | |
| C.4 pgx_stat_walwriter | |
| C.5 pgx_stat_sql | |
| C.6 pgx_stat_gmc | |
| Appendix D Tables Used by Data Masking | 140 |
| D.1 pgx_confidential_columns | |
| D.2 pgx_confidential_policies | |
| D.3 pgx_confidential_values | |
| Appendix E Tables Used by High-Speed Data Load | 142 |
| E.1 pgx_loader_state | |
| | |
| Appendix F Starting and Stopping the Web Server Feature of WebAdmin | |
| F.1 Starting the Web Server Feature of WebAdmin. | |
| F.2 Stopping the Web Server Feature of WebAdmin | |
| Appendix G WebAdmin Wallet | |
| G.1 Creating a Credential | |
| G.2 Using a Credential | 146 |

| Appendix H WebAdmin Disallow User Inputs Containing Hazardous Characters | |
|---|-----|
| Appendix I Collecting Failure Investigation Data | 148 |
| Appendix J Operation of Transparent data Encryption in File-based Keystores | 149 |
| J.1 Setting the Master Encryption Key | |
| J.2 Opening the Keystore | 149 |
| J.3 Encrypting a Tablespace | |
| J.4 Checking an Encrypted Tablespace | |
| J.5 Managing the Keystore | |
| J.5.1 Changing the Master Encryption Key | |
| J.5.2 Changing the Keystore Passphrase | |
| J.5.3 Enabling Automatic Opening of the Keystore | |
| J.5.4 Backing Up and Recovering the Keystore | |
| J.6 Backing Up and Restoring/Recovering the Database | |
| J.7 Importing and Exporting the Database | |
| J.8 Encrypting Existing Data | |
| J.9 Operations in Cluster Systems | |
| J.9.1 HA Clusters that do not Use Database Multiplexing | |
| J.9.2 Database Multiplexing Mode | |
| J.10 Security-Related Notes | 155 |
| J.11 Tips for Installing Built Applications | |
| Index | 156 |

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.

- 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.

Operation management using commands

You can use commands for configuring and operating the database and managing operations.



You cannot combine WebAdmin and server commands to perform the following operations:

- Use commands to operate an instance created using WebAdmin.
- Use WebAdmin to recover a database backed up using commands.

For instances created with WebAdmin, however, backup can be obtained with the pgx_dmpall command. Also, WebAdmin can perform recovery by using the backup obtained with the pgx_dmpall command.

- To operate an instance created using the initdb command in WebAdmin, the instance needs to be imported using WebAdmin.

Features used in each phase

The following table lists the features used in each phase for GUI-based operations and command-based operations.

| O | peration | 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. | A standby instance is created using the pg_basebackup command. |

| Operation | | Operation with the GUI | Operation with commands |
|-------------------------|----------------------------------|---|---|
| | | WebAdmin performs a base backup of the source instance and creates a standby instance. | |
| | 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 | | None. | This is defined 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. | The messages that are output to the database server log are monitored. |
| | Disk space | The status in the WebAdmin window can be checked. A warning will be displayed if the free space falls below 20%. | This is monitored using the df command of the operating system, for example. |
| | Connection status | None. | This can be checked referencing pg_stat_activity of the standard statistics view from psql or the application. |

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

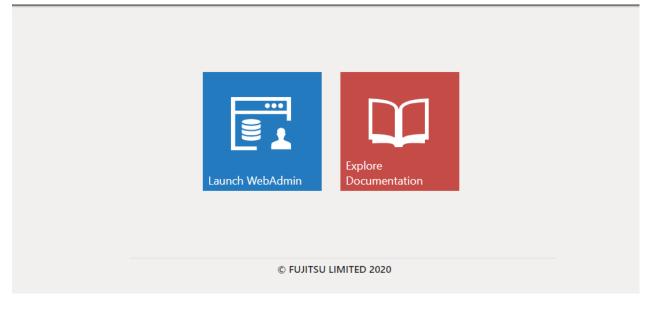
For a server with IP address "192.0.2.0" and port number "27515"

http://192.0.2.0:27515/

The startup URL window shown below is displayed. From this window you can log in to WebAdmin or access the product documentation.

FUJITSU Enterprise Postgres







- You must start the Web server feature of WebAdmin before using WebAdmin.
- Refer to "Appendix F Starting and Stopping the Web Server Feature of WebAdmin" for information on how to start the Web server feature of WebAdmin.

Log in to WebAdmin

Click [Launch WebAdmin] in the startup URL window to start WebAdmin and display the login window.

To log in, specify the following values:

- [User name]: User name (OS user account) of the instance administrator
- [Password]: Password corresponding to the user name

関 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 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.



- 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.

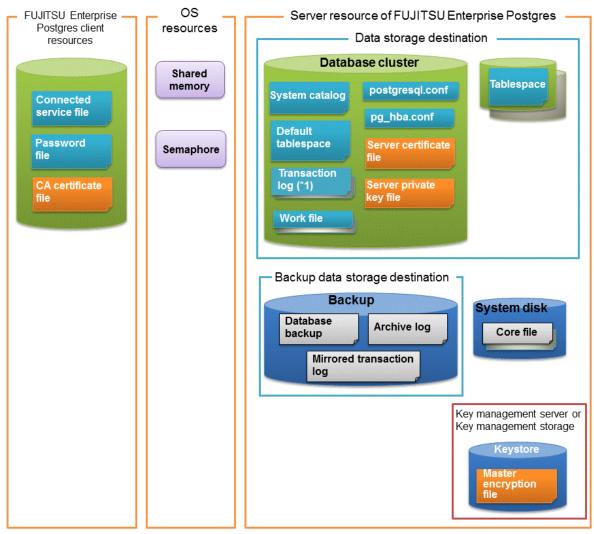
- Refer to "Configuring Environment Variables" in the Installation and Setup Guide for Client for information on the values to be set in the environment variables.

1.4 Operating Environment of FUJITSU Enterprise Postgres

This section describes the operating environment and the file composition of FUJITSU Enterprise Postgres.

1.4.1 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

| Туре | Role |
|---------------|---|
| Shared memory | Used when a database process exchanges information with an external |
| Semaphore | process. |

| Туре | Role |
|-------------------------|---|
| Connection service file | Specifies information, such as the host name, user name, and password, for connecting to FUJITSU Enterprise Postgres. |
| Password file | Securely manages the password for connecting to FUJITSU Enterprise Postgres. |
| CA certificate file | CA (certificate authority) certificate used for server authentication when encrypting communication data. |

| Туре | Role |
|---|--|
| Database cluster | Database storage area on the database storage disk. It is a collection of databases managed by an instance. |
| System catalog | Contains information required for the system to run, including the database definition information and the operation information created by the user |
| Default tablespace | Contains table files and index files stored by default. |
| Transaction log | Contains log information in case of a crash recovery or rollback. This is the same as the WAL (Write Ahead Log). |
| Work file | Work file used when executing applications or commands. |
| postgresql.conf | Contains information that defines the operating environment of FUJITSU Enterprise Postgres. |
| pg_hba.conf | FUJITSU Enterprise Postgres uses this file to authenticate individual client hosts. |
| Server certificate file | Contains information about the server certificate to be used when encrypting communication data and authenticating a server |
| Server private key file | Contains information about the server private key to be used when encrypting communication data and authenticating a server |
| Tablespace | Stores table files and index files in a separate area from the database cluster. Specify a space other than that under the database cluster. |
| Backup | Stores the data required for recovering the database when an error, such as disk failure, occurs. |
| Database backup | Contains the backup data for the database. |
| Archive log | Contains the log information for recovery. |
| Mirrored transaction log (mirrored WAL) | Enables a database cluster to be restored to the state immediately before an error even if both the database cluster and transaction log fail when performing backup/recovery operations using the pgx_dmpall command or WebAdmin. |
| Core file | FUJITSU Enterprise Postgres process core file that is output when an error occurs during a FUJITSU Enterprise Postgres process. |
| Key management server or key management storage | Server or storage where the master encryption key file is located. |
| Master encryption key file | Contains the master encryption key to be used when encrypting storage data. The master encryption key file is managed on the key management server or key management storage. |

1.4.2 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.

| Eile type | Paguirad | Quantity | How to creative the location | |
|------------------|-----------------|----------------|--|--|
| Table 1.4 Number | of files within | n a single ins | stance and how to specify their location | |

| File type | Required | Quantity | How to specify the location |
|------------------|----------|----------|---|
| Program files | Y | Multiple | Note that "< <i>x</i> >" indicates the product version. /opt/fsepv< <i>x</i> >server64 |
| Database cluster | Y | 1 | Specify using WebAdmin or server commands. |
| Tablespace | Y | Multiple | Specify a space other than that under the database cluster, using the DDL statement. |

| File type | Required | Quantity | How to specify the location |
|------------------------------------|----------|----------|---|
| 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) | Ν | 1 | Specify using postgresql.conf. |
| Server private key file (*1) | Ν | 1 | Specify using postgresql.conf. |
| Master encryption key file (*1) | Ν | 1 | Specify the directory created as the key store using postgresql.conf. |
| Connection service file (*1) | Ν | 1 | Specify using environment variables. |
| Password file (*1) | Ν | 1 | Specify using environment variables. |
| CA certificate file (*1) | Ν | 1 | Specify using environment variables. |

Y: Mandatory

N: Optional

*1: Set manually when using the applicable feature.

G Note

- Do not use an NFS for files used in FUJITSU Enterprise Postgres except when creating a database space in a storage device on a network.
- If anti-virus software is used, set scan exception settings for directories so that none of the files that comprise FUJITSU Enterprise Postgres are scanned for viruses. Alternatively, if the files that comprise FUJITSU Enterprise Postgres are to be scanned for viruses, stop FUJITSU Enterprise Postgres and perform the scan when tasks that use FUJITSU Enterprise Postgres are not operating.

•••••

1.5 Notes on Compatibility of Applications Used for Operations

When you upgrade FUJITSU Enterprise Postgres to a newer version, there may be some effect on applications due to improvements or enhancements in functionality.

Take this into account when creating applications so that you can maintain compatibility after upgrading to a newer version of FUJITSU Enterprise Postgres.



Refer to "Notes on Application Compatibility" in the Application Development Guide for details.

1.6 Notes on Upgrading Database Instances

When upgrading FUJITSU Enterprise Postgres 11 or newer database instances 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
- oracle_compatible
- pg_dbms_stats
- pg_hint_plan

For all databases except "template0", execute the following command to remove these extensions:

DROP EXTENSION extensionName;

Once the pg_upgrade operation is complete, for all databases except "template0", execute the following command to re-create these extensions as required:

CREATE EXTENSION extensionName;



- It is strongly recommended to back up the database using pg_dump before performing pg_upgrade or using DROP EXTENSION.
- If there are any columns created in the user tables using a data type from these extensions, then DROP EXTENSION will also drop these columns. Therefore, it is essential that alternate upgrade mechanisms are considered instead of pg_upgrade, in such scenarios. These may include pg_dump/pg_restore.

1.6.1 Additional Steps for upgrading to FUJITSU Enterprise Postgres with Vertical Clustered Index (VCI) Enabled

When upgrading FUJITSU Enterprise Postgres 11 instances that are using the VCI extension to FUJITSU Enterprise Postgres 12 using pg_upgrade, additional steps must be performed because of the incompatibility of the VCI extension between FUJITSU Enterprise Postgres 12 and FUJITSU Enterprise Postgres 11 or earlier.

Follow the procedure below in all databases in the FUJITSU Enterprise Postgres 11 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 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

関 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.

G 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

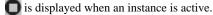
Start an instance by using the [Instances] tab in WebAdmin.

() is displayed when an instance is stopped.

To start a stopped instance, click **D**.

Stopping an instance

Stop an instance by using the [Instances] tab in WebAdmin.



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 2.
- 3. In the [Instances] tab, click **(**, and select the Immediate mode to stop the instance.

Checking the operating status of an instance

You can check the operating status of an instance by using the [Instances] tab. The following indicators are used to show the status of a resource.

| Status indicator | Explanation |
|------------------|--|
| | The resource is operating normally. |
| • | The resource is stopped. |
| • | There is an error in the resource. |
| • | An operation is in progress on this resource or the status is being checked. |
| <u> </u> | The resource is not operating optimally and needs intervention. |

If an instance stops abnormally, remove the cause of the stoppage and start the instance by using WebAdmin.

| FUJITSU Enterprise | Postgres | | | | English Y fsep Y FUĴÎTSU |
|--------------------|-------------------------------|-----------------------|--|--------------|--------------------------|
| Instances | + 🛓 🔌 🔟 | inst1 | | 6 6 K 14 1 | S 0 0 3 |
| My Wallet | Last refreshed 10 minutes ago | ✓ Summary | | | A |
| 🚯 Settings | Master10 fsepnode4:27872 | Title | inst1 (fsepnode4:21999) | | |
| | Standby10 fsepnode4:28745 | Server type | FUJITSU Enterprise Postgres AE Linux 6 | i4-bit | |
| | inst1 fsepnode4:21999 | Version | PostgreSQL 12.1 | | |
| | into | Server up time | | | •••• |
| | fsepnode4:24577 | Encoding | UTF8 | | |
| | | Instance type | Standalone | | |
| | | ✓ Storage summary | | | |
| | | Data storage status | Normal | | |
| | | Data storage path | /home/fsep/inst1/data | | |
| | | Backup storage status | Normal | | |
| | | Backup storage path | /home/fsep/inst1/backup/backup | Data storage | Backup storage |
| | | Backup time | 2019-09-11 11:15:41 | | • |

Figure 2.1 Example of operating status indicators

🌀 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



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.
```

.....



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.

Follow the procedure below to define a database using client commands.

An example of operations on the server is shown below.

1. Use psql command to connect to the postgres database. Execute psql postgres.

> psql postgres
psql (12.1)
Type "help" for help.

2. Create the database.

To create the database, execute the CREATE DATABASE databaseName; statement.

postgres=# CREATE DATABASE db01; CREATE DATABASE

3. Confirm that the database is created.

Execute \l+, and confirm that the name of the database created in step 2 is displayed.

postgres=# \l+

4. Disconnect from the postgres database.

Execute \q to terminate the psql command.

postgres=# \q

You can create a database using the createdb command.



Refer to "Creating a Database" in "Tutorial" in the PostgreSQL Documentation for information on creating a database using the createdb command.

Chapter 3 Backing Up the Database

This chapter describes how to back up the database.

Backup methods

The following backup methods enable you to recover data to a backup point or to the state immediately preceding disk physical breakdown or data logical failure.

- Backup using WebAdmin

This method enables you to back up data through intuitive window operations using the GUI.

WebAdmin is used for recovery.

- Backup using the pgx_dmpall command

Execute the pgx_dmpall command with a script to perform automatic backup.

To back up data automatically, you must register the process in the automation software of the operating system. Follow the procedure given in the documentation for your operating system.

The pgx_rcvall command is used for recovery.

🚮 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 13 Backup/ Recovery Using the Copy Command" for details.

Approximate backup time

The formula for deriving the approximate backup time when you use WebAdmin or the pgx_dmpall command is as follows:

backupTime = dataStorageDestinationUsage / diskWritePerformance x 1.5

- dataStorageDestinationUsage: Disk usage at the data storage destination
- *diskWritePerformance*: Maximum data volume (bytes/second) that can be written per second in the system environment where operation is performed
- 1.5: Coefficient to factor in tasks other than disk write (which is the most time-consuming step)

If using the copy command with the pgx_dmpall command, the backup time will depend on the implementation of the copy command.

G 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:
 - 5.7.5 Backing Up and Recovering the Keystore
 - 5.8 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.



- 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 "5.7.4 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</i> /bin/ pgx_walcopy.cmd "%p" " <i>backupDataStorageDestinationDirectory</i> / archived_wal/%f"' | Specify the path name of the command that will save the transaction log and the storage destination. |

Refer to "Appendix A Parameters" and "Write Ahead Log" under "Server Administration" in the PostgreSQL Documentation for information on the parameters.

Backup operation (file backup)

Use the pgx_dmpall command to perform file backup. You can even embed the pgx_dmpall command in OS automation software to perform backup.

The backup data is stored in the directory specified in the backup_destination parameter of postgresql.conf.

Specify the data storage destination in the -D option. If the -D option is omitted, the value of the PGDATA environment variable is used by default.

Example

> pgx_dmpall -D /database/inst1

G 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:

- 5.7.5 Backing Up and Recovering the Keystore

Backup status

Use the pgx_rcvall command to check the backup status.

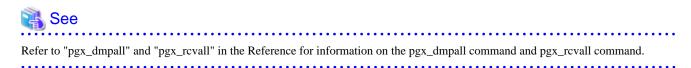
Specify the following values in the pgx_rcvall command:

- The -l option indicates backup data information.
- Specify the data storage destination in the -D option. If the -D option is omitted, the value of the PGDATA environment variable is used by default.

| > pgx_rcvall -l -D /database/instl | | |
|------------------------------------|----------|-----------------------------------|
| Date | Status | Dir |
| 2017-05-01 13:30:40 | COMPLETE | /backup/inst1/2017-05-01_13-30-40 |

If an error occurs and backup fails, a message is output to the system log.

In this case, the backup data is not optimized. Ensure that you check the backup result whenever you perform backup. If backup fails, remove the cause of failure and perform backup again.



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.

.....

Refer to "14.3.2 Using the pgx_rcvall Command" for information on using a restore point to recover the database.

G 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.

| 2 | See |
|-------|-----|
| - B J | |

Refer to "System Administration Functions" under "Functions and Operators" in the PostgreSQL Documentation for information on pg_create_restore_point.

Chapter 4 Configuring Secure Communication Using Secure Sockets Layer

If communication data transferred between a client and a server contains confidential information, encrypting the communication data can protect it against threats, such as eavesdropping on the network.

4.1 Configuring Communication Data Encryption

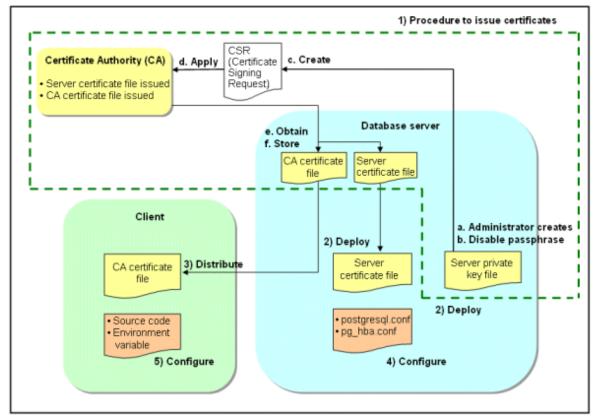
To encrypt communication data transferred between a client and a server, configure communication data encryption as described below. Communication data encryption not only protects the communication content, but it also guards against man-in-the-middle (MITM) attacks (for example, data and password theft through server impersonation).

Table 4.1 Configuration procedure

| Configuration procedure |
|---|
| 1) Issue a certificate |
| 2) Deploy a server certificate file and a server private key file |
| 3) Distribute a CA certificate file to the client |
| 4) Configure the operating environment for the database server |
| 5) Configure the operating environment for the client |

The following figure illustrates the environment for communication data encryption.

Figure 4.1 Environment for communication data encryption



4.1.1 Issuing a Certificate

For authenticating servers, you must acquire a certificate issued by the certificate authority (CA).

FUJITSU Enterprise Postgres supports X.509 standard PEM format files. If the certificate authority issues a file in DER format, use a tool such as the openssl command to convert the DER format file to PEM format.

The following provides an overview of the procedure. Refer to the procedure published by the public or independent certificate authority (CA) that provides the certificate file for details.

- a. Create a server private key file
- b. Disable the passphrase for the server private key file
- c. Create a CSR (signing request for obtaining a server certificate) from the server private key file
- d. Apply to the certificate authority (CA) for a server certificate
- e. Obtain a server certificate file and a CA certificate file from the certificate authority (CA)
- f. Store the server certificate file and the CA certificate file Note: If you lose or destroy the certificates, you will need to have them re-issued.

The above procedure enables you to prepare the following files:

- Server private key file
- Server certificate file
- CA certificate file

4.1.2 Deploying a Server Certificate File and a Server Private Key File

Create a directory on the local disk of the database server and store the server certificate file and the server private key file in it. Use the operating system features to set access privileges for the server certificate file and the server private key file so that only the database administrator has load privileges.

Back up the server certificate file and the server private key file in the event that data corruption occurs and store them securely.

4.1.3 Distributing a CA Certificate File to the Client

Create a directory on the local disk of the client and place the distributed CA certificate file there. Use the operating system features to set load privileges to protect the CA certificate file against accidental deletion.

4.1.4 Configuring the Operating Environment for the Database Server



Refer to "Secure TCP/IP Connections with SSL" under "Server Administration" in the PostgreSQL Documentation for details.

4.1.5 Configuring the Operating Environment for the Client

🝓 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, take one of the following actions:

- Create one server certificate, replicate it, and place a copy on each server used for database multiplexing.

If sslmode is set to verify-full, add all domain names in subjectAltName.

- Create server certificate 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.

There is only one master encryption key in the database cluster, which is encrypted and stored in the keystore.

Therefore, an attacker cannot read the master encryption key from the keystore.

Keystore management

FUJITSU Enterprise Postgres works in conjunction with the IBM Z Hardware Security Module (HSM) to provide hardware management of master encryption keys for robust security. The master encryption key is encrypted based on the master key in the HSM and is never leaked out over its lifetime. Use hardware-stored keystores to reduce deployment and operating costs for keystore management.

File-based keystores that do not work with the HSM are also possible. The master encryption key is then encrypted based on the passphrase that you specify and stored in the keystore. Refer to "Appendix J Operation of Transparent data Encryption in File-based Keystores" for information about the operation of transparent data encryption in file-based.

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 hardware-based encryption/decryption

Take advantage of the CPACF (CP Assist for Cryptographic Functions) in the IBM Z processor to minimize encryption and decryption overhead. 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.

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 "5.5 Encrypting a Tablespace" for details.

Backup data

The pgx_dmpall command and pg_basebackup command create backup data by copying the OS file. Backups of the encrypted data are, therefore, also encrypted. Information is protected from leakage even if the backup medium is stolen.

WAL and temporary files

WAL, which is created by updating encrypted tables and indexes, is encrypted with the same security strength as the update target. When large merges and sorts are performed, the encrypted data is written to a temporary file in encrypted format.

Streaming replication support

You can combine streaming replication and transparent data encryption. The data and WAL encrypted on the primary server is transferred to the standby server in its encrypted format and stored.



The following are not encrypted:

- pg_dump and pg_dumpall output files
- Files output by the COPY command
- Notification event payloads that communicate using the LISTEN or NOTIFY command

5.2 Preparing for HSM Collaboration

FUJITSU Enterprise Postgres manages master encryption keys with the IBM Z Crypto Express Adapter Card. To work with the Crypto Express Adapter Card, use openCryptoki, a PKCS#11 compliant API implementation. Use the openCryptoki CCA token or the EP 11 token.

FUJITSU Enterprise Postgres uses the slot assigned to each instance to access the token through a user pin. You must configure Crypto Express Adapter Cardo and openCryptoki to access the token.

Refer to the IBM documentation for the IBM Z Crypto Express Adapter Card configuration.

Refer to the openCryptoki project documentation for openCryptoki configuration. Note that the tokname attribute must be specified because FUJITSU Enterprise Postgres requires a unique token directory.

5.3 Setting the Master Encryption Key

To use transparent data encryption, you must create a keystore and set the master encryption key.

1. Set postgresql.conf parameters.

Set the keystore_location parameter to the directory where you want to store the keystore.

Specify a different location for each database cluster.

The shared_preload_libraries parameter also sets the extension to be enabled.

The tde_z.SLOT_ID parameter specifies the slot ID that was set in "5.2 Preparing for HSM Collaboration".

```
keystore_location = '/key/store/location'
shared_preload_libraries = 'tde_z'
tde_z.SLOT_ID = 5
```

When the token model "CCA" is used, it is necessary to take care of multi-coprocessor and multi-domain selection. Specify the postgresql.conf parameter for CCA configuration so that FUJITSU Enterprise Postgres can use the specific coprocessor and domain.

- tde_z.IBM_CCA_CSU_DEFAULT_ADAPTER: this parameter enables to change a default CCA coprocessor.

- tde_z.IBM_CCA_CSU_DEFAULT_DOMAIN: this parameter enables to select a single domain.

The values of these parameters are taken over to CCA service via CCA environment variable CSU_DEFAULT_ADAPTER or CSU_DEFAULT_DOMAIN.

```
tde_z.IBM_CCA_CSU_DEFAULT_ADAPTER = 'CRP01'
tde_z.IBM_CCA_CSU_DEFAULT_DOMAIN = '3'
```

Refer to "Appendix A Parameters" for information on postgresql.conf.

Refer to IBM documentation for CCA environment variable CSU_DEFAULT_ADAPTER, CSU_DEFAULT_DOMAIN, the multicoprocessor selection capabilities and domain selection capabilities.

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 database superuser.

SELECT pgx_set_master_key('user pin');

The argument should be the user pin set in "5.2 Preparing for HSM Collaboration".

5.4 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 in the token becomes accessible 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('user pin');

user pin is the user pin configured in "5.2 Preparing for HSM Collaboration".

Refer to "B.2 Transparent Data Encryption Control Functions" for information on the pgx_open_keystore function.

Note that, in the following cases, the user pin 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 --user-pin option in the pg_ctl command, and then start the instance. This will display the prompt for the user-pin to be entered, as shown below.

```
> pg_ctl --user-pin start
Enter User PIN:
The server is starting
>
```



When using an automatically opening keystore, you do not need to enter the user pin and you can automatically open the keystore when the database server starts. Refer to "5.7.4 Enabling Automatic Opening of the Keystore" for details.

.

5.5 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 and not set with encryption algorithm, the encryption algorithm can be set with the command below.

ALTER TABLESPACE db_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.



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.

5.6 Checking an Encrypted Tablespace

The pgx_tablespaces system view displays information about whether each tablespace has been encrypted, and about the encryption algorithm. Refer to "C.1 pgx_tablespaces" for information on strings.

You can discover which tablespaces have been encrypted by executing the following SQL statements.

However, when considering continued compatibility of applications, do not reference system catalogs (pg_tablespace) directly in SQL statements.

SELECT spcname, spcencalgo
FROM pg_tablespace ts, pgx_tablespaces tsx
WHERE ts.oid = tsx.spctablespace;





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

5.7 Managing the Keystore

This section describes how to manage the keystore and the master encryption key to guard against the threat of theft.

5.7.1 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 "5.3 Setting the Master Encryption Key" for details.

After changing the master encryption key, you must immediately back up the keystore and the entire opencryptoki token directory of the slot you assigned to FUJITSU Enterprise Postgres. (This token directory has same name as was set to the tokname attribute in "5.2 Preparing for HSM Collaboration").

5.7.2 Changing the HSM master key

The master encryption key is encrypted with the master key hidden by the HSM, and the master key is highly secure, but the HSM master key can be changed. Refer to the IBM documentation for instructions on how to do this. First, stop the database server. As soon as you have change the HSM master key, back up the entire opencryptoki token directory of the slot assigned to FUJITSU Enterprise Postgres.

5.7.3 Changing User Pins

You can change the user pin used for authentication when the keystore is opened.

Refer to the openCryptoki project documentation for instructions on how to do this.

You must then either reopen the keystore with the following SQL or restart the database server.

```
SELECT pgx_open_keystore('newUserpin');
```

As soon as you change the user pin, back up the entire opencryptoki token directory of the slot you assigned to FUJITSU Enterprise Postgres.

5.7.4 Enabling Automatic Opening of the Keystore

Auto-open keystores allow you to automatically open a keystore when you start an instance without entering user pins. To enable automatic keystore opening, specify the following postgresql.conf parameters:

Example

tde_z.USER_PIN = 'user pin'

Specify the user pin as set in "5.2 Preparing for HSM Collaboration".

5.7.5 Backing Up and Recovering the Keystore

Back up the keystore and the entire opencryptoki token directory of the slot allocated to FUJITSU Enterprise Postgres at the following times in case they are corrupted or lost.

- When the master encryption key is first configured
- When the master encryption key is changed
- When the database is backed up
- When the HSM master key is changed
- When a user pin is changed

関 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
```

> tar -cf token_directory_fep_20200501.tar /var/lib/opencryptoki/fep

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('user pin')" postgres
> cp -p /key/store/location/keystore.ks /keybackup/keystore_20200505.ks
> tar -cf token_directory_fep_20200505.tar /var/lib/opencryptoki/fep
```

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) and the entire opencryptoki token directory of the slot allocated to FUJITSU Enterprise Postgres. If there is no keystore containing the latest master encryption key, restore the keystore and the entire opencryptoki token directory of the slot assigned to FUJITSU Enterprise Postgres to their 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
> tar -xf token_directory_fep_20200505.tar
```

- > tai -xi token_directory_rep_z0z00305.tai
- 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
> tar -xf token_directory_fep_20200501.tar
> pgx_rcvall -B /backup/inst1 -D /database/inst1 --user-pin
```

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 --user-pin option prompts you to enter the user pin to open the keystore.

💦 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 "5.7.4 Enabling Automatic Opening of the Keystore" for information on how to enable automatic opening of the keystore.

5.8 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 and the entire opencryptoki token directory of the slot allocated to FUJITSU Enterprise Postgres 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.



When recoverying the database, it is necessary to restore the configuration of both Crypto Express Adapter Card and openCryptoki which contains the same key as was set at the time of taking backup so that the recovered database can cooperate with hardware security module.

Backup and recovery using WebAdmin

- Backup

WebAdmin backs up encrypted data.

Back up the key store and the entire opencryptoki token directory of the slot allocated to FUJITSU Enterprise Postgres after backing up the database.

- Recovery

Restore the keystore and the entire opencryptoki token directory of the slot allocated to FUJITSU Enterprise Postgres to their state at the time of database backup. Refer to "5.7.5 Backing Up and Recovering the Keystore" for details.

Enable automatic opening of the keystore in accordance with the procedure described in "5.7.4 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 and the entire opencryptoki token directory of the slot allocated to FUJITSU Enterprise Postgres after backing up the database.

- Recovery

Restore the keystore and the entire opencryptoki token directory of the slot allocated to FUJITSU Enterprise Postgres to their 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 --user-pin option specified. This will display the prompt for the user pin 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
> tar -cf token_directory_fep_20200501.tar /var/lib/opencryptoki/fep
```

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_rcvall -B /backup/inst1 -D /database/inst1 --user-pin
```

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 --user-pin option prompts you to enter the user pin 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 "5.9 Importing and Exporting the Database" below.

Back up the key store and the entire opencryptoki token directory of the slot allocated to FUJITSU Enterprise Postgres 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 and the entire opencryptoki token directory of the slot allocated to FUJITSU Enterprise Postgres after performing the backup.

- Restore

Restore the keystore and the entire opencryptoki token directory of the slot allocated to FUJITSU Enterprise Postgres to their 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 and the entire opencryptoki token directory of the slot allocated to FUJITSU Enterprise Postgres after performing the backup.

- Recovery

Restore the keystore and the entire opencryptoki token directory of the slot allocated to FUJITSU Enterprise Postgres to their 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 --user-pin option specified. This will display the prompt for the user pin to be entered.



- Refer to "pg_ctl" under "Reference" in the PostgreSQL Documentation for information on the pg_ctl command.
- Refer to "Reference" in the PostgreSQL Documentation for information on the following commands:
 - psql
 - pg_dump
 - pg_basebackup

- Refer to the Reference for information on the following commands:
 - pgx_rcvall
 - pgx_dmpall
 - pg_dumpall

5.9 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.

5.10 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.

5.11 Operations in Cluster Systems

This section describes how to use transparent data encryption on cluster systems such as streaming replication, and database multiplexing.

5.11.1 Database Multiplexing Mode

Note the following when using transparent data encryption in environments that use streaming replication, or database multiplexing with streaming replication.

HSM Master Key Configuration

Because the master encryption key is encrypted with the HSM master key, you must load the same HSM master key on all servers that will be part of the database multiplexing operation.

For instructions on loading the primary server's HSM master key on the standby server, refer to the IBM documentation.

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.

Copy procedure for the opencryptoki token directory

Execute the following operations as the system administrator (root).

- 1. On source machine
 - a. Backup the token directory of the source machine (retain file owners and groups) in the backup
- 2. On target machine
 - a. Update the opencryptoki.conf file

Add new Slot/tokname to use on the target machine

b. Restart Slot Manager daemon (pkcsslotd)

This will create the openCryptoki token folders including a new one for new slot

- c. Copy all the backed-up token directory files into this new token directory (preserving the file owners and groups)
- d. Be sure to remove any shared memory files in /dev/shm which are associated with the new tokname
- e. Restart Slot Manager daemon (pkcsslotd)
- f. New token should be ready to use.

(The SO PIN and User PIN are same as they were on the source machine)

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.

Building and starting a standby server

occurs during the following operations.

Before using the pg_basebackup command or pgx_rcvall command to build a standby server, copy the keystore file and the entire opencryptoki token directory of the slot assigned to FUJITSU Enterprise Postgres from the primary server to 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 --user-pin option in the pg_ctl command or pgx_rcvall command and enter the user pin, or use an automatically opening keystore.



- If you start a standby server without copying the openCryptoki token directory from the primary server to the standby server, an error
 - Accessing tables defined in an encrypted tablespace on the standby server
 - Changing the master encryption key on the primary server

- If an error occurs, you can recover by copying the opencryptoki token directory from the primary server to the standby server and restarting the standby server.

Changing the master encryption key

Change the master encryption key 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 are reflected in the keystore on the standby server.



Refer to "pgx_rcvall " in the Reference for information on pgx_rcvall command.

Refer to "pg_ctl" under "Reference" in the PostgreSQL Documentation for information on pg_ctl command.

Refer to "pg_basebackup" under "Reference" in the PostgreSQL Documentation for information on pg_basebackup command.

Refer to "High Availability, Load Balancing, and Replication" under "Server Administration" in the PostgreSQL Documentation for information on how to set up streaming replication.

5.12 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.
- The 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.

5.13 Tips for Installing Built Applications

With transparent data encryption, you can easily encrypt all the data in an application without modifying the application. Database administrators install built applications in the following manner. However, this procedure stores data to the default tablespace, so take necessary action if processing differs from the original design.

1. (Normal procedure) Create an owner and a database for the built application.

```
CREATE USER crm_admin ...;
CREATE DATABASE crm_db ...;
```

2. (Procedure for encryption) Create an encrypted tablespace to store the data for the built application.

```
SET tablespace_encryption_algorithm = 'AES256';
CREATE TABLESPACE crm_tablespace LOCATION '/crm/data';
```

3. (Procedure for encryption) Configure an encrypted tablespace as the default tablespace for the owner of the built application.

```
ALTER USER crm_admin SET default_tablespace = 'crm_tablespace';
ALTER USER crm_admin SET temp_tablespaces = 'crm_tablespace';
```

4. (Normal procedure) Install the built application. The application installer prompts you to enter the host name and the port number of the database server, the user name, and the database name. The installer uses the entered information to connect to the database server and execute the SQL script. For applications that do not have an installer, the database administrator must manually execute the SQL script.

Normally, the application's SQL script includes logic definition SQL statements, such as CREATE TABLE, CREATE INDEX, and GRANT or REVOKE, converted from the entity-relationship diagram. It does not include SQL statements that create databases, users, and tablespaces. Configuring the default tablespace of the users who will execute the SQL script deploys the objects generated by the SQL script to the tablespace.

Chapter 6 Data Masking

Data masking is a feature that can change the returned data for queries generated by applications, so that it can be referenced by users. For example, for a query of employee data, digits except the last four digits of an eight-digit employee number can be changed to "*" so that it can be used for reference.

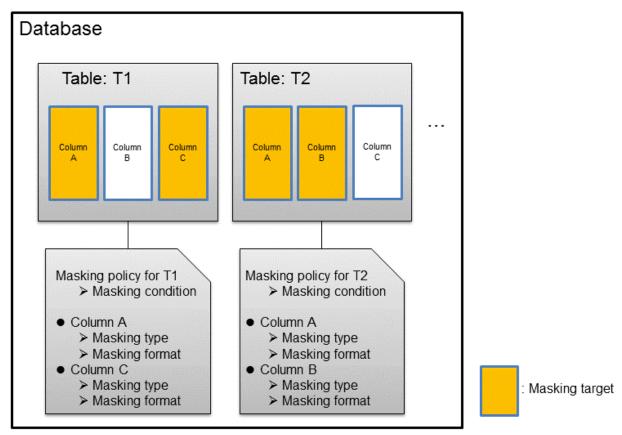


When using this feature, it is recommended that the changed data be transferred to another medium for users to reference. This is because, if users directly access the database to extract the masked data, there is a possibility that they can deduce the original data by analyzing the masking policy or query result to the masking target column.

6.1 Masking Policy

Masking policy is a method of changing data under specific conditions when it is returned for a query from an application. One masking policy can be created per table. You can configure masking target, masking type, masking condition and masking format in a masking policy.





G 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

G Note

- If a masking target is specified to INSERT...SELECT target columns, processing will be performed using data before change.

- If a masking target other than SELECT target columns is specified, processing will be performed using data before change.
- If a masking target is specified in a function where the data type will be converted, an error will occur.

6.1.2 Masking Type

Masking type is a method to change column data that is returned from queries. Specify the masking type in the function_type parameter. The following masking types can be specified and selected depending on the masking target data type.

Full masking

All the data in the specified column is changed. The changed value returned to the application that made the query varies depending on the column data type.

For example, 0 is used for a numeric type column and a space is used for a character type column.

Partial masking

The data in the specified column is partially changed. For example, digits except the last four digits of an employee number can be changed to "*".

Regular expression masking

The data in the specified column is changed via a search that uses a regular expression.

For example, for strings such as email address that can have variable length, "*" can be used to change characters preceding "@" by using a regular expression. Regular expression masking can only be used for character type data.



- If multiple valid masking targets are specified for a function, the masking type for the left-most masking target will be applied. For example, if "SELECT GREATEST(c1, c2) FROM t1" is executed for numeric type masking target c1 and c2, the masking type for c1 will be applied.

- When masking the data that includes multibyte characters, do not specify partial masking for masking type. The result may not be as expected.

6.1.3 Masking Condition

Masking condition refers to the conditions configured to perform masking. Specify the masking condition in the expression parameter. Changed or actual data can be displayed for different users by defining masking condition. An expression that returns a boolean type result needs to be specified in masking condition and masking is performed only when TRUE is returned. Refer to "Value Expressions" in the PostgreSQL Documentation for information on the expressions that can be specified. Note that expressions that include a column cannot

be specified.

For example, when masking data only for "postgres" users, specify 'current_user = "postgres" in the masking condition.



```
Specify '1=1' so the masking condition is always evaluated to be TRUE and masking is performed all the time.
```

6.1.4 Masking Format

Masking format is a combination of change method and displayed characters when the masking condition is met. Masking format varies depending on the masking type. The following describes the masking format.

Full masking

With full masking, all characters are changed to values as determined by the database. Changed characters can be referenced in the pgx_confidential_values table. Also, replacement characters can be changed using the pgx_update_confidential_values system management function.

Refer to "6.3 Data Types for Masking" for information on the data types for which data masking can be performed.

Partial masking

With partial masking, data is changed according to the content in the function_parameters parameter. The method of specifying function_parameters varies depending on the data type.

| Category | Method of specifying function_parameters | | | | | |
|----------------|--|--|--|--|--|--|
| Numeric type | 'replacementCharacter, startPosition, endPosition' | | | | | |
| | - replacementCharacter: Specify the number to display. Specify a value from 0 to 9. | | | | | |
| | - startPosition: Specify the start position of masking. Specify a positive integer. | | | | | |
| | - <i>endPosition</i> : Specify the end position of masking. Specify a positive integer that is greater than <i>startPosition</i> . | | | | | |
| | Example | | | | | |
| | Specify as below to change the values from the 1st to 5th digits to 9. | | | | | |
| | function_parameters := '9, 1, 5' | | | | | |
| | In this example, if the original data is "123456789", it will be changed to "999996789". | | | | | |
| Character type | 'inputFormat, outputFormat, replacementCharacter, startPosition, endPosition' | | | | | |
| | - <i>inputFormat</i> : Specify the current format of the data. Specify "V" for characters that will potentially be masked, and specify "F" for values such as spaces or hyphens that will not be masked. | | | | | |
| | - <i>outputFormat</i> : Define the method to format the displayed data. Specify "V" for characters that will potentially be masked. Any character to be output can be specified for each character "F" in <i>inputFormat</i> . If you want to output a single quotation mark, specify two of them consecutively. | | | | | |
| | - <i>replacementCharacter</i> . Specify any single character. If you want to output a single quotation mark, specify two of them consecutively. | | | | | |
| | - <i>startPosition</i> : Specify the position of "V" as the start position of masking. For example, to specify the position of the 4th "V" from the left, specify 4. Specify a positive integer. | | | | | |

| Category | Method of specifying function_parameters |
|---------------------|--|
| | - <i>endPosition</i> : Specify the position of "V" as an end position of masking. When working out the end position, do not include positions of "F". For example, to specify the position of the 11th "V" from the left, specify 11. Specify a positive integer that is greater than <i>startPosition</i> . |
| | Specify as below to mask a telephone number other than the first three digits using *. |
| | function_parameters := 'VVVFVVVVFVVVV, VVV-VVVV-VVVV, *, 4, 11' |
| | In this example, if the original data is "012-3156-7890", it will be changed to "012-****-***". |
| Date/timestamp type | 'MDYHMS' |
| | - M: Masks month. To mask month, enter the month from 1 to 12 after a lowercase letter m. Specify an uppercase letter M to not mask month. |
| | - D: Masks date. To mask date, enter the date from 1 to 31 after a lowercase letter d. If a value bigger than the last day of the month is entered, the last day of the month will be displayed. Specify an uppercase letter D to not mask date. |
| | - Y: Masks year. To mask year, enter the year from 1 to 9999 after a lowercase letter y. Specify an uppercase letter Y to not mask year. |
| | - H: Masks hour. To mask hour, enter the hour from 0 to 23 after a lowercase letter h. Specify an uppercase letter H to not mask hour. |
| | - M: Masks minute. To mask minute, enter the minute from 0 to 59 after a lowercase letter m. Specify an uppercase letter M to not mask minute. |
| | - S: Masks second. To mask second, enter the second from 0 to 59 after a lowercase letter s. Specify an uppercase letter S to not mask second. |
| | Example |
| | Specify as below to mask hour, minute, and second and display 00:00:00. |
| | function_parameters := 'MDYh0m0s0' |
| | In this example, if the original data is "2010-10-10 10:10", it will be changed to "2010-10-10 00:00:00". |



- Refer to "B.3.2 pgx_create_confidential_policy" for information on function_parameters.

- Refer to "6.3 Data Types for Masking" for information on the data types for which masking can be performed.

Regular expression masking

With regular expression masking, data is changed according to the content of the regexp_pattern, regexp_replacement and regexp_flags parameters. For regexp_pattern, specify the search pattern using a regular expression. For regexp_replacement, specify the replacement character to use when data matches the search pattern. For regexp_flags, specify the regular expression flags.

.....

💕 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".



- Refer to "POSIX Regular Expressions" in the PostgreSQL Documentation and check pattern, replacement, and flags for information on the values that can be specified for regexp_pattern, regexp_replacement, and regexp_flags.

- Refer to "6.3 Data Types for Masking" for information on the data types for which masking can be performed.



- When column data type is character(*n*) or char(*n*) and if the string length after change exceeds n, the extra characters will be truncated and only characters up to the nth character will be displayed.
- When column data type is character varying(*n*) or varchar(*n*) and if the string length after change exceeds the length before the change, the extra characters will be truncated and only characters up to the length before change will be displayed.

6.2 Usage Method

Preparation

The following preparation is required to use this feature.

1. Set the postgresql.conf file parameters.

Prepend "pgx_datamasking" to the shared_preload_libraries parameter.

- 2. Restart the instance.
- 3. Execute CREATE EXTENSION for the database that will use this feature.

The target database is described as "postgres" here.

Use the psql command to connect to the "postgres" database.

🞾 Example

```
postgres=# CREATE EXTENSION pgx_datamasking;
CREATE EXTENSION
```



You must always prepend "pgx_datamasking" to the "shared_preload_libraries" parameter.

📶 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

```
G 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: '1=1'

2. Confirm the displayed data

Confirm that the masking target data (column c1) has been correctly changed.

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



=9

- Refer to "B.3.2 pgx_create_confidential_policy" for information on the pgx_create_confidential_policy system management function.



- Only one masking policy can be created per table.
- All users can view the masking policy created, so do not grant the login privilege of the database where this feature is set to the users who refer to the changed data. Masking policies are defined in the "pgx_confidential_columns", "pgx_confidential_policies" and "pgx_confidential_values" tables.

6.2.2 Changing a Masking Policy

- 1. An example of the operation on the server is shown below.
- 2. Change a masking policy

Execute the pgx_alter_confidential_policy system management function to change a masking policy. The following values are changed in this example.

- Content of change: Add a masking target
- Masking target: Character type c2
- Masking type: PARTIAL
- Masking condition: 'VVVFVVVVFVVVV, VVV-VVVV, *, 4, 11'

3. Confirm the displayed data

Confirm that the masking target data has been correctly changed.

💦 See

- Refer to "B.3.1 pgx_alter_confidential_policy" for information on the pgx_alter_confidential_policy system management function.

6.2.3 Confirming a Masking Policy

An example of the operation on the server is shown below.

 Confirm information about a masking target where a masking policy is set Refer to the pgx_confidential_columns table to confirm the masking target where the masking policy is set.

```
postgres=# select * from pgx_confidential_columns;
schema_name | table_name | policy_name | column_name | function_type |
                   | regexp_pattern | regexp_replacement | regexp_flags |
function_parameters
column description
                  _____
+-----
public
                         c1
                                   FULL
       | t1
                | p1
```

```
        public
        | t1
        | p1
        | c2
        | PARTIAL
        | VVVFVVVVFVVVV, VVV-VVV-

        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=# se | lect * from p | gx_confidentia | l_policies; | | |
|---------------|---------------|----------------|-------------|----------|--------------------|
| schema_name | table_name | policy_name | expression | enable | policy_description |
| public | + | + p1 | + | + t | + |
| (1 row) | | | | | |

💦 See

- 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.

2. Confirm the displayed data

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

3. Enable a masking policy

Execute the pgx_enable_confidential_policy system management function to enable a masking policy.

4. Confirm the displayed data

Confirm that the masking target data has been correctly changed.

```
0 | 012-***-***
(3 row)
```



- Refer to "B.3.4 pgx_enable_confidential_policy" for information on the pgx_enable_confidential_policy system management function.

6.2.5 Deleting a Masking Policy

An example of the operation on the server is shown below.

1. Delete a masking policy

Execute the pgx_drop_confidential_policy system management function to delete a masking policy.

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

2. Confirm the displayed data

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

🝓 See

- Refer to "B.3.3 pgx_drop_confidential_policy" for information on the gx_drop_confidential_policy function.

6.3 Data Types for Masking

The data types for which data masking can be performed are shown below.

| Category | Data type | | | |
|--------------|------------------|--------------|-----------------|----------------------------------|
| | | Full masking | Partial masking | Regular expression masking |
| Numeric type | smallint | Υ | Y | Ν |
| | integer | Υ | Y | Ν |
| | bigint | Υ | Y | Ν |
| | decimal | Υ | Y | Ν |
| | numeric | Υ | Y | Ν |
| | float | Y Y | | Ν |
| | real | Y | Y | Ν |
| | double precision | Υ | Y | Ν |

| Category | Data type | Masking type | | | |
|---------------------|-------------------------------|--------------|-----------------|----------------------------------|--|
| | | Full masking | Partial masking | Regular expression masking | |
| Character type | character varying(<i>n</i>) | Υ | 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 | Ν | |
| | timestamp | Y | Y | Ν | |

6.4 Security Notes

- The logical replication is available which allows non-backed up clusters to subscribe to databases where data masking policies are enabled. Logical replication allows publisher and subscriber databases to have their own or the same data masking policies.

In this scenario, the user must disable data masking on the publisher database whenever a subscription is created. This ensures that subscribers are able to obtain the original data (initial copy) instead of the masked version. Then, it is the user's responsibility to set masking policies to each subscribed database.

- Take strong caution in publishing data masking's confidential tables (pgx_confidential_policies, pgx_confidential_columns, etc.) unless the user is publishing all tables of the database and wants to apply the same data masking's policies on the subscribed database for all of them.

Otherwise, as these confidential tables contain the masking policies for all tables of the database, confidential policies of unpublished tables may be unintentionally published. Additionally, it is not possible to apply different data masking policies on the subscriber database.

Chapter 7 Periodic Operations

This chapter describes the operations that must be performed periodically when running daily database jobs.

7.1 Configuring and Monitoring the Log

FUJITSU Enterprise Postgres enables you to output database errors and warnings to a log file.

This information is useful for identifying if errors have occurred and the causes of those errors.

By default, this information is output to the system log. It is recommended that you configure FUJITSU Enterprise Postgres to collect logs from its log files (for example, log_destination) before operating FUJITSU Enterprise Postgres.

Periodically monitor the log files to check if any errors have occurred.



- Refer to "Error Reporting and Logging" under "Server Administration" in the PostgreSQL Documentation for information on logs.

.....

- Refer to "Configuring Parameters" in the Installation and Setup Guide for Server for information on log settings when operating with WebAdmin.

7.2 Monitoring Disk Usage and Securing Free Space

When a database is used for an extended period, free space on the disk is continuously consumed and in some cases the disk space runs out. When this happens, database jobs may stop and no longer run.

You should, therefore, periodically monitor the usage of disk space, and delete obsolete files located in the disk.

Monitor the disk usage of the disk where the following directories are located:

- Data storage destination directory
- Transaction log storage destination (if the transaction log is stored in a different directory from the data storage destination directory)
- Backup data storage destination directory
- Tablespace storage destination directory

7.2.1 Monitoring Disk Usage

To check the disk usage, use the following operating system commands:

- df command

You can even use SQL statements to check tables and indexes individually.

Refer to "Determining Disk Usage" under "Server Administration" in the PostgreSQL Documentation for information on this method.

🔟 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. |

G 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.



Refer to "Connection Settings" under "Server Administration" in the PostgreSQL Documentation for information on the parameters.

7.4 Monitoring the Connection State of an Application

FUJITSU Enterprise Postgres does not immediately delete the updated or deleted data. If the VACUUM determines there are no transactions that reference the database, FUJITSU Enterprise Postgres collects obsolete data.

However, obsolete data is not collected if there are connections that have remained active for an extended period or connections occupying resources. In this case the database may expand, causing performance degradation.



Refer to "Routine Vacuuming" under "Server Administration" in the PostgreSQL Documentation for information on the VACUUM command.

.....

In such cases, you can minimize performance degradation of the database by monitoring problematic connections.

The following method is supported for monitoring connections that have been in the waiting status for an extended period:

- 7.4.1 Using the View (pg_stat_activity)

7.4.1 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 > | | | | | |
|--|---|--|--|--|--|--|
| <pre>cast(query_start + interval '60 minutes' as timestamp);</pre> | | | | | | |
| -[RECORD 1]+ | -[RECORD 1]+ | | | | | |
| datid | 13003 | | | | | |
| datname | db01 | | | | | |
| pid | 4638 | | | | | |
| usesysid | 10 | | | | | |
| usename | fsep | | | | | |
| application_name | apl01 | | | | | |
| client_addr | 192.33.44.15 | | | | | |
| client_hostname | | | | | | |
| client_port | 27500 | | | | | |
| backend_start | 2018-02-24 09:09:21.730641+09 | | | | | |
| xact_start | 2018-02-24 09:09:23.858727+09 | | | | | |
| query_start | 2018-02-24 09:09:23.858727+09 | | | | | |
| state_change | 2018-02-24 09:09:23.858834+09 | | | | | |
| wait_event_type | Client | | | | | |
| wait_event | ClientRead | | | | | |
| state | idle in transaction | | | | | |
| backend_xid | | | | | | |
| backend_xmin | | | | | | |
| query | begin; | | | | | |
| backend_type | client backend | | | | | |



- Refer to "Notes on Application Compatibility" in the Application Development Guide for information on maintaining application compatibility.
- Refer to "The Statistics Collector" under "Server Administration" in the PostgreSQL Documentation for information on pg_stat_activity.

7.5 Reorganizing Indexes

Normally, a database defines indexes in tables, but if data is frequently updated, indexes can no longer use free space in the disk efficiently. This situation can also cause a gradual decline in database access performance.

To rearrange used space on the disk and prevent the database access performance from declining, it is recommended that you periodically execute the REINDEX command to reorganize indexes.

Check the disk usage of the data storage destination using the method described in "7.2 Monitoring Disk Usage and Securing Free Space".



Because the REINDEX command retrieves the exclusive lock for an index being processed and locks writing of tables that are the source of the index, other processes that access these may stop while waiting to be locked.

Therefore, it is necessary to consider measures such as executing the command after the task is completed.



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

関 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 pg_catalog | pgbench_accounts_pkey pg_depend_depender_index | 0.6 | 2208 КВ 224 КВ | 13 MB 368 КВ |
| pg_catalog | pg_depend_reference_index | 0.58 | 216 КВ | 368 KB |



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

7.6 Monitoring Database Activity

FUJITSU Enterprise Postgres enables you to collect information related to database activity. By monitoring this information, you can check changes in the database status.

This information includes wait information for resources such as internal locks, and is useful for detecting performance bottlenecks. Furthermore, you should collect this information in case you need to request Fujitsu technical support for an investigation.

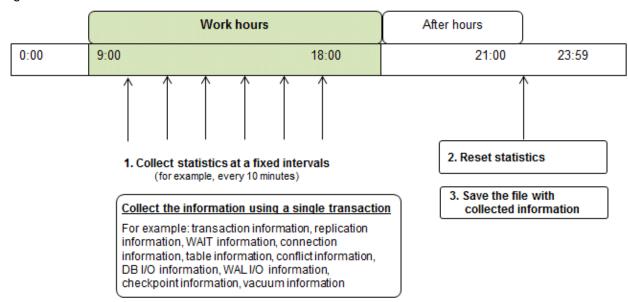


Figure 7.1 Overview of information collection

1. Collect statistics at fixed intervals during work hours.

Accumulate the collected information into a file.

Wherever possible, collect data from the various statistics views using a single transaction, because it enables you to take a snapshot of system performance at a given moment.

Refer to "7.6.1 Information that can be Collected" for information on the system views that can be collected.

2. Reset statistics after work hours, that is, after jobs have finished.

Refer to "7.6.3 Information Reset" for information on how to reset statistics.

3. Save the file with collected information.

Keep the file with collected information for at least two days, in order to check daily changes in performance and to ensure that the information is not deleted until you have sent a query to Fujitsu technical support.

Where jobs run 24 hours a day, reset statistics and save the file with collected information when the workload is low, for example, at night.



Statistics cumulatively add the daily database value, so if you do not reset them, the values will exceed the upper limit, and therefore will not provide accurate information.

The subsections below explain the following:

- Information that can be collected
- Collection configuration
- Information reset

7.6.1 Information that can be Collected

Information that can be collected is categorized into the following types:

- Information common to PostgreSQL
- Information added by FUJITSU Enterprise Postgres

Information common to PostgreSQL

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. Displays statistics related to WAL writing, in a single line. pgx_stat_walwriter 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. Displays statistics related to Global Meta Cache hit ration and used memory size. pgx_stat_gmc

information on the Global Meta Cache.

7.6.2 Collection Configuration

The procedure for configuring collection depends on the information content.

- Information common to PostgreSQL
- Information added by FUJITSU Enterprise Postgres

Information common to PostgreSQL



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

Refer to "C.6 pgx_stat_gmc" for detail. Also refer to Chapter 12 Global Meta Cache" for

Information added by FUJITSU Enterprise Postgres

Information added by FUJITSU Enterprise Postgres is collected by default.

To enable or disable information collection, change the configuration parameters in postgresql.conf. The following table lists the views for which you can enable or disable information collection, and the configuration parameters.

| View name | Parameter |
|-----------------|------------------|
| pgx_stat_lwlock | track_waits (*1) |
| pgx_stat_latch | |
| pgx_stat_sql | track_sql |
| pgx_stat_gmc | track_gmc |

Remarks: You cannot change the collection status for pgx_stat_walwriter.

*1: When executing the SQL statement with EXPLAIN ANALYZE, processing time may increase because of this information collection. It is recommended to set this parameter to "off" when executing EXPLAIN ANALYZE to check the processing time.

Refer to "Appendix A Parameters" for information on the parameters.

7.6.3 Information Reset

This section describes how to reset information.

Information added by FUJITSU Enterprise Postgres

You can reset information added by FUJITSU Enterprise Postgres by using the pg_stat_reset_shared function in the same way as for information common to PostgreSQL.

Configure the following parameters in the pg_stat_reset_shared function:

| Function | Type of return value | Description |
|----------------------------|----------------------|---|
| pg_stat_reset_shared(text) | void | Reset some cluster-wide statistics counters to zero, depending on the argument (requires superuser privileges). |
| | | Calling pg_stat_reset_shared('lwlock') will zero all counters shown in pgx_stat_lwlock. |
| | | Similarly, in the following cases, all values of the pertinent statistics counter are reset: |
| | | - If pg_stat_reset_shared('latch') is called: |
| | | All values displayed in pgx_stat_latch |
| | | - If pg_stat_reset_shared('walwriter') is called: |
| | | All values displayed in pgx_stat_walwriter |
| | | - If pg_stat_reset_shared('sql') is called: |
| | | All values displayed in pgx_stat_sql |
| | | - If pg_stat_reset_shared('gmc') is called: |
| | | All values except size column in pgx_stat_gmc |



Refer to "Statistics Functions" in "Monitoring Database Activity" under "Server Administration" in the PostgreSQL Documentation for information on other parameters of the pg_stat_reset_shared function.

Chapter 8 Streaming Replication Using WebAdmin

This chapter describes how to create a streaming replication cluster using WebAdmin.

Streaming replication allows the creation of one or more standby instances, which connect to the master instances and replicate the data using WAL records. The standby instance can be used for read-only operations.

WebAdmin can be used to create a streaming replication cluster. WebAdmin allows the creation of a cluster in the following configurations:

- Master-Standby Configuration: This configuration creates a master and standby instance together.
- Standby Only Configuration: This configuration creates a standby instance from an already existing instance.

関 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.

| FUJITSU Enterprise Po | stgres | | English ∽ fsepuser ∽ FUĴITSU | |
|-----------------------|-------------------------------|---|--|--|
| Instances | Create instance | | ✓ Ø | |
| My Wallet | | | пасаез тециней нейцэл | |
| Settings | Configuration type | Standby only configuration | | |
| | Server product type | FUJITSU Enterprise Postgres 12 | | |
| | Existing master configuration | | | |
| | Instance name | inst1 | | |
| | Host name | osboxes | | |
| | Port | 27500 | | |
| | New standby configuration | | | |
| | Location | Local | | |
| | Replication credential * | Replication credential Cred1 User name webadminser | Password | |
| | Instance name * | inst1s | | |
| | Instance port * | 27502 | | |
| | Host IP address * | 192.168.56.101 | | |
| | Data storage path * | /home/fsepuser/inst1s/data | | |
| | Backup | Enabled V | Backup storage path * /home/fsepuser/inst1s/backup | |
| | Transaction log path | /home/fsepuser/inst1s/transactionlog | | |
| | Encoding | UTF8 V | | |
| | Replication mode | Asynchronous V | | |
| | Application name * | hacluster | | |

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 \checkmark 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:

| E FUJITSU Enterprise | Postgres | | | | | English 🗸 | fsepuser 🗸 | FUJITSU |
|----------------------|--------------------------------|-----------------------|---|-------|--------------|-----------|----------------|---------|
| Instances | + <u>+</u> 🖉 🗉 | inst1s | | F. F. | < | □ 2 | 0 0 | æ |
| My Wallet | Last refreshed 4 minutes ago . | ✓ Summary | | | | | | * |
| Settings | inst1 osboxes:27500 | Title | inst1s (osboxes:27502) | | | | | |
| | inst1s osboxes:27502 | Server type | FUJITSU Enterprise Postgres AE Linux 64-bit | | | | | |
| | | Version | PostgreSQL 12.1 | | | | | |
| | | Server up time | 43 seconds | | | inst1 | | |
| | | Encoding | UTF8 | | | | | |
| | | Instance type | Asynchronous Standby | | | | | |
| | | Last WAL applied time | | | | inst1s | | |
| | | Application name | hacluster | | | | | |
| | | ✓ Storage summary | | | | | | . |
| | | Data storage status | Normal | | | | | |
| | | Data storage path | /home/fsepuser/inst1s/data | | | | \bigcirc | |
| | | Backup storage status | - | | | | | |
| | | Backup storage path | /home/fsepuser/inst1s/backup/backup | | Data storage | | Backup storage | |
| | | Backup time | - | | | | | |

🕝 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".

- 55 -

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 1.
- 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 **Co**.

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 \mathbf{O} .
- 6. Select the standby instance. [Instance type] will now show the updated status.

🌀 Note

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.
- 8. Select [Restart later] or [Restart now], and then click [Yes] to set up the standby instance.
- 9. Upon successful completion, the confirmation dialog box will be displayed.
- 10. Click [Close] to return to the instance details window.

The instance will become a standby instance, and will be part of the streaming replication cluster. The replication diagram will display the relationship between the standby instance and the remote instance. The user can change the replication relationship of the remote instance from asynchronous to synchronous (and vice versa) using the [Configuration] window.

Chapter 9 Installing and Operating the In-memory Feature

The in-memory feature enables fast aggregation using Vertical Clustered Index (VCI) and memory-resident feature.

VCI has a data structure suitable for aggregation, and features parallel scan and disk compression, which enable faster aggregation through reduced disk I/O.

The memory-resident feature reduces disk I/O that occurs during aggregation. It consists of the preload feature that reads VCI data to memory in advance, and the stable buffer feature that suppresses VCI data eviction from memory. The stable buffer feature secures the proportion specified by parameter in the shared memory for VCI.

This chapter describes how to install and operate the in-memory feature.

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

G Note

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.

💦 See

- 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 |
| 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. | Ν |

💕 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
```

.....



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

.



- Refer to "Appendix A Parameters" for information on all parameters for VCI. Refer also to default value for each parameter and details such as specification range in the same chapter. Refer to "Server Configuration" under "Server Administration" in the PostgreSQL documentation for information on shared_preload_libraries, session_preload_libraries, and max_worker_processes.

.

9.1.3.2 Installing the Extensions

Execute CREATE EXTENSION to install the VCI and pg_prewarm extensions. Both extensions need to be installed for each database.

- Installing VCI

db01=# CREATE EXTENSION vci;

- Installing pg_prewarm

db01=# CREATE EXTENSION pg_prewarm;

G 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);

G 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

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 |
| | polygon | N |
| | circle | N |
| Network address | cidr | N |
| | inet | N |
| | macaddr | N |
| | macaddr8 | N |

| Category | Data type | Supported by VCI? |
|---|-------------------------|-------------------|
| Bit string | bit(<i>n</i>) | Y |
| | bit varying(<i>n</i>) | Y |
| Text search | tsvector | Ν |
| | tsquery | N |
| UUID | uuid | Y |
| XML | xml | N |
| JSON | json | N |
| | jbson | N |
| Range | int4range | N |
| | int8range | N |
| | numrange | N |
| | tsrange | N |
| | tstzrange | Ν |
| | daterange | N |
| Object identifier | oid | N |
| | regproc | N |
| | regprocedure | N |
| | regoper | N |
| | regoperator | N |
| | regclass | N |
| | regtype | N |
| | regconfig | Ν |
| | 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 | |
| | SET TABLESPACE | if a VCI is specified. If the operation is required, delete the VCI using DROP INDEX, and re-create it using CREATE INDEX after completing the operation. | |
| | ALL IN TABLESPACE | | |
| ALTER OPERATOR | RENAME TO | The subcommands on the left cannot be performed | |
| CLASS | OWNER TO | if a VCI is specified. | |
| | SET SCHEMA | This operation is not supported in VCI. | |
| ALTER OPERATOR | ADD | | |
| FAMILY | DROP | | |
| | RENAME TO | | |
| | OWNER TO | | |
| | SET SCHEMA | | |
| ALTER TABLE | ALL IN TABLESPACE <i>name</i> [OWNED BY <i>roleName</i>] SET | A tablespace that contains a VCI cannot be specified. | |
| | TABLESPACE newTablespace | If the operation is required, delete the VCI using DROP INDEX, and re-create it using CREATE INDEX after completing the operation. | |
| | DROP [COLUMN] [IF EXISTS] colName [RESTRICT CASCADE] | A column that contains a VCI cannot be specified. If the operation is required, delete the VCI using | |
| | ALTER [COLUMN] <i>colName</i> [SET DATA] TYPE <i>dataType</i> [COLLATE <i>collation</i>] [USING <i>expr</i>] | DROP INDEX, and re-create it using CREAT INDEX after completing the operation. | |
| | CLUSTER ON indexName | A VCI cannot be specified. | |
| | REPLICA IDENTITY {DEFAULT USING INDEX <i>indexName</i> FULL NOTHING} | This operation is not supported in VCI. | |
| CLUSTER | - | A table that contains a VCI and VCI cannot be specified. | |
| | | If the operation is required, delete the VCI using DROP INDEX, and re-create it using CREATE INDEX after completing the operation. | |
| CREATE INDEX | UNIQUE | The subcommands on the left cannot be performed | |
| | CONCURRENTLY | if a VCI is specified. | |
| | | This operation is not supported in VCI. | |

| Command | Subcommand | Description |
|---------------------------|---------------------------|--|
| | [ASC DESC] | |
| | [NULLS { FIRST LAST }] | |
| | WITH | |
| | WHERE | |
| | INCLUDE | |
| CREATE OPERATOR CLASS | - | A VCI cannot be specified. |
| CREATE OPERATOR FAMILY | - | This operation is not supported in VCI. |
| CREATE TABLE | EXCLUDE | |
| DROP INDEX | CONCURRENTLY | The subcommands on the left cannot be performed if a VCI is specified. |
| | | This operation is not supported in VCI. |
| REINDEX | - | A VCI cannot be specified. |
| | | This command is not required as VCI uses daemon's automatic maintenance to prevent disk space from increasing. |

Client application command

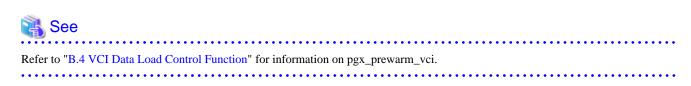
- Operations that cannot be performed on relations containing a VCI

| Command | Description |
|-----------|---|
| clusterdb | Clustering cannot be performed for tables that contain a VCI. |
| reindexdb | VCIs cannot be specified on theindex option. |

9.2.2 Data Preload Feature

The first aggregation using VCI immediately after an instance is started may take time, because the VCI data has not been loaded to buffer. Therefore, use the preload feature to load the VCI data to buffer in advance when performing VCI aggregation after an instance is started. When using the preload feature, execute the function pgx_prewarm_vci to each VCI created with CREATE INDEX.

db01=# SELECT pgx_prewarm_vci('idx_vci');



Chapter 10 Parallel Query

FUJITSU Enterprise Postgres enhances parallel queries, by taking into consideration the aspects below:

- CPU load calculation
- Increase of workers during runtime

10.1 CPU Load Calculation

There may be a case when the user tries to execute a parallel query but there is not enough CPU available.

Adding dynamic workers at this stage will provide no benefits - instead, it may add overhead due to context switching.

FUJITSU Enterprise Postgres takes into consideration the current CPU load when deciding on the number of workers for parallel query.

10.2 Increase of Workers during Runtime

This FUJITSU Enterprise Postgres enhancement allows systems to allocate additional workers during query execution (if there are workers available at the time). This improves query performance, which could otherwise starve of CPU if there were fewer or no workers when the query started.



The ability to increase workers during runtime is available only with parallel sequence scan.

Chapter 11 High-Speed Data Load

High-speed data load uses the pgx_loader command to load data from files at high speed into FUJITSU Enterprise Postgres.



This feature is not available in single-user mode. This is because in single-user mode instances run in a single process, and it cannot start parallel workers.

11.1 Installing High-Speed Data Load

This section describes how to install high-speed data load.

Installation flow

- 1. Deciding whether to Install
- 2. Estimating Resources
- 3. Setup

11.1.1 Deciding whether to Install

The feature achieves high speed data load by executing the COPY FROM command in parallel. If the database system is unable to use sufficient resources due to the feature using more resources than the COPY FROM command of PostgreSQL, load performance may be inferior to that of the COPY FROM command of PostgreSQL. Therefore, determine if the feature will be effective by considering the factors below before deciding to install.

Database server memory

If the value of shared_buffers in postgresql.conf is small, fewer data pages are cached to the shared memory of the database server. This will result in multiple parallel workers more often having to wait for write exclusive locks to the same data page. Moreover, the smaller the number of data pages, the more often the table expands. During table expansion, access to the table is exclusive (standby event name: extend), so write time increases. To cater for that, increase the value of shared_buffers.



The standby event name is stored in the wait_event column of the pg_stat_activity view. Refer to "wait_event Description" in "The Statistics Collector" in the PostgreSQL Documentation for details.

Frequency of checkpoints

If checkpoints are issued at short intervals, write performance is reduced. If the messages below are output to the server log during data writes, increase the values of max_wal_size and checkpoint_timeout in postgresql.conf to reduce the frequency of checkpoints.

💕 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.

G 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 |
|---------------------------|---|-----------|
| max_prepared_transactions | Add the number of transactions that can be prepared by parallel workers during data load to the parameter's current value. | Mandatory |
| | The resulting value must be equal to or greater than the maximum number of parallel workers used with this feature. | |
| max_worker_processes | Number of parallel workers to perform data load. | Mandatory |
| max_parallel_workers | Add the maximum number of parallel workers to be used in a parallel query by this feature to the parameter's current value. | Mandatory |
| | The resulting value must be equal to or greater than the number of parallel workers used with this feature. | |

💕 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
```



As shown in the example above, set the value of max_prepared_transactions, max_worker_processes and max_parallel_workers multiplied by the number of instances of this feature executed simultaneously.

The table below lists the postgresql.conf parameters that must also be checked.

| Parameter | Setting | Required |
|----------------------------|---|-----------|
| dynamic_shared_memory_type | Implementation of dynamic shared memory to be used by the instance. | Mandatory |
| | The default value is recommended. | |



```
Refer to "Resource Consumption" in the PostgreSQL Documentation for information on the parameters.
```

11.1.3.2 Installing the Extension

Execute CREATE EXTENSION to install the high-speed data load extension. The extension needs to be installed on each database.

💕 Example

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

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



- Only superusers can install the high-speed data load extension.
- The high-speed data load extension can only be installed on the public schema.

11.2 Using High-Speed Data Load

This section describes how to use high-speed data load.

11.2.1 Loading Data

To load data from a file into a FUJITSU Enterprise Postgres table, execute the pgx_loader command in load mode.

💕 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
```



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.

.



.

- Refer to "pgx_loader" in the Reference for information on the command.
- Refer to "COPY" in the PostgreSQL Documentation for information on the deployment destination and access privileges for external files.

11.2.2 Recovering from a Data Load that Ended Abnormally

If a system interruption such as a server failure occurs while high-speed data load is being performed, transactions prepared using this feature may be changed to the in-doubt state. At that point, resources occupied by the transaction will be locked, and access to the relevant resources from other transactions will be blocked, rendering them unusable.

In such cases, check transactions that are in an in-doubt state, and resolve them.

Checking for in-doubt transactions

This section describes how to check for in-doubt transactions.

1. Refer to the pgx_loader_state table in the pgx_loader schema.

Retrieve the global transaction identifier (gid column) of in-doubt transactions. In-doubt transactions will contain "rollback" in the column "state".

💕 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.

2. Refer to the pg_prepared_xacts system view.

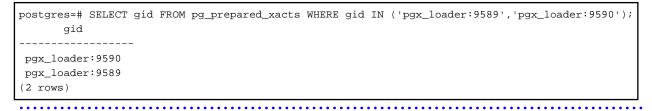
Check if the in-doubt transactions retrieved above exist.



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

.

.





Refer to "E.1 pgx_loader_state" for information on the pgx_loader_state table.

Resolving in-doubt transactions

Execute the pgx_loader command in recovery mode to resolve in-doubt transactions.

After executing the pgx_loader command in recovery mode, perform the procedure described in "Checking for in-doubt transactions" to check if the in-doubt transactions have been resolved.

💕 Example

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

\$ pgx_loader recovery -t tbl



The recovery mode of the pgx_loader command only resolves transactions prepared by high-speed data load. For transactions prepared by an application using distributed transactions other than this feature, follow the procedure described in "14.13 Actions in Response to Error in a Distributed Transaction".

11.3 Removing High-Speed Data Load

This section describes how to remove high-speed data load.

11.3.1 Removing the Extension

Execute DROP EXTENSION to remove the high-speed data load extension. The extension needs to be removed on each database.

💕 Example

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

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



- The information required for operation of high-speed data load is stored in the pgx_loader_state table of the pgx_loader schema. Do not remove the high-speed data load extension if the pgx_loader_state table is not empty.

- Only superusers can remove the high-speed data load extension.

- The high-speed data load extension can only be removed on the public schema.

Chapter 12 Global Meta Cache

The Global Meta Cache (GMC) feature loads a meta cache into shared memory using the pgx_global_metacache parameter. This reduces the amount of memory required throughout the system.

12.1 Usage

Describes how to use the Global Meta Cache feature.

12.1.1 Deciding Whether to Enable the Global Meta Cache Feature

Global Meta Cache is a mechanism for sharing meta caches between processes, so it works well on systems with a high number of resources accessed and SQL connections. The number of resources is primarily the number of tables accessed by a process, the number of indexes, or the total number of all columns in all tables accessed.

In particular, consider using Global Meta Cache if the total size of the meta cache for each process exceeds the amount of installed memory, or takes up a large portion of that memory, thereby squeezing memory allocations to the database cache or the Operating system file cache. Using Global Meta Cache may slightly increase the time it takes to execute a single SQL to reference a meta cache on shared memory (For example, it will not exceed 5% of the execution time in a situation where all the data is located in the database cache), but you can expect a greater benefit from being able to allocate more memory, such as for the database cache.

If performance degradation using Global Meta Cache is not acceptable, you may want to limit the number of tables accessed by a process.

12.1.2 Estimating Memory for Global Meta Cache

To enable the Global Meta Cache feature, the pgx_global_metacache parameter must specify an upper limit on the size of the shared memory (Hereinafter, the GMC area) dedicated to Global Meta Cahche. Ideally, this upper limit should be the size estimated in "Appendix A Parameters". Values lower than this can still work, but refer to "12.1.3 How the GMC Memory Area Is Used" on using the GMC area to understand the disadvantages.

12.1.3 How the GMC Memory Area Is Used

At startup, the memory for the GMC area is not used much, but the GMC area grows as new meta caches are placed in the GMC area. If it does, it discards any meta caches that the system determines are not heavily used and places a new one in the GMC area.

Therefore, the GMC area will work even if it is smaller than the estimate, but the meta cache will be regenerated if the discarded meta cache needs to be reused. Note that if this happens frequently, it will degrade overall performance.

With this in mind, it may not be a problem if, for example, the tables to be accessed are different depending on the time zone, and the degradation of the time zone immediately after the change is acceptable.

In any case, be sure to test and tune the system thoroughly before running it.

12.1.4 Enabling the Global Meta Cache Feature

To Enable the Global Meta Cache feature edit the postgresql.conf file and set the pgx_global_metacahe parameter. Restarting the instance after editing the postgresql.conf file is required. Refer to "Appendix A Parameters" for information on the parameters.

| Parameter Name | Description |
|----------------------|--|
| pgx_global_metacache | Specify the maximum amount of memory for the GMC area on shared memory. When it's set to 0 (default value), the Global Meta Cache feature is disabled. When enabled, the minimum value allowed is 10MB. |

When the cache is created, if the total amount of meta caches on shared memory exceeds the value specified by pgx_global_metacache, the inactive, unreferenced meta caches are removed from the GMC area. Note that if all GMC are in use and the cache cannot be created in the GMC area, the cache is temporarily created in the local memory of the backend process.



Here is an example postgresql.conf configuration:

pgx_global_metacache = 800 MB

Wait Events

The Global Meta Cache feature may cause wait events. Wait events are identified in the wait_event column of the pg_stat_activity view. GMC specific wait events are described below.

[GMC Feature Wait Events]

| Wait Event Type | Wait Event Name | Description |
|-----------------|------------------------|--|
| LWLock | global_metacache | Waiting to find, add, and remove meta caches in the GMC area. |
| | global_metacache_sweep | Waiting to select a meta cache that can be deleted when GMC space is low. If the GMC is fully referencing and there is no deletable meta cache, it is waiting for the reference to be removed and a deletable |
| | | meta cache to be selected. |



If global_metacache_sweep is happened frequently, increase the pgx_global_metacache setting.



Refer to "Viewing Statistics" in the PostgreSQL Documentation for information on the pg_stat_activity view.

12.1.5 Estimating Resources

Refer to "Global Meta Cache Memory Requirements" in the Installation and Setup Guide for Server for formulas to estimate the amount of memory used by the Global Meta Cache feature.

12.2 Statistics

Describes the statistics for the Global Meta Cache feature.

12.2.1 System View

You can check the cache hit ratio and size of the GMC area in the system view pgx_stat_gmc. Refer to "C.6 pgx_stat_gmc" for information on the columns.

If the cache hit ratio is low and the current memory usage is close to pgx_global_metacache, increase the pgx_global_metacache setting because performance may be degraded.

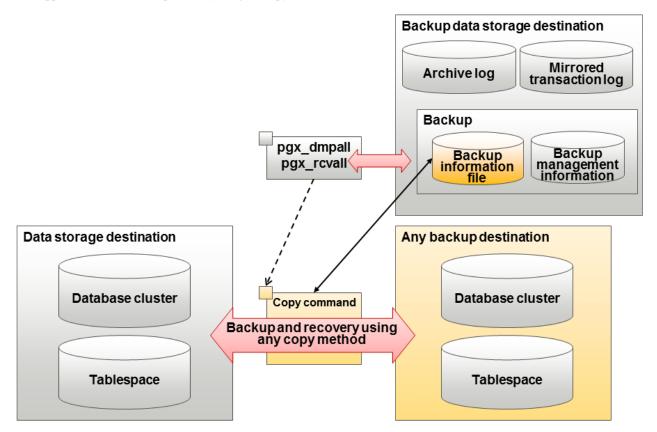
Refer to "7.6 Monitoring Database Activity" in the Operations Guide for information on the statistics.

Chapter 13 Backup/Recovery Using the Copy Command

By using a copy command created by the user, the pgx_dmpall command and the pgx_rcvall command can perform backup to any destination and can perform recovery from any destination using any copy method.

Copy commands must be created in advance as executable scripts for the user to implement the copy process on database clusters and tablespaces, and are called when executing the pgx_dmpall and pgx_rcvall commands.

This appendix describes backup/recovery using the copy command.



関 Point

It is also possible to back up only some tablespaces using the copy command. However, database resources not backed up using the copy command are still backed up to the backup data storage destination.



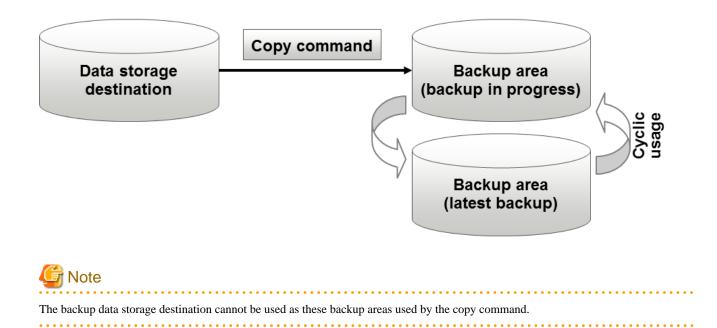
Both the backup data storage destination and the optional backup destination are necessary for recovery - if they are located in secondary media, combined management of these is necessary.

13.1 Configuration of the Copy Command

This section describes the configuration of the copy command for backup and recovery.

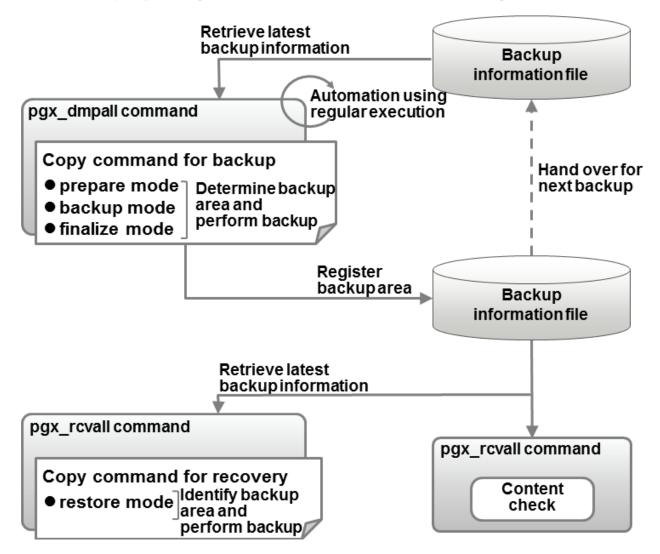
Cyclic usage of the backup area

Prepare two backup areas for the copy command in case an issue affects the data storage destination during backup. The copy command performs backup while cyclically using these backup areas.



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.



The user can use any method to hand over backup information between modes within the copy command, such as creating temporary files.

Configuration of the copy command for recovery

The pgx_rcvall command calls the copy command for recovery for the mode below. It is therefore necessary for the copy command for recovery to implement the required processing for the mode.

- restore mode

Any copy method can be used to implement restore from the backup destination retrieved using the copy command for backup.



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

💕 Example

Using a bash script

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

```
;;
restore)
processingRequiredForRestoreMode
;;
esac
```

関 Point

A sample batch file that backs up the database cluster and tablespace directory to a specific directory is supplied to demonstrate how to write a copy command.

The sample is stored in the directory below:

/installDir/share/copy_command.archive.sh.sample

13.2 Backup Using the Copy Command

To perform backup using the copy command, in addition to performing the standard backup procedure, it is also necessary to create a copy command, and then execute the pgx_dmpall command specifying it. This section describes the procedure specific to using the copy command.

Preparing for backup

You must prepare for backup before actually starting the backup process.

Perform the following procedure:

1. Determine the database resources to be backed up

Determine the database resources to be backed up using the copy command. The copy command can back up the following resources:

- Database cluster
- Tablespace

To back up only some tablespaces, create a file listing them. This file is not necessary to back up all tablespaces.

Example

To back up only tablespaces tblspc1 and tblspc2

| tblspc1 | |
|---------|--|
| tblspc2 | |

2. Prepare a backup area

Prepare a backup area to save the database resources to be backed up, as determined in step 1.

3. Create the copy command

Create the copy commands for backup and recovery. Refer to "13.4 Copy Command Interface" for details.

Performing backup

Execute the pgx_dmpall command with the -Y option specifying the full path of the copy command for backup created in step 3 of preparation for backup.

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

| 1 | Example |
|---|---------|
| | |

| <pre>\$ pgx_dmpall -D /database/inst1 -Y '/database/command/backup.sh'</pre> | |
|--|--|
| exclude-copy-cluster -P '/database/command/tablespace_list.txt' | |

関 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.



```
- 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 -1 -D /database/inst1
Date Status Dir Resources backed up by the copy command
2017-05-01 13:30:40 COMPLETE /backup/inst1/2017-05-01_13-30-40 pg_data,dbspace,indexspace

13.3 Recovery Using the Copy Command

To perform recovery using the copy command, in addition to performing the standard recovery procedure, it is also necessary to create a copy command, and then execute the pgx_rcvall command specifying it. This section describes the procedure specific to using the copy command.

Determining the backup area of the latest backup

Check the backup information file to determine the backup area used for the latest backup, and confirm that it is in a recoverable state.

Execute the pgx_rcvall command with the --view-results-of-copying option to output the content of the backup information file.

\$ pgx_rcvall -D /database/inst1 --view-results-of-copying

Perform recovery

Execute the pgx_rcvall command with the -Y option specifying the full path of the copy command for recovery created in step 3 of the preparation for backup described in "13.2 Backup Using the Copy Command".

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



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

関 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.



Refer to "pgx_rcvall" in the Reference for information on the command.

13.4 Copy Command Interface

The following types of copy command are available:

- Copy command for backup
- Copy command for recovery

This appendix describes the interface of each copy command.

13.4.1 Copy Command for Backup

Feature

User command called from the pgx_dmpall command.

Format

The syntax for calling the copy command from the pgx_dmpall command is described below.

If the operation mode is "prepare"

```
copyCommandName prepare 'pathOfBackupInfoFile' 'pathOfBackupTargetListFile'
```

If the operation mode is "backup"

copyCommandName backup

If the operation mode is "finalize"

copyCommandName finalize 'pathOfBackupInfoFile'

Argument

- Operation mode

| Mode | Description | |
|---------|---|--|
| prepare | Implements the preparation process for backing up using the copy command. | |
| | Called before the PostgreSQL online backup mode is started. | |
| backup | Implements the backup process. | |
| | Called during the PostgreSQL online backup mode. | |

| Mode | Description | |
|----------|--|--|
| finalize | Implements the backup completion process. | |
| | Called after the PostgreSQL online backup mode is completed. | |

- Full path of the backup information file

Full path of the backup information file of the latest backup, enclosed in 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

1 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 pgx_dmpall command continues processing. | |
| Other than 0 | Abnormal end | |
| | The pgx_dmpall command terminates in error. | |

Description

- The copy command operates with the privileges of the operating system user who executed the pgx_dmpall command. Therefore, grant copy command execution privileges to users who will execute the pgx_dmpall command. Additionally, have the copy command change users as necessary.
- To write to the backup information file, use a method such as redirection from the copy command.
- Because the copy command is called for each mode, implement all processing for each one.
- To copy multiple resources simultaneously, have the copy command copy them in parallel.



- The backup information file and backup target list file cannot be deleted. Additionally, the privileges cannot be changed.
- Standard output and standard error output of the copy command are output to the terminal where the pgx_dmpall command was executed.

- If the copy command becomes unresponsive, the pgx_dmpall command will also become unresponsive. If the copy command is deemed to be unresponsive by the operating system, use an operating system command to forcibly stop it.
- Output the copy command execution trace and the result to a temporary file, so that if it terminates in error, the cause can be investigated at a later time.
- For prepare mode only, it is possible to use the PostgreSQL client application to access the database using the copy command. For all other modes, do not execute FUJITSU Enterprise Postgres commands or PostgreSQL applications.
- Enable the fsync parameter in postgresql.conf, because data on the shared memory buffer needs to have been already written to disk when backup starts.

13.4.2 Copy Command for Recovery

Feature

User command called from the pgx_rcvall command.

Format

The syntax for calling the copy command from the pgx_rcvall command is described below.

```
copyCommandName restore 'pathOfBackupInfoFile' 'pathOfBackupTargetListFile'
```

.

Argument

- Operation mode

| Mode | Description |
|---------|-------------------|
| restore | Performs restore. |

- Full path of the backup information file

Full path of the backup information file, enclosed in single quotation marks.

- Full path of the backup target list file

Full path of the file containing the resources to be restored using the copy command, enclosed in single quotation marks.

The access permissions for the backup information file and backup target list file are as below.

| Backup information file | Backup target list file | |
|--|--|--|
| Can be viewed by the instance administrator only | Can be viewed by the instance administrator only | |

Return value

| Return value | Description | |
|--------------|--|--|
| 0 | Normal end | |
| | The pgx_rcvall command continues processing. | |
| Other than 0 | Abnormal end | |
| | The pgx_rcvall command terminates in error. | |

Description

- The copy command operates with the privileges of the operating system user who executed the pgx_rcvall command. Therefore, grant copy command execution privileges to users who will execute the pgx_rcvall command. Additionally, have the copy command change users as necessary.
- The copy command is called once only in restore mode.
- To copy multiple resources simultaneously, have the copy command copy them in parallel.



- The backup information file and backup target list file cannot be deleted. Additionally, the privileges cannot be changed.
- Standard output and standard error output of the copy command are output to the terminal where the pgx_rcvall command was executed.

- If the copy command becomes unresponsive, the pgx_rcvall command will also become unresponsive. If the status of the copy command is deemed to be unresponsive by the operating system, use an operating system command to forcibly stop it.
- Output the copy command execution trace and the result to a temporary file, so that if it terminates in error, the cause can be investigated at a later time.
- Do not execute FUJITSU Enterprise Postgres commands or PostgreSQL applications in the copy command.
- There may be files and directories not required for recovery using the archive log included in the backup, such as postmaster.pid, pg_wal/*subdirectory* and pg_replslot in the database cluster. If such unnecessary files and directories exist, have the copy command delete them after the restore.

- 83 -

Chapter 14 Actions when an Error Occurs

This chapter describes the actions to take when an error occurs in the database or an application, while FUJITSU Enterprise Postgres is operating.

Depending on the type of error, it may be necessary to recover the database cluster. The recovery process recovers the following resources:

- Data storage destination
- Transaction log storage destination (if the transaction log is stored in a separate disk from the data storage destination)
- Backup data storage destination

🌀 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 "14.1 Recovering from Disk Failure (Hardware)", and take actions accordingly.

- If the disk is not defective

Refer to "14.14 I/O Errors Other than Disk Failure", and take actions accordingly.

A few examples of errors generated even if the disk is not defective include:

- Network error with an external disk
- Errors caused by power failure or mounting issues

Determining the cause of an error

If an error occurs, refer to the WebAdmin message and the server log, and determine the cause of the error.



Approximate recovery time

The formulas for deriving the approximate recovery time of resources in each directory are given below.

If using the copy command with the pgx_rcvall command, the recovery time will depend on the implementation of the copy command.

- Data storage destination or transaction log storage destination

Recovery time = (usageByTheDataStorageDestination + usageByTheTransactionLogStorageDestination) / diskWritePerformance x 1.5

- usageByTheDataStorageDestination: Disk space used by the database cluster
- usageByTheTransactionLogStorageDestination: Disk space used by the transaction log stored outside the database cluster
- *diskWritePerformance*: Measured maximum data volume (bytes/second) that can be written per second in the system environment where the operation is performed
- 1.5: Coefficient assuming the time excluding disk write, which is the most time-consuming step

- Backup data storage destination

Recovery time = usageByTheBackupDataStorageDestination / diskWritePerformance x 1.5

- usageByTheBackupDataStorageDestination: Disk space used by the backup data
- *diskWritePerformance*: Measured maximum data volume (bytes/second) that can be written per second in the system environment where the operation is performed
- 1.5: Coefficient assuming the time excluding disk write, which is the most time-consuming step

14.1 Recovering from Disk Failure (Hardware)

This section describes how to recover database clusters to a point immediately before failure, if a hardware failure occurs in the data storage disk or the backup data storage disk.

There are two methods of recovery:

- 14.1.1 Using WebAdmin
- 14.1.2 Using Server Command



Back up the database cluster after recovering it. Backup deletes obsolete archive logs (transaction logs copied to the backup data storage destination), freeing up disk space and reducing the recovery time.

14.1.1 Using WebAdmin

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



Recovery operation cannot be performed on an instance that is part of a streaming replication cluster in standby mode.

If disk failure occurs on a standby instance, it may be necessary to delete and re-create the instance.

Recovery operation can be performed on an instance that is part of a streaming replication cluster in "Master" mode. If a recovery operation is performed on a master instance, it will break the replication cluster and streaming replication will stop between the master instance and all its standby instances. In such an event, the standby instances can be promoted to standalone instances or can be deleted and re-created.

If failure occurred in the data storage disk or the transaction log storage disk

Follow the procedure below to recover the data storage disk or the transaction log storage disk.

1. Stop applications

Stop applications that are using the database.

2. Stop the instance

Stop the instance. Refer to "2.1.1 Using WebAdmin" for information on how to stop an instance. WebAdmin automatically stops instances if recovery of the database cluster is performed without stopping the instance.

3. Recover the failed disk

Replace the disk, and then recover the volume configuration information.

4. Create a tablespace directory

If a tablespace was defined after backup, create a directory for it.

5. Recover the keystore, and enable automatic opening of the keystore

Do the following if the data in the database has been encrypted:

- Restore the keystore to its state at the time of the database backup.
- Enable automatic opening of the keystore.

6. Recover the database cluster

Log in to WebAdmin, and in the [Instances] tab, click [Solution] for the error message in the lower-right corner.

7. Run recovery

In the [Restore Instance] dialog box, click [Yes].

Instance restore is performed. An instance is automatically started when recovery is successful.



WebAdmin does not support recovery of hash index. If you are using a hash index, then after recovery, execute the REINDEX command to rebuild it. Use of hash indexes is not recommended.

8. Resume applications

Resume applications that are using the database.



WebAdmin may be unable to detect disk errors, depending on how the error occurred.

If this happens, refer to "14.10.3 Other Errors" to perform recovery.

If failure occurred on the backup data storage disk

Follow the procedure below to recover the backup data storage disk.

1. Recover the failed disk

Replace the disk, and then recover the volume configuration information.

2. Recover the backup data

Log in to WebAdmin, and in the [Instances] tab, click [Solution] for the error message.

3. Run backup

Perform backup to enable recovery of the backup data. In the [Backup] dialog box, click [Yes]. The backup is performed. An instance is automatically started when backup is performed.

関 Point

If you click [Recheck the status], the resources in the data storage destination and the backup data storage destination are reconfirmed. As a result, the following occurs:

- If an error is not detected

The status of the data storage destination and the backup data storage destination returns to normal, and it is possible to perform operations as usual.

- If an error is detected

An error message is displayed in the message list again. Click [Solution], and resolve the problem by following the resolution for the cause of the error displayed in the dialog box.

14.1.2 Using Server Command

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

If failure occurred on the data storage disk or the transaction log storage directory

Follow the procedure below to recover the data storage disk or the transaction log storage directory.

1. Stop applications

Stop applications that are using the database.

2. Stop the instance

Stop the instance, refer to "2.1.2 Using Server Commands" for details.

If the instance fails to stop, refer to "14.11 Actions in Response to Failure to Stop an Instance".

3. Recover the failed disk

Replace the disk, and then recover the volume configuration information.

- 4. Create a storage destination directory
 - If failure occurred on the data storage disk Create a data storage destination directory. If a tablespace was defined, also create a directory for it.
 - If failure occurred on the translation log storage disk Create a transaction log storage destination directory.

Example

To create a data storage destination directory:

```
$ mkdir /database/inst1
$ chown fsepuser:fsepuser /database/inst1
$ chmod 700 /database/inst1
```

🐴 See

Refer to "Preparing Directories to Deploy Resources" under "Setup" in the Installation and Setup Guide for Server for information on how to create a storage directory.

5. Recover the keystore, and enable automatic opening of the keystore

When the data in the database has been encrypted, restore the keystore to its state at the time of the database backup. Configure automatic opening of the keystore as necessary.

6. Recover the database cluster

Recover the database cluster using the backup data.

Specify the following in the pgx_rcvall command:

- Specify the data storage location in the -D option. If the -D option is omitted, the value of the PGDATA environment variable is used by default.
- Specify the backup data storage location in the -B option.

Example

> pgx_rcvall -D /database/inst1 -B /backup/inst1



If recovery fails, remove the cause of the error in accordance with the displayed error message and then re-execute the pgx_rcvall command.

If the message "pgx_rcvall: an error occurred during recovery" is displayed, then the log recorded when recovery was executed is output after this message. The cause of the error is output in around the last fifteen lines of the log, so remove the cause of the error in accordance with the message and then re-execute the pgx_rcvall command.

The following message displayed during recovery is output as part of normal operation of pgx_rcvall command (therefore the user does not need not be concerned).

FATAL: The database system is starting

7. Start the instance

Start the instance.

Refer to "2.1.2 Using Server Commands" for information on how to start an instance.

8. Resume applications

Resume applications that are using the database.

If failure occurred on the backup data storage disk

The procedure for recovering the backup data storage disk is described below.

There are two methods of taking action:

- Performing recovery while the instance is active
- Stopping the instance before performing recovery

The following table shows the different steps to be performed depending on whether you stop the instance.

| No | Step | Instance stopped | | |
|----|--|------------------|-----|--|
| | | No | Yes | |
| 1 | Confirm that transaction log mirroring has stopped | Y | Ν | |
| 2 | Stop output of archive logs | Y | Ν | |
| 3 | Stop applications | Ν | Y | |
| 4 | Stop the instance | Ν | 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 | Ν | |
| 8 | Resume transaction log mirroring | Y | Ν | |
| 9 | Start the instance | Ν | Y | |
| 10 | Run backup | Y | Y | |
| 11 | Resume applications | Ν | 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
```

2. Stop output of archive logs

Transaction logs may accumulate during replacement of backup storage disk, and if the data storage disk or the transaction log storage disk becomes full, there is a risk that operations may not be able to continue.

To prevent this, use the following methods to stop output of archive logs.

- Changing archive_command

Specify a command that will surely complete normally, such as "echo skipped archiving WAL file %f" or "/bin/true", so that archive logs will be regarded as having been output.

If you specify echo, a message is output to the server log, so it may be used as a reference when you conduct investigations.

- Reload the configuration file

Execute the pg_ctl reload command or the pg_reload_conf SQL function to reload the configuration file.

If you simply want to stop output of errors without the risk that operations will not be able to continue, specify an empty string (") in archive_command and reload the configuration file.

3. Recover the failed disk

Replace the disk, and then recover the volume configuration information.

4. Create a backup data storage destination

Create a backup data storage destination.

Example

```
$ mkdir /database/inst1
$ chown fsepuser:fsepuser /database/inst1
$ chmod 700 /database/inst1
```

Refer to "3.2.2 Using Server Commands" for information on how to create a backup data storage destination.

5. Resume output of archive logs

Return the archive_command setting to its original value, and reload the configuration file.

6. Resume transaction log mirroring

Execute the pgx_resume_wal_multiplexing SQL function.

Example

```
SELECT pgx_resume_wal_multiplexing()
```

7. Run backup

Use the pgx_dmpall command to back up the database cluster.

Specify the following value in the pgx_dmpall command:

- Specify the data storage destination in the -D option. If the -D option is omitted, the value of the PGDATA environment variable is used by default.

Example

> pgx_dmpall -D /database/inst1

If an instance has been stopped

1. Stop applications

Stop applications that are using the database.

2. Stop the instance

Stop the instance. Refer to "2.1.2 Using Server Commands" for details.

If the instance fails to stop, refer to "14.11 Actions in Response to Failure to Stop an Instance".

3. Recover the failed disk

Replace the disk, and then recover the volume configuration information.

4. Create a backup data storage destination

Create a backup data storage destination.

Example

```
# mkdir /backup/inst1
# chown fsepuser:fsepuser /backup/inst1
# chmod 700 /backup/inst1
```

Refer to "3.2.2 Using Server Commands" for details.

5. Start the instance

Start the instance. Refer to "2.1.2 Using Server Commands" for information on how to start an instance.

6. Run backup

Use the pgx_dmpall command to back up the database cluster.

Specify the following value in the pgx_dmpall command:

- Specify the data storage destination in the -D option. If the -D option is omitted, the value of the PGDATA environment variable is used by default.

Example

> pgx_dmpall -D /database/inst1

7. Resume applications

Resume applications that are using the database.



- Refer to "pgx_rcvall" and "pgx_dmpall" in the Reference for information on the pgx_rcvall command and pgx_dmpall command.
- Refer to "Write Ahead Log" under "Server Administration" in the PostgreSQL Documentation for information on archive_command.
- Refer to "B.1 WAL Mirroring Control Functions" for information on pgx_resume_wal_multiplexing.

14.2 Recovering from Data Corruption

If data in a disk is logically corrupted and the database does not operate properly, you can recover the database cluster to its state at the time of backup.

There are two methods of recovery:

- 14.2.1 Using WebAdmin
- 14.2.2 Using the pgx_rcvall Command



- Back up the database cluster after recovering it. Backup deletes obsolete archive logs (transaction logs copied to the backup data storage destination), freeing up disk space and reducing the recovery time.
- If you recover data to a point in the past, a new time series (database update history) will start from that recovery point. When recovery is complete, the recovery point is the latest point in the new time series. When you subsequently recover data to the latest state, the database update is re-executed on the new time series.

14.2.1 Using WebAdmin

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

Refer to "14.1.1 Using WebAdmin" for details.

14.2.2 Using the pgx_rcvall Command

Recover the database cluster by specifying in the pgx_rcvall command the date and time of the backup you want to read from. Then reexecute the transaction as required to recover the data.

Follow the procedure below to recover the data storage disk.

1. Stop applications

Stop applications that are using the database.

2. Stop the instance

Stop the instance. Refer to "2.1.2 Using Server Commands" for information on how to stop an instance.

If the instance fails to stop, refer to "14.11 Actions in Response to Failure to Stop an Instance".

3. Confirm the backup date and time

Execute the pgx_rcvall command to confirm the backup data saved in the backup data storage destination, and determine a date and time prior to data corruption.

Specify the following values in the pgx_rcvall command:

- Specify the data storage destination in the -D option. If the -D option is omitted, the value of the PGDATA environment variable is used by default.
- Specify the backup storage directory in the -B option.
- The -l option displays the backup data information.

Example

4. Recover the keystore, and enable automatic opening of the keystore

When the data in the database has been encrypted, restore the keystore to its state at the time of the database backup. Configure automatic opening of the keystore as necessary.

5. Recover the database cluster

Use the pgx_rcvall command to recover the database cluster.

Specify the following values in the pgx_rcvall command:

- Specify the data storage destination in the -D option. If the -D option is omitted, the value of the PGDATA environment variable is used by default.
- Specify the backup storage directory in the -B option.

- Specify the recovery date and time in the -e option.

Example

In the following examples, "May 20, 2017 10:00:00" is specified as the recovery time.

> pgx_rcvall -D /database/inst1 -B /backup/inst1 -e '2017-05-20 10:00:00'



If recovery fails, remove the cause of the error in accordance with the displayed error message and then re-execute the pgx_rcvall command.

If the message "pgx_rcvall: an error occurred during recovery" is displayed, then the log recorded when recovery was executed is output after this message. The cause of the error is output in around the last fifteen lines of the log, so remove the cause of the error in accordance with the message and then re-execute the pgx_rcvall command.

The following message displayed during recovery is output as part of normal operation of pgx_rcvall command (therefore the user does not need not be concerned).

FATAL: The database system is starting

6. Start the instance

Start the instance. Refer to "2.1.2 Using Server Commands" for information on how to start an instance.

If necessary, re-execute transaction processing from the specified recovery time, and then resume database operations.



The pgx_rcvall command cannot accurately recover a hash index. If you are using a hash index, wait for the instance to start and then

execute the REINDEX command for the appropriate index.

7. Resume applications

Resume applications that are using the database.



Refer to "pgx_rcvall" in the Reference for information on the pgx_rcvall command.

14.3 Recovering from an Incorrect User Operation

This section describes how to recover database clusters when data has been corrupted due to erroneous user operations.

There are two methods of recovery:

- 14.3.1 Using WebAdmin
- 14.3.2 Using the pgx_rcvall Command



- Back up the database cluster after recovering it. Backup deletes obsolete archive logs (transaction logs copied to the backup data storage destination), freeing up disk space and reducing the recovery time.
- If you recover data to a point in the past, a new time series (database update history) will start from that recovery point. When recovery is complete, the recovery point is the latest point in the new time series. When you subsequently recover data to the latest state, the database update is re-executed on the new time series.

- An effective restore point is one created on a time series for which you have made a backup. That is, if you recover data to a point in the past, you cannot use any restore points set after that recovery point. Therefore, once you manage to recover your target past data, make a backup.

14.3.1 Using WebAdmin

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

🌀 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.



WebAdmin cannot accurately recover a hash index. If you are using a hash index, then after recovery, execute the REINDEX command for the appropriate index.

.

6. Resume database operations

If necessary, re-execute transaction processing from the backup point to when an erroneous operation was performed, and then resume database operations.

14.3.2 Using the pgx_rcvall Command

The pgx_rcvall command recovers database clusters to the restore point created with the server command. Refer to "Setting a restore point" in "3.2.2 Using Server Commands" for information on how to create a restore point.

Follow the procedure below to recover the data in the data storage disk.

1. Stop applications

Stop applications that are using the database.

2. Stop the instance

Stop the instance. Refer to "2.1.2 Using Server Commands" for information on how to stop an instance.

If the instance fails to stop, refer to "14.11 Actions in Response to Failure to Stop an Instance".

3. Confirm the restore point

Execute the pgx_rcvall command to confirm the backup data saved in the backup data storage destination, and use a restore point recorded in an arbitrary file, as explained in "3.2.2 Using Server Commands", to determine a restore point prior to the erroneous operation.

Specify the following values in the pgx_rcvall command:

- Specify the data storage destination in the -D option. If the -D option is omitted, the value of the PGDATA environment variable is used by default.
- Specify the backup data storage destination in the -B option.
- The -l option displays the backup data information.

Example

```
> pgx_rcvall -D /database/inst1 -B /backup/inst1 -1
Date Status Dir
2017-05-01 10:00:00 COMPLETE /backup/inst1/2017-05-01_10-00-00
```

4. Recover the keystore, and enable automatic opening of the keystore

When the data in the database has been encrypted, restore the keystore to its state at the time of the database backup. Configure automatic opening of the keystore as necessary.

5. Recover the database cluster

Use the pgx_rcvall command to recover the database cluster.

Specify the following values in the pgx_rcvall command:

- Specify the data storage destination in the -D option. If the -D option is omitted, the value of the PGDATA environment variable is used by default.
- Specify the backup data storage destination in the -B option.
- The -n option recovers the data to the specified restore point.

Example

The following example executes the pgx_rcvall command with the restore point "batch_20170503_1".

|--|



If recovery fails, remove the cause of the error in accordance with the displayed error message and then re-execute the pgx_rcvall command.

If the message "pgx_rcvall: an error occurred during recovery" is displayed, then the log recorded when recovery was executed is output after this message. The cause of the error is output in around the last fifteen lines of the log, so remove the cause of the error in accordance with the message and then re-execute the pgx_rcvall command.

The following message displayed during recovery is output as part of normal operation of pgx_rcvall (therefore the user does not need not be concerned).

FATAL: The database system is starting

6. Start the instance

Start the instance.

Refer to "2.1.2 Using Server Commands" for information on how to start an instance.

G Note

The pgx_rcvall command cannot accurately recover a hash index. If you are using a hash index, wait for the instance to start and then execute the REINDEX command for the appropriate index.

7. Restart operation of the database

If necessary, re-execute transaction processing from the specified recovery time to the point when an erroneous operation was performed, and then resume database operations.



Refer to "pgx_rcvall" in the Reference for information on the pgx_rcvall command.

14.4 Actions in Response to an Application Error

If there is a connection from a client that has been in the waiting state for an extended period, you can minimize performance degradation of the database by closing the problematic connection.

The following methods are available for identifying a connection to be closed:

- view(pg_stat_activity) (refer to "14.4.1 When using the view (pg_stat_activity)")
- ps command (refer to "14.4.2 Using the ps Command")

Use the system management function (pg_terminate_backend) to disconnect connections.

14.4.1 When using the view (pg_stat_activity)

When using the view (pg_stat_activity), follow the procedure below to close a connection.

1. Use psql command to connect to the postgres database.

```
> psql postgres
psql (12.1)
Type "help" for help.
```

2. Close connections from clients that have been in the waiting state for an extended period.

Use pg_terminate_backend() to close connections that have been trying to connect for an extended period.

However, when considering continued compatibility of applications, do not reference or use system catalogs and functions directly in SQL statements. Refer to "Notes on Application Compatibility" in the Application Development Guide for details.

Example

The following example closes connections where the client has been in the waiting state for at least 60 minutes.

```
select pid,usename,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
```



- Refer to "System Administration Functions" under "The SQL Language" in the PostgreSQL Documentation for information on pg_terminate_backend.

.....

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

.

14.4.2 Using the ps Command

Follow the procedure below to close a connection using a standard Unix tool (ps command).

1. Execute the ps command.

Note that "<*x*>" indicates the product version.

| > ps axw | fo use | r,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</x> |
| fsepuser | 20518 | 20517 | ? | _ postgres: logger process |
| fsepuser | 20520 | 20517 | ? | _ postgres: checkpointer process |
| fsepuser | 20521 | 20517 | ? | _ postgres: writer process |
| fsepuser | 20522 | 20517 | ? | _ postgres: wal writer process |
| fsepuser | 20523 | 20517 | ? | _ postgres: autovacuum launcher process |
| fsepuser | 20524 | 20517 | ? | _ postgres: archiver process |
| fsepuser | 20525 | 20517 | ? | _ postgres: stats collector process |
| 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.

14.5 Actions in Response to an Access Error

If access is denied, grant privileges allowing the instance administrator to operate the following directories, and then re-execute the operation. Also, refer to the event log and the server log, and confirm that the file system has not been mounted as read-only due to a disk error. If the file system has been mounted as read-only, mount it properly and then re-execute the operation.

- Data storage destination
- Tablespace storage destination
- Transaction log storage destination
- Backup data storage destination

🚯 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.

14.6 Actions in Response to Insufficient Space on the Data Storage Destination

If the data storage destination runs out of space, check if the disk contains any unnecessary files and delete them so that operations can continue.

If deleting unnecessary files does not solve the problem, you must migrate data to a disk with larger capacity.

There are two methods of migrating data:

- 14.6.1 Using a Tablespace
- 14.6.2 Replacing the Disk with a Larger Capacity Disk

14.6.1 Using a Tablespace

FUJITSU Enterprise Postgres enables you to use a tablespace to change the storage destination of database objects, such as tables and indexes, to a different disk.

The procedure is as follows:

1. Create a tablespace

Use the CREATE TABLESPACE command to create a new tablespace in the prepared disk.

2. Modify the tablespace

Use the ALTER TABLE command to modify tables for the newly defined tablespace.



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

14.6.2 Replacing the Disk with a Larger Capacity Disk

Before replacing the disk with a larger capacity disk, migrate resources at the data storage destination using the backup and recovery features.

There are two methods of performing backup and recovery:

- 14.6.2.1 Using WebAdmin

- 14.6.2.2 Using Server Commands

The following sections describe procedures that use each of these methods to replace the disk and migrate resources at the data storage destination.



- Before replacing the disk, stop applications and instances that are using the database.
- It is recommended that you back up the database cluster following recovery. Backup deletes obsolete archive logs (transaction logs copied to the backup data storage destination), freeing up disk space and reducing the recovery time.

14.6.2.1 Using WebAdmin

Follow the procedure below to replace the disk and migrate resources at the data storage destination by using WebAdmin.

1. Back up files

If the disk at the data storage destination contains any required files, back up the files. It is not necessary to back up the data storage destination.

2. Stop applications

Stop applications that are using the database.

3. Back up the database cluster

Back up the latest resources at the data storage destination. Refer to "3.2.1 Using WebAdmin" for details.

4. Stop the instance

Stop the instance. Refer to "2.1.1 Using WebAdmin" for information on how to stop an instance.

5. Replace with a larger capacity disk

Replace the disk. Then, recover the volume configuration information.

6. Recover the database cluster

Log in to WebAdmin, and perform recovery operations. Refer to steps 4 ("Create a tablespace directory ") to 7 ("Run recovery") under "If failure occurred in the data storage disk or the transaction log storage disk" in "14.1.1 Using WebAdmin" for information on the procedure. An instance is automatically started when recovery is successful.

7. Resume applications

Resume applications that are using the database.

8. Restore the files

Restore the files backed up in step 1.

14.6.2.2 Using Server Commands

Follow the procedure below to replace the disk and migrate resources at the data storage destination by using server commands.

1. Back up files

If the disk at the data storage destination contains any required files, back up the files. It is not necessary to back up the data storage destination.

2. Stop applications

Stop applications that are using the database.

3. Back up the database cluster

Back up the latest resources at the data storage destination. Refer to "3.2.2 Using Server Commands" for details.

4. Stop the instance

After backup is complete, stop the instance. Refer to "2.1.2 Using Server Commands" for information on how to stop an instance. If the instance fails to stop, refer to "14.11 Actions in Response to Failure to Stop an Instance".

5. Replace with a larger capacity disk

Replace the disk. Then, recover the volume configuration information.

6. Create a data storage destination

Create a data storage destination. If a tablespace was defined, also create a directory for it.

Example

```
$ mkdir /database/inst1
$ chown fsepuser:fsepuser /database/inst1
$ chmod 700 /database/inst1
```

7. Recover the keystore, and enable automatic opening of the keystore

When the data in the database has been encrypted, restore the keystore to its state at the time of the database backup. Configure automatic opening of the keystore as necessary.

8. Recover the database cluster

Use the pgx_rcvall command to recover the database cluster.

- Specify the data storage destination in the -D option. If the -D option is omitted, the value of the PGDATA environment variable is used by default.
- Specify the backup storage directory in the -B option.

Example

> pgx_rcvall -D /database/inst1 -B /backup/inst1



If recovery fails, remove the cause of the error in accordance with the displayed error message and then re-execute the pgx_rcvall command.

.....

If the message "pgx_rcvall: an error occurred during recovery" is displayed, then the log recorded when recovery was executed is output after this message. The cause of the error is output in around the last fifteen lines of the log, so remove the cause of the error in accordance with the message and then re-execute the pgx_rcvall command.

The following message displayed during recovery is output as part of normal operation of pgx_rcvall (therefore the user does not need not be concerned).

FATAL: The database system is starting

祸 See

Refer to "pgx_rcvall" in the Reference for information on the pgx_rcvall command.

9. Start the instance

Start the instance.

Refer to "2.1.2 Using Server Commands" for information on how to start an instance.



The pgx_rcvall command cannot accurately recover a hash index. If you are using a hash index, wait for the pgx_rcvall command to end and then execute the REINDEX command for the appropriate index.

10. Resume applications

Resume applications that are using the database.

11. Restore files

Restore the files backed up in step 1.

14.7 Actions in Response to Insufficient Space on the Backup Data Storage Destination

If space runs out on the backup data storage destination, check if the disk contains any unnecessary files and delete them, and then make a backup as required.

If deleting unnecessary files does not solve the problem, take the following action:

- 14.7.1 Temporarily Saving Backup Data
- 14.7.2 Replacing the Disk with a Larger Capacity Disk

14.7.1 Temporarily Saving Backup Data

This method involves temporarily moving backup data to a different directory, saving it there, and securing disk space on the backup data storage destination so that a backup can be made normally.

Use this method if you need time to prepare a larger capacity disk.

If space runs out on the backup data storage destination, archive logs can no longer be stored in the backup data storage destination. As a result, transaction logs continue to accumulate in the data storage destination or the transaction log storage destination.

If action is not taken soon, the transaction log storage destination will become full, and operations may not be able to continue.

To prevent this, secure space in the backup data storage destination, so that archive logs can be stored.

There are two methods of taking action:

- 14.7.1.1 Using WebAdmin
- 14.7.1.2 Using Server Commands

14.7.1.1 Using WebAdmin

Follow the procedure below to recover the backup data storage disk.

1. Temporarily save backup data

Move backup data to a different directory and temporarily save it, and secure space in the backup data storage destination directory.

The reason for saving the backup data is so that the data in the data storage destination can be recovered even if it is corrupted before you perform recovery. If there is no disk at the save destination and you consider that there is no risk of corruption at the data storage destination, delete the backup data.

The following example saves backup data from the backup data storage destination directory (/backup/inst1) under /mnt/usb/backup.

Example

```
> mkdir /mnt/usb/backup/
> mv /backup/inst1/* /mnt/usb/backup/
```

2. Back up the database cluster

Back up the latest resources at the data storage destination. Refer to "3.2.1 Using WebAdmin" for details.

3. Delete temporarily saved backup data

If backup completes normally, the temporarily saved backup data becomes unnecessary and is deleted.

The following example deletes backup data that was temporarily saved in /mnt/usb.

Example

> rm -rf /mnt/usb/backup

14.7.1.2 Using Server Commands

The following describes the procedure for recovering the backup storage disk.

There are two methods of taking action:

- Performing recovery while the instance is active
- Stopping the instance before performing recovery

The following table shows the different steps to be performed depending on whether you stop the instance.

| No | Step | Instance stopped | |
|----|--------------------------------------|------------------|-----|
| | | No | Yes |
| 1 | Stop transaction log mirroring | Y | N |
| 2 | Stop output of archive logs | Y | Ν |
| 3 | Stop applications | Ν | Y |
| 4 | Stop the instance | Ν | 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 | Ν | Y |
| 9 | Run backup Y Y | | Y |
| 10 | Resume applications | N | Y |
| 11 | Delete temporarily saved backup data | Y | Y |

Y: Required

N: Not required

The procedure is as follows:

Performing recovery while the instance is active

1. Stop transaction log mirroring

Stop transaction log mirroring.

2. Stop output of archive logs

Transaction logs may accumulate during replacement of backup storage disk, and if the data storage disk or the transaction log storage disk becomes full, there is a risk that operations may not be able to continue.

To prevent this, use the following methods to stop output of archive logs.

- Changing the archive_command parameter

Specify a command that will surely complete normally, such as "echo skipped archiving WAL file %f" or "/bin/true", so that archive logs will be regarded as having been output.

If you specify echo, a message is output to the server log, so it may be used as a reference when you conduct investigations.

- Reloading the configuration file

Run the pg_ctl reload command or the pg_reload_conf SQL function.

If you simply want to stop output of errors without the risk that operations will not be able to continue, specify an empty string (") in archive_command and reload the configuration file.

3. Temporarily save backup data

Move backup data to a different directory and temporarily save it, and secure space in the backup data storage destination directory.

The reason for saving the backup data is so that the data in the data storage destination can be recovered even if it is corrupted before you perform the next step. If there is no disk at the save destination and you consider that there is no risk of corruption at the data storage destination, delete the backup data.

The following example saves backup data from the backup data storage destination directory (/backup/inst1) under /mnt/usb/backup.

Example

```
> mkdir /mnt/usb/backup/
> mv /backup/inst1/* /mnt/usb/backup/
```

4. Resume output of archive logs

Return the archive_command setting to its original value, and reload the configuration file.

5. Resume transaction log mirroring

Execute the pgx_resume_wal_multiplexing SQL function.

Example

```
SELECT pgx_resume_wal_multiplexing()
```

6. Run backup

Use the pgx_dmpall command to back up the database cluster.

Specify the following option in the pgx_dmpall command:

- Specify the directory of the data storage destination in the -D option. If the -D option is omitted, the value of the PGDATA environment variable is used by default.

Example

> pgx_dmpall -D /database/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 "14.11 Actions in Response to Failure to Stop an Instance".

3. Temporarily save backup data

Move backup data to a different directory and temporarily save it, and secure space in the backup data storage destination directory.

The reason for saving the backup data is so that the data in the data storage destination can be recovered even if it is corrupted before you perform recovery. If there is no disk at the save destination and you consider that there is no risk of corruption at the data storage destination, delete the backup data.

The following example saves backup data from the backup data storage destination directory (/backup/inst1) under/mnt/usb/backup.

Example

```
> mkdir /mnt/usb/backup/
> mv /backup/inst1/* /mnt/usb/backup/
```

4. Start the instance

Start the instance. Refer to "2.1.2 Using Server Commands" for information on how to start an instance.

5. Run backup

Use the pgx_dmpall command to back up the database cluster.

Specify the following value in the pgx_dmpall command:

- Specify the data storage destination in the -D option. If the -D option is omitted, the value of the PGDATA environment variable is used by default.

Example

> pgx_dmpall -D /database/inst1

6. Resume applications

Resume applications that are using the database.

7. Delete temporarily saved backup data

If backup completes normally, the temporarily saved backup data becomes unnecessary and is deleted.

The following example deletes backup data that was temporarily saved in /mnt/usb.

Example

> rm -rf /mnt/usb/backup



- Refer to "pgx_rcvall" and "pgx_dmpall" in the Reference for information on the pgx_rcvall command and pgx_dmpall command.

- Refer to "Write Ahead Log" under "Server Administration" in the PostgreSQL Documentation for information on archive_command.
- Refer to "B.1 WAL Mirroring Control Functions" for information on the pgx_is_wal_multiplexing_paused and pgx_resume_wal_multiplexing.

14.7.2 Replacing the Disk with a Larger Capacity Disk

This method involves replacing the disk at the backup data storage destination with a larger capacity disk, so that it does not run out of free space again. After replacing the disk, back up data to obtain a proper backup.

There are two methods of performing backup:

- 14.7.2.1 Using WebAdmin

- 14.7.2.2 Using Server Commands



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

14.7.2.1 Using WebAdmin

Follow the procedure below to recover the backup storage disk.

1. Back up files

If the disk at the backup data storage destination contains any required files, back up the files. It is not necessary to back up the backup data storage destination.

2. Temporarily save backup data

Save the backup data to a different directory.

The reason for saving the backup data is so that the data in the data storage destination can be recovered even if it is corrupted before you perform the next step. If there is no disk at the save destination and you consider that there is no risk of corruption at the data storage destination, delete the backup data.

The following example saves backup data from the backup data storage destination directory (/backup/inst1) under /mnt/usb/backup.

Example

```
> mkdir /mnt/usb/backup/
> mv /backup/inst1/* /mnt/usb/backup/
```

3. Replace with a larger capacity disk

Replace the disk. Then, recover the volume configuration information.

4. Run backup

Log in to WebAdmin, and perform recovery operations. Refer to steps 2 ("Recover the backup data") and 3 ("Run backup") under "If failure occurred on the backup storage disk" in "14.1.1 Using WebAdmin".

5. Restore files

Restore the files backed up in step 1.

6. Delete temporarily saved backup data

If backup completes normally, the temporarily saved backup data becomes unnecessary and is deleted.

The following example deletes backup data that was temporarily saved in /mnt/usb.

Example

> rm -rf /mnt/usb/backup

14.7.2.2 Using Server Commands

The procedure for recovering the backup data storage disk is described below.

There are two methods of taking action:

- Performing recovery while the instance is active
- Stopping the instance before performing recovery

The following table shows the different steps to be performed depending on whether you stop the instance.

| No | Step Instance stoppe | | 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 | Ν |
| 4 | Stop output of archive logs | Y | Ν |
| 5 | Stop applications | Ν | 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 | Ν |
| 10 | Resume transaction log mirroring | Y | Ν |
| 11 | Start the instance | Ν | 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/instl/* /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 "14.11 Actions in Response to Failure to Stop an Instance".

5. Replace with a larger capacity disk

Replace the disk. Then, recover the volume configuration information.

6. Create a backup data storage destination

Create a backup data storage destination.

Example

```
# mkdir /backup/inst1
# chown fsepuser:fsepuser /backup/inst1
# chmod 700 /backup/inst1
```

Refer to "3.2.2 Using Server Commands" for details.

7. Start the instance

Start the instance. Refer to "2.1.2 Using Server Commands" for information on how to start an instance.

8. Run backup

Use the pgx_dmpall command to back up the database cluster.

Specify the following value in the pgx_dmpall command:

- Specify the data storage destination in the -D option. If the -D option is omitted, the value of the PGDATA environment variable is used by default.

Example

> pgx_dmpall -D /database/inst1

9. Resume applications

Resume applications that are using the database.

10. Restore files

Restore the files backed up in step 1.

11. Delete temporarily saved backup data

If backup completes normally, the temporarily saved backup data becomes unnecessary and is deleted.

The following example deletes backup data that was temporarily saved in /mnt/usb.

Example

> rm -rf /mnt/usb/backup



- Refer to "pgx_rcvall" and "pgx_dmpall" in the Reference for information on the pgx_rcvall command and pgx_dmpall command.

- Refer to "Write Ahead Log" under "Server Administration" in the PostgreSQL Documentation for information on archive_command.
- Refer to "B.1 WAL Mirroring Control Functions" for information on the pgx_is_wal_multiplexing_paused and pgx_resume_wal_multiplexing.

14.8 Actions in Response to Insufficient Space on the Transaction Log Storage Destination

If the transaction log storage destination runs out of space, check if the disk contains any unnecessary files and delete them so that operations can continue.

If deleting unnecessary files does not solve the problem, you must migrate data to a disk with larger capacity.

14.8.1 Replacing the Disk with a Larger Capacity Disk

Before replacing the disk with a larger capacity disk, migrate resources at the transaction log storage destination using the backup and recovery features.

There are two methods of performing backup and recovery:

- 14.8.1.1 Using WebAdmin
- 14.8.1.2 Using Server Commands

The following sections describe procedures that use each of these methods to replace the disk and migrate resources at the transaction log storage destination.



- Before replacing the disk, stop applications that are using the database.
- It is recommended that you back up the database cluster following recovery. Backup deletes obsolete archive logs (transaction logs copied to the backup data storage destination), freeing up disk space and reducing the recovery time.

14.8.1.1 Using WebAdmin

Follow the procedure below to replace the disk and migrate resources at the transaction log storage destination by using WebAdmin.

1. Back up files

If the disk at the transaction log storage destination contains any required files, back up the files. It is not necessary to back up the transaction log storage destination.

2. Back up the database cluster

Back up the latest data storage destination resources and transaction log storage destination resources (refer to "3.2.1 Using WebAdmin" for details).

3. Stop applications

Stop applications that are using the database.

4. Stop the instance

Stop the instance. Refer to "2.1.1 Using WebAdmin" for information on how to stop an instance. WebAdmin automatically stops instances if recovery of the database cluster is performed without stopping the instance.

5. Replace with a larger capacity disk

Replace the disk. Then, recover the volume configuration information.

6. Create a tablespace directory

If a tablespace was defined after backing up, create a directory for it.

7. Recover the keystore, and enable automatic opening of the keystore

Do the following if the data in the database has been encrypted:

- Restore the keystore to its state at the time of the database backup.
- Enable automatic opening of the keystore.
- 8. Recover the database cluster

Log in to WebAdmin, and perform recovery operations. Refer to steps 4 ("Create a tablespace directory ") to 7 ("Run Recovery") under " If failure occurred in the data storage disk or the transaction log storage disk " in "14.1.1 Using WebAdmin" for information on the procedure. An instance is automatically started when recovery is successful.

9. Resume applications

Resume applications that are using the database.

10. Restore files

Restore the files backed up in step 1.

14.8.1.2 Using Server Commands

Follow the procedure below to replace the disk and migrate resources at the transaction log storage destination by using server commands.

1. Back up files

If the disk at the transaction log storage destination contains any required files, back up the files. It is not necessary to back up the transaction log storage destination.

2. Back up the database cluster

Use server commands to back up the latest data storage destination resources and transaction log storage destination resources. Refer to "3.2.2 Using Server Commands" for information on how to perform backup.

3. Stop applications

Stop applications that are using the database.

4. Stop the instance

After backup is complete, stop the instance. Refer to "2.1.2 Using Server Commands" for information on how to stop an instance.

If the instance fails to stop, refer to "14.11 Actions in Response to Failure to Stop an Instance".

5. Replace with a larger capacity disk

Replace the disk. Then, recover the volume configuration information.

6. Create a transaction log storage destination

Create a transaction log storage destination. If a tablespace was defined, also create a directory for it.

Example

```
# mkdir /tranlog/inst1
# chown fsepuser:fsepuser /tranlog/inst1
# chmod 700 /tranlog/inst1
```

7. Recover the keystore, and enable automatic opening of the keystore

When the data in the database has been encrypted, restore the keystore to its state at the time of the database backup. Configure automatic opening of the keystore as necessary.

8. Recover the database cluster

Use the pgx_rcvall command to recover the database cluster.

- Specify the data storage destination in the -D option. If the -D option is omitted, the value of the PGDATA environment variable is used by default.
- Specify the backup storage directory in the -B option.

Example

> pgx_rcvall -D /database/inst1 -B /backup/inst1



If recovery fails, remove the cause of the error in accordance with the displayed error message and then re-execute the pgx_rcvall command.

If the message "pgx_rcvall: an error occurred during recovery" is displayed, then the log recorded when recovery was executed is output after this message. The cause of the error is output in around the last fifteen lines of the log, so remove the cause of the error in accordance with the message and then re-execute the pgx_rcvall command.

The following message displayed during recovery is output as part of normal operation of pgx_rcvall command (therefore the user does not need not be concerned).

FATAL: The database system is starting

🐴 See

Refer to "pgx_rcvall" in the Reference for information on the pgx_rcvall command.

9. Start the instance

Start the instance.

Refer to "2.1.2 Using Server Commands" for information on how to start an instance.



The pgx_rcvall command cannot accurately recover a hash index. If you are using a hash index, wait for the instance to start and then execute the REINDEX command for the appropriate index.

10. Resume applications

Resume applications that are using the database.

11. Restore files

Restore the files backed up in step 1.

14.9 Errors in More Than One Storage Disk

If an error occurs in the storage destination disks or resources are corrupted, determine the cause of the error from system logs and server logs and remove the cause.

If errors occur in either of the following combinations, you cannot recover the database.

Recreate the instance, and rebuild the runtime environment.

| Data storage disk | Transaction log storage disk | Backup data storage disk |
|----------------------|---------------------------------|-----------------------------|
| Error | - | Error |
| - | Error | Error |



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

14.10 Actions in Response to Instance Startup Failure

If an instance fails to start, refer to the system log and the server log, and determine the cause of the failure.

If using WebAdmin, remove the cause of the error. Then, click [Solution] and [Recheck the status] and confirm that the instance is in the normal state.

The following sections describe common causes of errors and the actions to take.

14.10.1 Errors in the Configuration File

If you have directly edited the configuration file using a text editor or changed the settings using WebAdmin, refer to the system log and the server log, confirm that no messages relating to the files below have been output.

- postgresql.conf
- pg_hba.conf



Refer to the following for information on the parameters in the configuration file:

- "Configuring Parameters" in the Installation and Setup Guide for Server

- "Appendix A Parameters"
- "Server Configuration" and "Client Authentication" under "Server Administration" in the PostgreSQL Documentation

.....

14.10.2 Errors Caused by Power Failure or Mounting Issues

If mounting is cancelled after restarting the server, for example, because the disk device for each storage destination disk was not turned on, or because automatic mounting has not been set, then starting an instance will fail. Refer to "14.14.2 Errors Caused by Power Failure or Mounting Issues", and take actions accordingly.

14.10.3 Other Errors

This section describes the recovery procedure to be used if you cannot take any action or the instance cannot start even after you have referred to the system log and the server log.

There are two methods of recovery:

- 14.10.3.1 Using WebAdmin
- 14.10.3.2 Using Server Commands

Note that recovery will not be possible if there is an error at the backup data storage destination. If the problem cannot be resolved, contact Fujitsu technical support.

14.10.3.1 Using WebAdmin

Follow the procedure below to perform recovery.

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 "14.2.1 Using WebAdmin" for details.

14.10.3.2 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 "14.2.2 Using the pgx_rcvall Command" for details.

14.11 Actions in Response to Failure to Stop an Instance

If an instance fails to stop, refer to the system log and the server log, and determine the cause of the failure.

If the instance cannot stop despite taking action, perform the following operation to stop the instance.

There are two methods of recovery:

- 14.11.1 Using WebAdmin
- 14.11.2 Using Server Commands

14.11.1 Using WebAdmin

In the [Instances] tab, click 🔲 and select the Fast stop mode or the Immediate stop mode to stop the instance. Forcibly terminate the server process from WebAdmin if the instance cannot be stopped.

Refer to "2.1.1 Using WebAdmin" for information on the stop modes.

14.11.2 Using Server Commands

There are three methods:

- Stopping the Instance Using the Fast Mode

If backup is in progress, then terminate it, roll back all executing transactions, forcibly close client connections, and then stop the instance.

- Stopping the Instance Using the Immediate Mode

Forcibly terminate the instance immediately. A crash recovery is run when the instance is restarted.

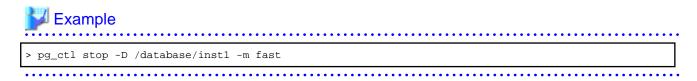
- Forcibly Stopping the Server Process

Reliably stops the server process when the other methods are unsuccessful.

14.11.2.1 Stopping the Instance Using the Fast Mode

Specify "-m fast" in the pg_ctl command to stop the instance.

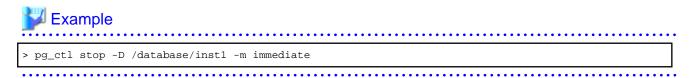
If the instance fails to stop when you use this method, stop the instance as described in "14.11.2.2 Stopping the Instance Using the Immediate Mode" or "14.11.2.3 Forcibly Stopping the Server Process".



14.11.2.2 Stopping the Instance Using the Immediate Mode

Specify "-m immediate " in the pg_ctl command to stop the instance.

If the instance fails to stop when you use this method, stop the instance as described in "14.11.2.3 Forcibly Stopping the Server Process".



14.11.2.3 Forcibly Stopping the Server Process

If both the Fast mode and the Immediate mode fail to stop the instance, use the kill command or the kill parameter of the pg_ctl command to forcibly stop the server process.

The procedure is as follows:

1. Execute the ps command

Note that "<*x*>" indicates the product version.

```
> ps axwfo user,pid,ppid,tty,command | grep postgres
fsepuser 19174 18027 pts/1
                                           \_ grep postgres
fsepuser 20517
                   1 ?
                              /opt/fsepv<x>server64/bin/postgres -D /database/inst1
fsepuser 20518 20517 ?
                               \_ postgres: logger process
fsepuser 20520 20517 ?
                               \_ postgres: checkpointer process
fsepuser 20521 20517 ?
                               \ postgres: writer process
fsepuser 20522 20517 ?
                               \_ postgres: wal writer process
fsepuser 20523 20517 ?
                               \_ postgres: autovacuum launcher process
fsepuser 20524 20517 ?
                               \ postgres: archiver process
fsepuser 20525 20517 ?
                               \_ postgres: stats collector process
```

The process ID (20517) indicates the server process.

2. Forcibly stop the server process

As instance manager, forcibly stop the server process.

Using the pg_ctl command

> pg_ctl kill SIGQUIT 20517

Using the kill command

> kill -s SIGQUIT 20517

14.12 Actions in Response to Failure to Create a Streaming Replication Standby Instance

When creating a streaming replication standby instance using WebAdmin, if the instance creation fails, refer to the system log and the server log, and determine the cause of the failure.

When an error occurs in the creation of the standby instance using WebAdmin, it is unlikely that the partially created standby instance can be resumed to complete the operation.

In such a scenario, fix the cause of the error, delete the partially created standby instance, and then create a new standby instance. This recommendation is based on the following assumptions:

- As the instance is yet to be created completely, there are no applications connecting to the database.
- The standby instance is in error state and is not running.
- There are no backups for the standby instance and as a result, it cannot be recovered.

🐴 See

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

14.13 Actions in Response to Error in a Distributed Transaction

If a system failure (such as server failure) occurs in an application that uses distributed transactions, 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

2103 | **374cc221-f6dc-4b73-9d62-d4fec9b430cd** | 2017-05-06 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
```

14.14 I/O Errors Other than Disk Failure

Even if a disk is not defective, the same input-output error messages, as those generated when the disk is defective, may be output.

A few examples of such errors are given below. The appropriate action for each error is explained respectively.

- 14.14.1 Network Error with an External Disk
- 14.14.2 Errors Caused by Power Failure or Mounting Issues

14.14.1 Network Error with an External Disk

This is an error that occurs in the network path to/from an external disk.

Determine the cause of the error by checking the information in the system log and the server log, the disk access LED, network wiring, and network card status. Take appropriate action to remove the cause of the error, for example, replace problematic devices.

14.14.2 Errors Caused by Power Failure or Mounting Issues

These are errors that occur when the disk device is not turned on, automatic mounting of the disk was not set, or mounting was accidentally cancelled.

In this case, check the information in the system log and the server log, the disk access LED, and whether the disk is mounted correctly. If problems are detected, take appropriate action.

If mounting has been cancelled, it is possible that mounting was accidentally cancelled, or automatic mounting at the time of starting the operating system is not set. In this case, set the mounting to be performed automatically.

14.15 Anomaly Detection and Resolution

The following operations performed via the command line interface will result in an anomaly in WebAdmin:

- Changes to the port and backup_destination parameters in postgresql.conf
- Changes to Mirroring Controller configuration of cluster replication added via WebAdmin

This section describes when WebAdmin checks for such anomalies, and what takes place when an anomaly is detected.

14.15.1 Port Number and Backup Storage Path Anomalies

An anomaly occurs when the value of [Port number] and/or [Backup storage path] in WebAdmin is different from the value of its corresponding parameter in postgresql.conf - port and backup_destination, respectively.

WebAdmin checks for anomalies when an instance is selected for viewing or any instance operation is performed. Anomalies will be identified for the selected instance only.

The following occurs when an anomaly is detected in port number and/or backup storage path:

- All instance operation buttons are disabled, except for "Edit instance", "Refresh instance", and "Delete Mirroring Controller"
- A red error status indicator is displayed on the instance icon
- For an anomaly specific to backup storage path, a red error status indicator is displayed on the [Backup storage] disk icon, and [Backup storage status] is set to "Error"
- The message, "WebAdmin has detected an anomaly with...", is displayed in the [Message] section along with an associated [Solution] button

Click [Solution]. The [Anomaly Error] dialog box is displayed.

Anomaly Error

| WebAdmin has detected an anomaly due to change of Port number in postgresql.conf. The instance operations are disabled and will be available only after addressing this anomaly. |
|--|
| The instance may be restarted as part of the anomaly resolution. |
| Select one of the following options to resolve this anomaly: |
| Recheck the status of the anomaly condition |
| Navigate to the "Edit instance" page |
| Reset to override Port number in postgresql.conf |
| Override Port number in WebAdmin |
| |
| |
| |
| |
| OK Cancel |
| |

Select the required option, click [OK], and then resolve the anomaly error.

Refer to "Editing instance information" in the Installation and Setup Guide for Server for information on the [Edit instance] page.

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

14.15.2 Mirroring Controller Anomalies

The following conditions will cause a Mirroring Controller anomaly:

- The Mirroring Controller management folder or configuration files have been deleted
- The permissions to the Mirroring Controller management folder or configuration files have been changed such that:
 - The instance administrator's access to Mirroring Controller configuration is denied
 - Users other than an instance administrator have access privileges to Mirroring Controller configuration files

WebAdmin checks for anomalies when Mirroring Controller status check is performed.

The following occurs when a Mirroring Controller anomaly is detected:

- All Mirroring Controller functionality is disabled for the replication cluster, except for "Delete Mirroring Controller"
- [Mirroring Controller status] is set to "Error"
- Either of the following messages is displayed in the [Message] section

"Failed to access the Mirroring Controller management folder or configuration files '*path*'. Mirroring Controller functionality has been disabled. Consider deleting Mirroring Controller and adding it again."

"Failed to find the Mirroring Controller management folder or configuration files '*path*'. Mirroring Controller functionality has been disabled. Consider deleting Mirroring Controller and adding it again."

Appendix A Parameters

This appendix describes the parameters to be set in the postgresql.conf file of FUJITSU Enterprise Postgres.

The postgresql.conf file is located in the data storage destination.

- core_directory (string)

This parameter specifies the directory where the corefile is to be output. If this parameter is omitted, the data storage destination is used by default. This parameter can only be set when specified on starting an instance. It cannot be changed dynamically, while an instance is active.

- core_contents (string)

This parameter specifies the contents to be included in the corefile.

- full: Outputs all contents of the server process memory to the corefile.
- none: Does not output a corefile.
- minimum: Outputs only non-shared memory server processes to the corefile. This reduces the size of the corefile. However, in some cases, this file may not contain sufficient information for examining the factor that caused the corefile to be output.

If this parameter is omitted, "minimum" is used by default. This parameter can only be set when specified on starting an instance. It cannot be changed dynamically, while an instance is active.

- keystore_location (string)

This parameter specifies the directory that stores the keystore file. Specify a different location from other database clusters. This parameter can only be set when specified on starting an instance. It cannot be changed dynamically, while an instance is active.

- tde_z.IBM_CCA_CSU_DEFAULT_ADAPTER (string)

This parameter specifies the value of the environment variable CSU_DEFAULT_ADAPTER for IBM Common Cryptographic Architecture (CCA) configuration. If this parameter is omitted, will follow the current CCA configuration. This parameter can only be set when specified on starting an instance. It cannot be changed dynamically, while an instance is active.

Information

Refer to IBM documentaion for the environment variable CSU_DEFAULT_ADAPTER.

- tde_z.IBM_CCA_CSU_DEFAULT_DOMAIN (string)

This parameter specifies the value of the environment variable CSU_DEFAULT_DOMAIN for IBM Common Cryptographic Architecture (CCA) configuration. If this parameter is omitted, will follow the current CCA configuration. This parameter can only be set when specified on starting an instance. It cannot be changed dynamically, while an instance is active.

Information

Refer to IBM documentaion for the environment variable CSU_DEFAULT_DOMAIN.

- tde_z.SLOT_ID (string)

Specifies the slot ID assigned to the instance of FUJITSU Enterprise Postgres. This parameter can only be set when specified on starting an instance. It cannot be changed dynamically, while an instance is active.

- tde_z.USER_PIN (string)

User PIN to open a keystore. Specifies the user PIN set to the slot ID assigned to the instance of FUJITSU Enterprise Postgres. 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.

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.

- vci.max_parallel_degree (numerical value)

This parameter specifies the maximum number of background workers used for parallel scan. The number of workers for the entire instance is limited by max_worker_processes, so add the value specified here to max_worker_processes.

A value from -8388607 to 8388607 can be specified.

- Integer (1 or greater): Parallel scan is performed using the specified degree of parallelism.
- 0: Stops the parallel scan process.
- Negative number: The specified value minus the maximum number of CPUs obtained from the environment is used as the degree of parallelism and parallel scan is performed.

If this parameter is omitted or a value outside this range is specified, 0 will be used.

- vci.shared_work_mem (numerical value)

This parameter specifies the maximum memory size used for VCI parallel scan.

- Minimum value: 32 MB
- Maximum value: Maximum value that can be expressed as a 4-byte signed integer

If this parameter is omitted or a value outside this range is specified, 1 GB will be used.

Parameters for the Global Meta Cache

- pgx_global_metacache (numerical value)

Specifies the memory size of the GMC area.

Specify a value calculated by the formula below.

A value lower than the calculated value will still work, but the meta cache may not be able to fit into the GMC area.

In this case, the system will discard the meta cache it thinks it is no longer needed, but if it is needed again, the meta cache will need to be expanded and will not perform well.

If the value is less than 10 MB and is set to a nonzero value that disables the feature, the database startup fails because the Global Meta Cache feature cannot operate.

A setting of 0 disables the Global Meta Cache feature. The default is 0.

Changing this setting requires restarting the database.

```
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.



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 W | Al mirroring | control | functions |
|-------------|--------------|---------|-----------|
| | | COLLIN | runcuons |

| Name | Return type | Description |
|----------------------------------|----------------|--|
| pgx_pause_wal_multiplexing() | void | Stops WAL multiplexing |
| pgx_resume_wal_multiplexing() | void | Resumes WAL multiplexing |
| pgx_is_wal_multiplexing_paused() | boolean | Returns true if WAL multiplexing has stopped |

If WAL multiplexing has not been configured, these functions return an error. Setting the backup_destination parameter in postgresql.conf configures WAL multiplexing.

Only superusers can execute these functions.

B.2 Transparent Data Encryption Control Functions

The following table lists the functions that can be used for transparent data encryption.

Table B.2 Transparent data encryption control functions

| Name | Return type | Description |
|--|----------------|---------------------------------|
| pgx_open_keystore(user pin passphrase) | void | Opens the keystore |
| pgx_set_master_key(user pin passphrase) | void | Sets the master encryption key |
| pgx_set_keystore_passphrase(<i>oldPassphrase</i> , <i>newPassphrase</i>) | void | Changes the keystore passphrase |

The pgx_open_keystore function uses the specified user pin or passphrase to open the keystore. When the keystore is opened, it enables access to the master encryption key. In this way, you can access the encrypted data and create encrypted tablespaces. If you are using a file-based key store, and 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 user pin is a string of 4 to 8 bytes.

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 r | masking control | functions |
|------------------|-----------------|-----------|
|------------------|-----------------|-----------|

| Name | Return type | Description |
|--------------------------------|----------------|--|
| pgx_alter_confidential_policy | boolean | Changes masking policies |
| pgx_create_confidential_policy | boolean | Creates masking policies |
| pgx_drop_confidential_policy | boolean | Deletes masking policies |
| pgx_enable_confidential_policy | boolean | Enables or disables masking policies |
| pgx_update_confidential_values | boolean | Changes replacement characters when full masking is specified for masking type |

B.3.1 pgx_alter_confidential_policy

Description

Changes masking policies

Format

The format varies depending on the content to be changed. The format is shown below.

- Common format

```
common_arg:
[schema_name := 'schemaName',]
table_name := 'tableName',
policy_name := 'policyName'
```

- 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_COLU MN' |
| | column_name | varchar(63) | Masking target name | Mandatory |
| | function_type | varchar(63) | Masking type | 'FULL' |
| | | | - 'FULL': Full masking | |
| | | | - 'PARTIAL': Partial masking | |
| | | | - 'REGEXP': Regular expression masking | |
| Partial masking | function_parameters | varchar(1024) | Masking format for partial masking | Mandatory |
| Regular expression | regexp_pattern | varchar(1024) | Search pattern for regular expression masking | Mandatory |
| masking | regexp_replacement | varchar(1024) | Replacement character/string for regular expression masking | Mandatory |
| | regexp_flags | varchar(20) | Regular expression flags | NULL |

- Delete a masking target from a masking policy

| Masking type for which an argument can be specified | Argument | Data type | Description | Default value |
|--|-------------|-------------|---------------------|---------------|
| All | action | varchar(63) | 'DROP_COLUMN' | Mandatory |
| | column_name | varchar(63) | Masking target name | Mandatory |

- Change the masking condition

| Masking type for which an argument can be specified | Argument | Data type | Description | Default value |
|--|------------|---------------|---------------------------------|---------------|
| All | action | varchar(63) | 'MODIFY_EXPRESSION' | Mandatory |
| | expression | varchar(1024) | Masking condition to be changed | Mandatory |

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

| Masking type for which an argument can be specified | Argument | Data type | Description | Default value |
|--|---------------------|---------------|--|---------------|
| All | action | varchar(63) | 'MODIFY_COLUMN' | Mandatory |
| | column_name | varchar(63) | Masking target name | Mandatory |
| | function_type | varchar(63) | Masking type | 'FULL' |
| | | | - 'FULL': Full masking | |
| | | | - 'PARTIAL': Partial masking | |
| | | | 'REGEXP': Regular expression masking | |
| Partial masking | function_parameters | varchar(1024) | Masking format for partial masking | Mandatory |
| Regular expression | regexp_pattern | varchar(1024) | Search pattern for regular expression masking | Mandatory |
| masking | regexp_replacement | varchar(1024) | Replacement character/string for regular expression masking | Mandatory |
| | regexp_flags | varchar(20) | Regular expression flags | NULL |

- Change the masking policy description

| Masking type for which an argument can be specified | Argument | Data type | Description | Default value |
|--|--------------------|---------------|----------------------------|---------------|
| All | action | varchar(63) | 'SET_POLICY_DESCRIPTION' | Mandatory |
| | policy_description | varchar(1024) | Masking policy description | Mandatory |

- Change the masking target description

| Masking type for which an argument can be specified | Argument | Data type | Description | Default value |
|--|--------------------|---------------|----------------------------|---------------|
| All | action | varchar(63) | 'SET_COLUMN_DESCRIPTION' | Mandatory |
| | column_name | varchar(63) | Masking target name | Mandatory |
| | column_description | varchar(1024) | Masking target description | Mandatory |

Details about whether arguments can be omitted are as follows.

| Argument | | Mandatory or optional | | | | | | | | |
|-------------|------|-----------------------|-------|--------|--------|------|--------|-------|------------|-------------|
| | ADI | D_COL | UMN | DROP_C | MODIFY | MOE | DIFY_C | OLUM | SET_POLICY | SET_COLUMN |
| | | | | OLUMN | _EXPRE | | Ν | | _DESCRIPTI | _DESCRIPTIO |
| | Full | Partia | Regul | | SSION | Full | Partia | Regul | ON | N |
| | mas | I. | ar | | | mas | 1 | ar | | |
| | king | maski | expre | | | king | maski | expre | | |
| | | ng | ssion | | | | ng | ssion | | |
| | | | maski | | | | | maski | | |
| | | | ng | | | | | ng | | |
| schema_name | Y | Y | Y | Y | Y | Y | Y | Y | Y | Y |
| table_name | Ν | Ν | N | Ν | Ν | N | Ν | Ν | Ν | Ν |

| Argument | | Mandatory or optional | | | | | | | | |
|---------------------|---------------------|----------------------------|--|-----------------|-------|---------------------|----------------------------|------|------------|---------------------------|
| | ADI | ADD_COLUMI | | DROP_C OLUMN | | MOE | NFY_C | OLUM | _DESCRIPTI | SET_COLUMN _DESCRIPTIO |
| | Full mas king | Partia I maski ng | Regul ar expre ssion maski ng | | SSION | Full mas king | Partia I maski ng | ar | ON | Ν |
| policy_name | N | N | N | Ν | Ν | N | N | N | Ν | Ν |
| action | Y | Y | Y | Ν | Ν | N | N | N | Ν | Ν |
| column_name | N | N | N | Ν | - | Ν | N | N | - | Ν |
| function_type | Y | N | N | - | - | Y | N | N | - | - |
| expression | - | - | - | - | N | - | - | - | - | - |
| policy_description | - | - | - | - | - | - | - | - | Ν | - |
| column_description | - | - | - | - | - | - | - | - | - | Ν |
| function_parameters | - | N | - | - | - | - | N | - | - | - |
| regexp_pattern | - | - | N | - | - | - | - | N | - | - |
| regexp_replacement | - | - | N | - | - | - | - | N | - | - |
| regexp_flags | - | - | Y | - | - | - | - | Y | - | - |

Y: Can be omitted; N: Cannot be omitted; -: Ignored when specified

Return value

| Return value | Description | | |
|--------------|------------------|--|--|
| TRUE | Ended normally | | |
| FALSE | Ended abnormally | | |

Execution example 1

Adding masking policy p1 to masking target c2

Execution example 2

Deleting masking target c1 from masking policy p1

```
postgres=# select pgx_alter_confidential_policy(table_name := 't1', policy_name := 'p1', action :=
'DROP_COLUMN', column_name := 'c1');
pgx_alter_confidential_policy
------
t
(1 row)
```

Execution example 3

Changing the masking condition for masking policy p1

Execution example 4

Changing the content of masking policy p1 set for masking target c2

Execution example 5

Changing the description of masking policy p1

Execution example 6

Changing the description of masking target c2

Description

- The arguments for the pgx_alter_confidential_policy system management function can be specified in any order.
- The action parameters below can be specified. When action parameters are omitted, ADD_COLUMN is applied.

| Parameter | Description |
|------------------------|---|
| ADD_COLUMN | Adds a masking target to a masking policy. |
| DROP_COLUMN | Deletes a masking target from a masking policy. |
| MODIFY_EXPRESSION | Changes expression. |
| MODIFY_COLUMN | Changes the content of a masking policy set for a masking target. |
| SET_POLICY_DESCRIPTION | Changes policy_description. |
| SET_COLUMN_DESCRIPTION | Changes column_description. |

- The function_parameters argument is enabled when the function_type is PARTIAL. If the function_type is other than PARTIAL, it will be ignored.

- The arguments below are enabled when the function_type is REGEXP. If the function_type is other than REGEXP, these arguments will be ignored.
 - regexp_pattern
 - regexp_replacement
 - regexp_flags

🐴 See

- Refer to "String Constants" in the PostgreSQL Documentation for information on the strings to specify for arguments.
- Refer to "POSIX Regular Expressions" in the PostgreSQL Documentation and check pattern, replacement, and flags for information on the values that can be specified for regexp_pattern, regexp_replacement, and regexp_flags.

B.3.2 pgx_create_confidential_policy

Description

Creates masking policies

Format

The format varies depending on the masking type. The format is shown below.

```
pgx_create_confidential_policy(
[schema_name := 'schemaName',]
                    := 'tableName',
table_name

      table_name
      := 'table

      policy_name
      := 'pol

      expression
      := 'expr'

      [. enable
      := 'pol

                       := 'policyName',
[, enable
                      := 'policyStatus']
[, policy_description
                              := 'policyDesc']
[, column_name := 'colName'
                             := 'FULL']
      [, function_type
      [, function_type := 'PARTIAL', partialOpt] |
      [, function_type := 'REGEXP', regexpOpt]
                                    := 'colDesc']
      [, column_description
])
partialOpt:
function_parameters := 'maskingFmt'
regexp0pt:
                   := 'regexpPattern',
regexp_pattern
regexp_replacement := 'regexpRepl.
[, regexp_flags := 'regexpFlags']
                               := 'regexpReplacementChar',
```

Argument

Details are as follows.

| Masking type for which an argument can be specified | Argument | Data type | Description | Default value |
|---|-------------|-------------|--|---------------|
| All | schema_name | varchar(63) | Schema name of table for which the masking policy is created | 'public' |
| | table_name | varchar(63) | Name of table for which the masking policy is created | Mandatory |

| Masking type for which an argument can be specified | Argument | Data type | Description | Default value |
|---|----------------------------------|---------------|---|---------------|
| | policy_name | varchar(63) | Masking policy name | Mandatory |
| | expression | varchar(1024) | Masking condition | Mandatory |
| | enable | boolean | Masking policy status | 't' |
| | | | - 't': Enabled | |
| | | | - 'f': Disabled | |
| | policy_description | varchar(1024) | Masking policy description | NULL |
| | column_name | varchar(63) | Masking target name | NULL |
| | function_type | varchar(63) | Masking type | 'FULL' |
| | | | - 'FULL': Full masking | |
| | | | - 'PARTIAL': Partial masking | |
| | | | - 'REGEXP': Regular expression masking | |
| | column_description varchar(1024) | | Masking target description | NULL |
| Partial masking | function_parameters | varchar(1024) | Masking format for partial masking | Mandatory |
| Regular expression masking | regexp_pattern | varchar(1024) | Search pattern for regular expression masking | Mandatory |
| | regexp_replacement | varchar(1024) | Replacement character/string for regular expression masking | Mandatory |
| | regexp_flags | varchar(20) | Regular expression flags | NULL |

Details about whether arguments can be omitted are as follows.

| Argument | | Mandatory or op | tional |
|---------------------|--------------|-----------------|-------------------------------|
| | Full masking | Partial masking | Regular expression masking |
| schema_name | Υ | Υ | Y |
| table_name | Ν | Ν | Ν |
| policy_name | Ν | Ν | Ν |
| expression | Ν | N | Ν |
| enable | Y | Y | Y |
| policy_description | Y | Y | Y |
| column_name | Y | Y | Y |
| function_type | Y | Υ | Y |
| column_description | Y | Υ | Y |
| function_parameters | - | N | - |
| regexp_pattern | - | - | Ν |
| regexp_replacement | - | - | Ν |
| regexp_flags | - | - | Y |

Y: Can be omitted; N: Cannot be omitted; -: Ignored when specified

Return value

| Return value | Description | | |
|--------------|------------------|--|--|
| TRUE | Ended normally | | |
| FALSE | Ended abnormally | | |

Execution example 1

Creating masking policy p1 that does not contain a masking target

Execution example 2

Creating masking policy p1 that contains masking target c1 of which the masking type is full masking

```
postgres=# select pgx_create_confidential_policy(schema_name := 'public', table_name := 't1',
policy_name := 'p1', expression := '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

Execution example 4

Creating masking policy p1 that contains masking target c3 of which the masking type is regular expression masking

Description

- The arguments for the pgx_create_confidential_policy system management function can be specified in any order.
- If column_name is omitted, only masking policies that do not contain masking target will be created.
- One masking policy can be created for each table. Use the pgx_alter_confidential_policy system management function to add a masking target to a masking policy.

- The function_parameters argument is enabled when the function_type is PARTIAL. If the function_type is other than PARTIAL, it will be ignored.
- The arguments below are enabled when the function_type is REGEXP. If the function_type is other than REGEXP, these arguments will be ignored.
 - regexp_pattern
 - regexp_replacement
 - regexp_flags



If a table for which a masking policy is to be applied is deleted, delete the masking policy as well.

```
<table-of-contents> See
```

- Refer to "String Constants" in the PostgreSQL Documentation for information on the strings to specify for arguments.
- Refer to "POSIX Regular Expressions" in the PostgreSQL Documentation and check pattern, replacement, and flags for information on the values that can be specified for regexp_pattern, regexp_replacement, and regexp_flags.

B.3.3 pgx_drop_confidential_policy

Description

Deletes masking policies

Format

```
pgx_drop_confidential_policy(
[schema_name := 'schemaName', ]
table_name := 'tableName',
policy_name := 'policyName'
)
```

Argument

Details are as follows.

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

Details about whether arguments can be omitted are as follows.

| Argument | Mandatory or optional |
|-------------|-----------------------|
| schema_name | Υ |
| table_name | Ν |
| policy_name | Ν |

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 := 'tl', policy_name := 'pl');
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.

For which a masking policy is to be applied is deleted, delete the masking policy as well.

Refer to "String Constants" in the PostgreSQL Documentation for information on the strings to specify for arguments.

B.3.4 pgx_enable_confidential_policy

Description

Enables or disables masking policies

Format

```
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 |
|-------------|-------------|---|---------------|
| schema_name | varchar(63) | Schema name of table for which a masking policy is enabled or disabled | 'public' |
| table_name | varchar(63) | Name of table for which a masking policy is enabled or disabled Mandatory | |
| policy_name | varchar(63) | Masking policy name Mandatory | |
| enable | boolean | Masking policy status | Mandatory |
| | | - 't': Enabled | |

| Argument | Data type | Description | Default value |
|----------|-----------|-----------------|---------------|
| | | - 'f': Disabled | |

Details about whether arguments can be omitted are as follows.

| Argument | Mandatory or optional |
|-------------|-----------------------|
| schema_name | Y |
| table_name | Ν |
| policy_name | Ν |
| enable | Ν |

Y: Can be omitted; N: Cannot be omitted

Return value

| Return value | Description |
|--------------|------------------|
| TRUE | Ended normally |
| FALSE | Ended abnormally |

Execution example

Enabling masking policy p1

```
postgres=# select pgx_enable_confidential_policy(table_name := 't1', policy_name := 'p1', enable :=
't');
pgx_enable_confidential_policy
-------
t
(1 row)
```

Description

The arguments for the pgx_enable_confidential_policy system management function can be specified in any order.

👪 See

Refer to "String Constants" in the PostgreSQL Documentation for information on the strings to specify for arguments.

B.3.5 pgx_update_confidential_values

Description

Changes replacement characters when full masking is specified for masking type

Format

```
pgx_update_confidential_values(
[number_value := 'numberValue']
[, char_value := 'charValue']
[, varchar_value := 'varcharValue']
[, date_value := 'dateValue']
[, ts_value := 'tsValue']
)
```

Argument

Details are as follows.

| Argument | Data type | Description | |
|---------------|------------|---|--|
| number_value | integer | Replacement character in numeric type | |
| char_value | varchar(1) | Replacement character in char type | |
| varchar_value | varchar(1) | Replacement character in varchar type | |
| date_value | date | Replacement character in date type | |
| ts_value | timestamp | Replacement character in timestamp type | |

Return value

| Return value | Description |
|--------------|------------------|
| TRUE | Ended normally |
| FALSE | Ended abnormally |

Execution example

Using '*' as a replacement character in char type and varchar type

```
postgres=# select pgx_update_confidential_values(char_value := '*', varchar_value := '*');
    pgx_update_confidential_values
    t
    (1 row)
```

Description

- The arguments for the pgx_update_confidential_values system management function can be specified in any order.
- Specify one or more arguments for the pgx_update_confidential_values system management function. A replacement character is not changed for an omitted argument.



Refer to "String Constants" in the PostgreSQL Documentation for information on the strings to specify for arguments.

B.4 VCI Data Load Control Function

The table below lists the function that loads VCI data to buffer cache.

Table B.4 VCI data load control function

| Name | Return type | Description |
|-------------------------------------|-------------|-------------------------------------|
| pgx_prewarm_vci(vci_index regclass) | int8 | Loads the VCI data to buffer cache. |

pgx_prewarm_vci loads the specified VCI data to buffer cache and returns the number of blocks of the loaded VCI data.

The aggregation process using VCI may take time immediately after an instance is started, because the VCI data has not been loaded to buffer cache. Therefore, the first aggregation process can be sped up by executing pgx_prewarm_vci after an instance is started.

The amount of memory required for preloading is the number of blocks returned by pgx_prewarm_vci multiplied by the size of one block.

This function can only be executed if the user has reference privilege to the VCI index and execution privilege to the pg_prewarm function.

B.5 High-Speed Data Load Control Functions

The table below lists the functions that can be used for high-speed data load.

| 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 |

Table B.5 High-speed data load control functions

The pgx_loader command executes the above functions internally.

Appendix C System Views

This appendix describes how to use the system views in FUJITSU Enterprise Postgres.

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

C.1 pgx_tablespaces

The pgx_tablespaces catalog provides information related to the encryption of tablespaces.

| Column | Туре | References | Description | |
|---------------|------|-------------------|---------------------------------|--|
| spctablespace | oid | pg_tablespace.oid | Tablespace OID | |
| spcencalgo | text | | Tablespace encryption algorithm | |

The spcencalgo string displays one of the following values:

- none: Tablespace is not encrypted
- AES128: AES with key length of 128 bits
- AES256: AES with key length of 256 bits

C.2 pgx_stat_lwlock

The pgx_stat_lwlock view displays statistics related to lightweight locks, with each type of content displayed on a separate line.

Table C.1 pgx_stat_lwlock view

| Column | Туре | Description | |
|-----------------|---|--|--|
| lwlock_name | name Name of the lightweight lock | | |
| total_waits | bigint | Number of waits caused by the lightweight lock | |
| total_wait_time | double precision Number of milliseconds spent in waits cause lightweight lock | | |
| stats_reset | timestamp with timezone Last time at which this statistics was reset | | |

C.3 pgx_stat_latch

The pgx_stat_latch view displays statistics related to latches, with each type of wait information within FUJITSU Enterprise Postgres displayed on a separate line.

| Column | Туре | Description | |
|-----------------|-------------------------|--|--|
| latch_name | name | Name of the latch | |
| total_waits | bigint | Number of waits caused a wait | |
| total_wait_time | double precision | Number of milliseconds spent in waits caused by the late | |
| stats_reset | timestamp with timezone | Last time at which this statistic was reset | |

C.4 pgx_stat_walwriter

The pgx_stat_walwriter view displays statistics related to WAL writing, in a single line.

Table C.3 pgx_stat_walwriter view

| Column | Туре | Description |
|------------------|-------------------------|--|
| dirty_writes | bigint | Number of times old WAL buffers were written to the disk because the WAL buffer was full when WAL records were added |
| writes | bigint | Number of WAL writes |
| write_blocks | bigint | Number of WAL write blocks |
| total_write_time | double precision | Number of milliseconds spent on WAL writing |
| stats_reset | timestamp with timezone | Last time at which this statistic was reset |

C.5 pgx_stat_sql

The pgx_stat_sql view displays statistics related to SQL statement executions, with each type of SQL statement displayed on a separate line.

Table C.4 pgx_stat_sql view

| Column | Туре | Description | |
|--------------------------|-------------------------|--|--|
| selects | bigint | Number of SELECT statements executed | |
| | | In database multiplexing mode, this number includes the SELECT statements executed in Mirroring Controller. Mirroring Controller executes the SELECT statement using the interval specified for the heartbeat_interval of the server definition file (milliseconds). | |
| inserts | bigint | Number of INSERT statements executed | |
| deletes | bigint | Number of DELETE statements executed | |
| updates | bigint | Number of UPDATE statements executed | |
| selects_with_parallelism | bigint | Number of times parallel scan was used in SELECT statements | |
| inserts_with_parallelism | bigint | Not used | |
| deletes_with_parallelism | bigint | Not used | |
| updates_with_parallelism | bigint | Not used | |
| copies_with_parallelism | bigint | Not used | |
| declares | bigint | Number of DECLARE statements executed (number of cursor OPENs) | |
| fetches | bigint | Number of FETCH statements executed | |
| checkpoints | bigint | Number of CHECKPOINT statements executed | |
| clusters | bigint | Number of CLUSTER statements executed | |
| copies | bigint | Number of COPY statements executed | |
| reindexes | bigint | Number of REINDEX statements executed | |
| truncates | bigint | Number of TRUNCATE statements executed | |
| locks | bigint | Number of times a lock occurred | |
| stats_reset | timestamp with timezone | Last time at which this statistic was reset | |

C.6 pgx_stat_gmc

The pgx_stat_gmc view provides information about the GMC areas.

Table C.5 pgx_stat_gmc view

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

Appendix D Tables Used by Data Masking

This appendix explains tables used by the data masking feature.



These tables are updated by the data masking control function, so do not use SQL statements to directly update these tables.

D.1 pgx_confidential_columns

This table provides information on masking target for which masking policies are set.

| Column | Туре | Description | |
|---------------------|---------------|---|--|
| schema_name | varchar(63) | Schema name of table for which a masking policy is applied | |
| table_name | varchar(63) | Name of table for which a masking policy is applied | |
| policy_name | varchar(63) | Masking policy name | |
| column_name | varchar(63) | Masking target name | |
| function_type | varchar(63) | Masking type | |
| | | - 'FULL': Full masking | |
| | | - 'PARTIAL': Partial masking | |
| | | - 'REGEXP': Regular expression masking | |
| function_parameters | varchar(1024) | Masking format for partial masking | |
| regexp_pattern | varchar(1024) | Search pattern for regular expression masking | |
| regexp_replacement | varchar(1024) | Replacement character/string for regular expression masking | |
| regexp_flags | varchar(20) | Regular expression flags | |
| column_description | varchar(1024) | Masking target description | |

Execution example

| | e table_n rameters | ame policy | | ms; m_name functic regexp_replace | 1 | flags |
|-------------------------------------|-------------------------|--------------|-----------|---|-----------|-------------------------|
| + + public | + t1 | + pl | + cl | ++ FULL | I | + |
| public *, 4, 11 (2 row) | t1 | p1 | c2 | PARTIAL | VVVFVV | VVFVVVV, VVV-VVVV-VVVV, |

D.2 pgx_confidential_policies

This table provides information on masking policies.

| Column | Туре | Description | |
|-------------------------|------|--|--|
| schema_name varchar(63) | | Schema name of table for which a masking policy is applied | |

| Column | Туре | Description | |
|--------------------|---------------|---|--|
| table_name | varchar(63) | Name of table for which a masking policy is applied | |
| policy_name | varchar(63) | Masking policy name | |
| expression | varchar(1024) | Masking condition | |
| enable | boolean | Masking policy status | |
| | | - 't': Enabled | |
| | | - 'f': Disabled | |
| policy_description | varchar(1024) | Masking policy description | |

Execution example

D.3 pgx_confidential_values

This table provides information on replacement characters when full masking is specified for masking type.

| Column | Data type | Description | Default value |
|-----------------|------------|----------------|-----------------------|
| number_value | integer | Numeric | 0 |
| char_value | varchar(1) | char type | Spaces |
| varchar_value | varchar(1) | varchar type | Spaces |
| date_value | date | date type | '1970-01-01' |
| timestamp_value | timestamp | timestamp type | '1970-01-01 00:00:00' |

Execution example

| number_value | char_value | <pre>k_confidential_va varchar_value +</pre> | date_value | - | |
|--------------|------------|--|------------|---------------------|--|
| | • | | | 1970-01-01 00:00:00 | |

Appendix E Tables Used by High-Speed Data Load

This appendix describes the tables used by high-speed data load.

E.1 pgx_loader_state

The pgx_loader_state table provides information about transactions prepared by high-speed data load.

| Column | Туре | Description |
|--------------|---------|--|
| id | serial | Unique identifier. |
| | | This value is assigned from the pgx_loader_state_id_seq sequence. |
| gid | text | Global transaction identifier assigned to a transaction. |
| state | text | State of the transaction. |
| | | The value can be one of the following: |
| | | - commit: The prepared transaction has been committed. |
| | | - rollback: The prepared transaction is in in-doubt state. |
| master_pid | integer | Process ID of the backend process (master process) that executed the pgx_loader control function. |
| role_oid | integer | Role identifier (OID). A prepared transaction can only be completed by the same user who executed the original transaction or by a superuser. |
| relation_oid | integer | Object identifier (OID). |



The pgx_loader_state table and pgx_loader_state_id_seq sequence are updated by high-speed data load. Do not update these database objects directly using SQL.

.

Appendix F Starting and Stopping the Web Server Feature of WebAdmin

To use WebAdmin for creating and managing a FUJITSU Enterprise Postgres instance on a server where FUJITSU Enterprise Postgres is installed, you must first start the Web server feature of WebAdmin.

- Using WebAdmin in a single-server configuration

You must start the Web server on the server on which FUJITSU Enterprise Postgres and WebAdmin are installed.

- Using WebAdmin in a multiserver configuration

You must start the Web server on all servers on which WebAdmin has been installed.

This appendix describes how to start and stop the Web server feature of WebAdmin.

Note that "<*x*>" in paths indicates the product version.

🐴 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.

G 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.

| Instances | New credential | | ¥ (|
|-----------|--------------------|-------|---------------------------|
| My Wallet | | | * indicates required fiel |
| Settings | Credential name * | Cred1 | |
| | User name * | fsep | |
| | Password * | | |
| | Confirm password * | | |
| | | | |
| | | | |
| | | | |

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 \checkmark 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.

| E FUJITSU Enterprise Po | ostgres | | English ∨ fsep ∨ FUj́ | ÎTSU |
|-------------------------|---------------------------------------|---|--|------|
| Instances | Create instance | | ✓ Ø | ŋ |
| My Wallet | | | * indicates required field(| 5) |
| Settings | Configuration type | Standalone configuration | | |
| | Server product type | FUJITSU Enterprise Postgres 12 | | |
| | Standalone configuration | | | . |
| | Location | Remote | | |
| | Instance name * | inst2 | | |
| | Instance port * | 27502 | | |
| | Host name * | StandbyHost | | |
| | Operating system credential * | Operating system credential Cred1 User name fsep | Password | |
| | Data storage path * | /database/fsep/inst2/data | | |
| | Backup | Enabled v | Backup storage path * //database/fsep/inst2/backup | |
| | Transaction log path | /database/fsep/inst2/transactionlog | | |
| | Encoding | UTF8 T | | |
| | WAL file size (MB) | 8 | | |
| | Remote WebAdmin port | | | |
| | Remote WebAdmin port for standalone * | 25515 | | |

When "Cred1" is selected in [Operating system credential], the user name and password are automatically populated from the credential.

Appendix H WebAdmin Disallow User Inputs Containing Hazardous Characters

WebAdmin considers the following as hazardous characters, which are not allowed in user inputs.

| (pipe sign)

& (ampersand sign)

; (semicolon sign)

\$ (dollar sign)

% (percent sign)

@ (at sign)

' (single apostrophe)

" (quotation mark)

' (backslash-escaped apostrophe)

" (backslash-escaped quotation mark)

<> (triangular parenthesis)

() (parenthesis)

+ (plus sign)

CR (Carriage return, ASCII 0x0d)

LF (Line feed, ASCII 0x0a)

, (comma sign)

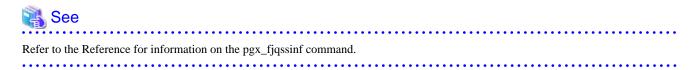
 $\ (backslash)$

Appendix I Collecting Failure Investigation Data

If the cause of an error that occurs while building the environment or during operations is unclear, data must be collected for initial investigation.

This appendix describes how to collect data for initial investigation.

Use the pgx_fjqssinf command to collect data for initial investigation.



Appendix J Operation of Transparent data Encryption in File-based Keystores

This appendix explains about the operation of transparent data encryption in file-based.

J.1 Setting the Master Encryption Key

To use transparent data encryption, you must create a keystore and set the master encryption key.

1. In the keystore_location parameter of postgresql.conf, specify the directory to store the keystore.

Specify a different location for each database cluster.

keystore_location = '/key/store/location'

Refer to "Appendix A Parameters" for information on postgresql.conf.

After editing the postgresql.conf file, either start or restart the instance.

- Using WebAdmin

Refer to "2.1.1 Using WebAdmin", and restart the instance.

- Using the pg_ctl command

Specify the following in the pg_ctl command:

- Specify "restart" as the mode.
- Specify the data storage destination directory in the -D option. If the -D option is omitted, the value of the PGDATA environment variable is used by default.
- Specify the -w option. This means that the command returns after waiting for the instance to start. If the -w option is not specified, it may not be possible to determine if the starting of the instance completed successfully or if it failed.

Example

> pg_ctl restart -w -D /database/inst1

2. Execute an SQL function, such as the one below, to set the master encryption key. This must be performed by the superuser. Execute it as the database superuser.

SELECT pgx_set_master_key('passphrase');

The value "passphrase" is the passphrase that will be used to open the keystore. The master encryption key is protected by this passphrase, so avoid specifying a short simple string that is easy to guess.

.

Refer to "B.2 Transparent Data Encryption Control Functions" for information on the pgx_set_master_key function.



Note that if you forget the passphrase, you will not be able to access the encrypted data. There is no method to retrieve a forgotten passphrase and decrypt data. Do not, under any circumstances, forget the passphrase.

The pgx_set_master_key function creates a file with the name keystore.ks in the keystore storage destination. It also creates a master encryption key from random bit strings, encrypts it with the specified passphrase, and stores it in keystore.ks. At this point, the keystore is open.

J.2 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
>
```

関 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 "J.5.3 Enabling Automatic Opening of the Keystore" for details.

J.3 Encrypting a Tablespace

Refer to "5.5 Encrypting a Tablespace".

J.4 Checking an Encrypted Tablespace

Refer to "5.6 Checking an Encrypted Tablespace".

J.5 Managing the Keystore

This section describes how to manage the keystore and the master encryption key to guard against the threat of theft.

J.5.1 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 "J.1 Setting the Master Encryption Key" for details.

After changing the master encryption key, you must immediately back up the keystore.

J.5.2 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.

J.5.3 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.



- To use WebAdmin for recovery, you must enable automatic opening of the keystore.

- Refer to "5.8 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

J.5.4 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

関 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, 2017.

> pgx_dmpall -D /database/inst1 > cp -p /key/store/location/keystore.ks /keybackup/keystore_20170501.ks

Specify the following in the pgx_dmpall command:

- Specify the data storage destination in the -D option. If the -D option is omitted, the value of the PGDATA environment variable is used by default.
- Change the master encryption key, and back up the keystore on May 5, 2017.

```
> psql -c "SELECT pgx_set_master_key('passphrase')" postgres
> cp -p /key/store/location/keystore.ks /keybackup/keystore_20170505.ks
```

Specify the following in the psql command:

- Specify the SQL function that sets the master encryption key in the -c option.
- Specify the name of the database to be connected to as the argument.

If the keystore is corrupted or lost, restore the keystore containing the latest master encryption key. If there is no keystore containing the latest master encryption key, restore the keystore to its state at the time of database backup, and recover the database from the database backup. This action recovers the keystore to its latest state.

💕 Example

- Restore the keystore containing the latest master encryption key as of May 5, 2017.

> cp -p /keybackup/keystore_20170505.ks /key/store/location/keystore.ks

- If there is no backup of the keystore containing the latest master encryption key, recover the keystore by restoring the keystore that was backed up along with the database on 1 May 2017.

> cp -p /keybackup/keystore_20170501.ks /key/store/location/keystore.ks > pgx_rcvall -B /backup/inst1 -D /database/inst1 --keystore-passphrase

Specify the following in the pgx_rcvall command:

- Specify the data storage directory in the -D option. If the -D option is omitted, the value of the PGDATA environment variable is used by default.

- Specify the backup data storage directory in the -B option.
- The --keystore-passphrase option prompts you to enter the passphrase to open the keystore.

If you have restored the keystore, repeat the process of enabling automatic opening of the keystore. This ensures that the contents of the automatically opening keystore (keystore.aks) are identical to the contents of the restored keystore.

It is recommended that you do not back up the automatically opening keystore file, keystore.aks. If the database backup medium and the backup medium storing the automatically opening keystore are both stolen, the attacker will be able to read the data even without knowing the passphrase.

If the automatically opening keystore is corrupted or lost, you must again enable automatic opening. The keystore.aks file will be recreated from keystore.ks at this time.

💦 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 "J.5.3 Enabling Automatic Opening of the Keystore" for information on how to enable automatic opening of the keystore.

J.6 Backing Up and Restoring/Recovering the Database

Refer to "5.8 Backing Up and Restoring/Recovering the Database".

You do not need to back up the Opencryptoki token directory.

J.7 Importing and Exporting the Database

Refer to "5.9 Importing and Exporting the Database".

J.8 Encrypting Existing Data

Refer to "5.10 Encrypting Existing Data".

J.9 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.

J.9.1 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.

J.9.2 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.

関 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.

.....



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.

J.10 Security-Related Notes

Refer to "5.12 Security-Related Notes".

J.11 Tips for Installing Built Applications

Refer to "5.13 Tips for Installing Built Applications".

Index

| [A] | |
|---|--------|
| Actions in Response to Instance Startup Failure | 111 |
| All user data within the specified tablespace | 22 |
| Approximate backup time | 14 |
| Approximate recovery time | 84 |
| Automatically opening the keystore | 32,154 |

[B]

| Backing Up and Recovering the Keystore 27,152 |
|---|
| Backing Up and Restoring/Recovering the Database28,153 |
| Backup/Recovery Using the Copy Command75 |
| Backup and recovery using the pgx_dmpall and pgx_rcvall |
| commands |
| backup cycle |
| Backup data |
| Backup operation |
| Backup operation (file backup) |
| Backup status |
| Backup using the backup information file |
| Backup Using the Copy Command78 |
| backup_destination (string)119 |
| Building and starting a standby server |

[C]

| Changing a Masking Policy41 |
|---|
| Changing the HSM master key27 |
| Changing the Keystore Passphrase150 |
| Changing the Master Encryption Key |
| Changing the master encryption key |
| Changing the master encryption key and the passphrase 155 |
| Changing User Pins |
| Checking an Encrypted Tablespace |
| Checking backup status |
| Checking the operating status of an instance10,12 |
| Collecting Failure Investigation Data |
| Configuration of the Copy Command75 |
| Configuration of the copy command for backup77 |
| Configuration of the copy command for recovery77 |
| Confirming a Masking Policy41 |
| Continuous archiving and point-in-time recovery |
| Copy Command for Backup80 |
| Copy Command for Recovery |
| Copy Command Interface80 |
| copy procedure for the opencryptoki token directory |
| core_contents (string)118 |
| core_directory (string)118 |
| Creating a Masking Policy |
| Cyclic usage of the backup area75 |

[D]

| Data Masking | j. |
|--|----|
| Data Types for Masking43 | ; |
| Deleting a Masking Policy | ; |
| Determining the backup area of the latest backup79 |) |

| [E] |
|---|
| Enabling and Disabling a Masking Policy |
| Enabling Automatic Opening of the Keystore |
| Encrypting a Tablespace |
| Encrypting Existing Data |
| |
| Encryption mechanisms |
| Errors in More Than One Storage Disk111 |
| |
| [F] |
| Faster hardware-based encryption/decryption22 |
| File system level backup and restore |
| |
| [H] |
| High-Speed Data Load |
| HSM Master Key Configuration |
| TISM Waster Key Configuration |
| ΓI] |
| •• |
| If failure occurred in the data storage disk or the transaction log |
| storage disk |
| If failure occurred on the backup data storage disk86,88 |
| If failure occurred on the data storage disk or the transaction log |
| storage directory |
| Importing and Exporting the Database |
| |
| Installing and Operating the In-memory Feature |
| |
| [K] |
| Keystore management |
| keystore_location (string)118 |
| |
| [L] |
| Logging in to WebAdmin |
| log in |
| |
| [M] |
| Managing the Keystore |
| |
| Masking Condition |
| Masking Format |
| Masking Policy |
| Masking Target |
| Masking Type |
| Monitoring Database Activity |
| Womtoring Database Activity |
| [O] |
| |
| Opening the Keystore |
| Operating FUJITSU Enterprise Postgres1 |
| Operation of Transparent data Encryption in File-based |
| Keystores149 |
| |
| [P] |
| Parallel Query |
| Performing backup |
| |
| |
| Perform recovery |

| pgx_stat_sql view138 |
|---|
| pgx_stat_walwriter view138 |
| pgx_tablespaces137 |
| Placement and automatic opening of the keystore file153 |
| Placing the keystore file |
| Preparing for backup |
| Preparing for HSM Collaboration |

[R]

| Recovery Using the Copy Command | 9 |
|---|---|
| reserve_buffer_ratio (numerical value)119 | 9 |

[S]

| Scope of encryption22 |
|---|
| search_path (string)119 |
| Security-Related Notes |
| Security Notes |
| Setting a restore point |
| Setting the Master Encryption Key 23,149 |
| Starting and Stopping the Web Server Feature of WebAdmin143 |
| Starting an instance |
| Startup URL for WebAdmin |
| Stopping an instance |
| Streaming replication support |
| Streaming Replication Using WebAdmin53 |
| Strong encryption algorithms |
| System Administration Functions 122 |
| System Views137 |

[T]

| tablespace_encryption_algorithm (string)118 |
|---|
| Tables Used by Data Masking140 |
| tde_z.IBM_CCA_CSU_DEFAULT_ADAPTER (string)118 |
| tde_z.IBM_CCA_CSU_DEFAULT_DOMAIN (string)118 |
| tde_z.SLOT_ID (string)118 |
| tde_z.USER_PIN (string)118 |
| Tips for Installing Built Applications |
| track_gmc (string)121 |
| track_sql (string)119 |
| track_waits (string)119 |
| Transparent Data Encryption Control Functions |
| Two-layer encryption key and the keystore |
| |

[U]

| User environment | 2 |
|------------------------|---|
| Using Server Commands1 | 1 |

[V]

| vci.control_max_workers (numerical value) | 120 |
|--|-----|
| vci.cost_threshold (numeric) | 119 |
| vci.enable (string) | 120 |
| vci.log_query (string) | 120 |
| vci.maintenance_work_mem (numerical value) | 120 |
| vci.max_local_ros (numerical value) | 120 |
| vci.max_parallel_degree (numerical value) | 120 |
| vci.shared_work_mem (numerical value) | 121 |
| | |

[W]

| WAL Mirroring Control Functions | 122 |
|---------------------------------|-----|
| WebAdmin Wallet | 145 |