

Fujitsu Enterprise Postgres
18SP1 for x86

Operation Guide

Windows/Linux



Fujitsu Enterprise Postgres
18SP1 for x86

Operation Guide

Linux



J2UL-3057-02ENZO(00)

March-2026

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 Using Transparent Data Encryption with Key Management Systems as Keystores](#)

Describes the operation of transparent data encryption when a key management system is used as a keystore.

[Chapter 7 Policy-based Login Security](#)

Describes how to apply login security policies to users.

[Chapter 8 Data Masking](#)

Describes the data masking feature.

[Chapter 9 Periodic Operations](#)

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

[Chapter 10 Building a Cluster System Using Patroni](#)

Describes how to build a cluster system using Patroni.

[Chapter 11 Autoscale of Replica Servers through Amazon EC2 Auto Scaling Integration](#)

Describes the autoscale of replica servers through Amazon EC2 Auto Scaling Integration.

[Chapter 12 Streaming Replication Using WebAdmin](#)

Describes how to create a streaming replication cluster using WebAdmin.

[Chapter 13 Installing and Operating the In-memory Feature](#)

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

[Chapter 14 Parallel Query](#)

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

[Chapter 15 High-Speed Data Load](#)

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

[Chapter 16 Global Meta Cache](#)

Describes how to use Global Meta Cache feature.

[Chapter 17 Local Meta Cache Limit](#)

Describes how to use Local Meta Cache Limit feature.

[Chapter 18 Backup/Recovery Using the Copy Command](#)

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

[Chapter 19 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 Catalogs](#)

Describes the system catalog used by Fujitsu Enterprise Postgres.

[Appendix D System Views](#)

Describes the system view used by Fujitsu Enterprise Postgres.

[Appendix E Tables Used by Transparent Data Encryption](#)

Describes the tables used by the transparent data encryption feature.

[Appendix F Tables Used by Data Masking](#)

Describes the tables used by the data masking feature.

[Appendix G Tables Used by Aggressive Freeze for Tuples](#)

Describes the tables used by the aggressive freeze for tuples.

[Appendix H Tables Used by High-Speed Data Load](#)

Describes the tables used by high-speed data load.

[Appendix I Starting and Stopping the Web Server Feature of WebAdmin](#)

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

[Appendix J WebAdmin Wallet](#)

Describes how to use the Wallet feature of WebAdmin.

[Appendix K WebAdmin Disallow User Inputs Containing Hazardous Characters](#)

Describes characters not allowed in WebAdmin.

[Appendix L Copy Command Samples that Use the Advanced Copy Feature of the ETERNUS Disk Array](#)

Describes copy command samples that use the advanced copy feature of the ETERNUS disk array.

[Appendix M Collecting Failure Investigation Data](#)

Describes how to collect information for initial investigation.

Export restrictions

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

Issue date and version

Edition 2.0: March 2026
Edition 1.0: December 2025

Copyright

Copyright 2015-2026 Fujitsu Limited

Contents

Chapter 1 Operating Fujitsu Enterprise Postgres.....	1
1.1 Operating Methods.....	1
1.2 Starting WebAdmin.....	2
1.2.1 Logging in to WebAdmin.....	2
1.3 Starting pgAdmin.....	3
1.3.1 Starting pgAdmin.....	3
1.3.2 Adding an Instance.....	3
1.3.3 Connecting/Disconnecting an Instance.....	4
1.3.4 Changing the pgAdmin Display Language.....	5
1.4 Operations Using Commands.....	5
1.5 Operating Environment of Fujitsu Enterprise Postgres.....	6
1.5.1 Operating Environment.....	6
1.5.2 File Composition.....	8
1.5.3 Precautions when Using Antivirus Software.....	8
1.6 Notes on Compatibility of Applications Used for Operations.....	9
1.7 Notes on Upgrading Database Instances.....	9
1.7.1 Additional Steps for upgrading to Fujitsu Enterprise Postgres with Vertical Clustered Index (VCI) Enabled.....	10
1.7.2 Additional Steps to upgrading to Fujitsu Enterprise Postgres with Transparent Data Encryption (TDE) Enabled.....	11
Chapter 2 Starting an Instance and Creating a Database.....	12
2.1 Starting and Stopping an Instance.....	12
2.1.1 Using WebAdmin.....	12
2.1.2 Using Server Commands.....	13
2.2 Creating a Database.....	15
2.2.1 Using pgAdmin.....	15
2.2.2 Using Client Commands.....	15
Chapter 3 Backing Up the Database.....	16
3.1 Periodic Backup.....	18
3.2 Backup Methods.....	18
3.2.1 Using WebAdmin.....	18
3.2.2 Using Server Commands.....	18
Chapter 4 Configuring Secure Communication Using Secure Sockets Layer.....	22
4.1 Configuring Communication Data Encryption.....	22
4.1.1 Issuing a Certificate.....	23
4.1.2 Deploying a Server Certificate File and a Server Private Key File.....	23
4.1.3 Distributing a CA Certificate File to the Client.....	23
4.1.4 Configuring the Operating Environment for the Database Server.....	23
4.1.5 Configuring the Operating Environment for the Client.....	23
4.1.6 Performing Database Multiplexing.....	24
Chapter 5 Protecting Storage Data Using Transparent Data Encryption.....	25
5.1 Protecting Data Using Encryption.....	25
5.2 Setting the Master Encryption Key.....	26
5.3 Opening the Keystore.....	27
5.4 Encrypting a Tablespace.....	27
5.5 Checking an Encrypted Tablespace.....	28
5.6 Managing the Keystore.....	29
5.6.1 Changing the Master Encryption Key.....	29
5.6.2 Changing the Keystore Passphrase.....	29
5.6.3 Enabling Automatic Opening of the Keystore.....	30
5.6.4 Backing Up and Recovering the Keystore.....	30
5.7 Backing Up and Restoring/Recovering the Database.....	32
5.8 Importing and Exporting the Database.....	34
5.9 Encrypting Existing Data.....	34

5.10 Operations in Cluster Systems.....	34
5.10.1 HA Clusters that do not Use Database Multiplexing.....	35
5.10.2 Database Multiplexing Mode.....	35
5.11 Security-Related Notes.....	36
5.12 Tips for Installing Built Applications.....	37
Chapter 6 Using Transparent Data Encryption with Key Management Systems as Keystores.....	38
6.1 Protecting Data Using Encryption.....	38
6.2 Setting the Master Encryption Key.....	40
6.3 Opening the Keystore.....	41
6.4 Encrypting a Tablespace.....	41
6.5 Checking an Encrypted Tablespace.....	42
6.6 Managing the Keystore.....	42
6.6.1 Changing the Master Encryption Key	42
6.6.2 Enabling Automatic Opening of the Keystore.....	42
6.6.3 Changing Credentials for Key Management Systems.....	42
6.6.4 Verifying the Master Encryption Key.....	42
6.6.5 Changes to the Key Management System.....	43
6.7 Backing Up and Restoring/Recovering the Database.....	43
6.8 Importing and Exporting the Database.....	44
6.9 Encrypting Existing Data.....	45
6.10 Operations in Cluster Systems.....	45
6.10.1 HA Clusters that do not Use Database Multiplexing.....	45
6.10.2 Database Multiplexing Mode.....	45
6.11 Security-Related Notes.....	46
6.12 Tips for Installing Built Applications.....	46
6.13 Reference Information for Linking Key Management Systems Using Sample Plugins.....	46
Chapter 7 Policy-based Login Security.....	48
7.1 Advance Preparation.....	48
7.2 Changing the Contents of the default Profile.....	48
7.3 Creating and Assigning Profiles.....	49
7.4 Actions to be Taken in the Event of Deviation from Policy.....	49
7.5 Settings in Streaming Replication Configuration.....	50
7.6 Backup and Recovery.....	50
7.7 Profile parameters.....	51
7.8 Worker Processes.....	54
Chapter 8 Data Masking.....	55
8.1 Masking Policy.....	55
8.1.1 Masking Target.....	55
8.1.2 Masking Type.....	56
8.1.3 Masking Condition.....	56
8.1.4 Masking Format.....	56
8.2 Usage Method.....	59
8.2.1 Creating a Masking Policy.....	60
8.2.2 Changing a Masking Policy.....	60
8.2.3 Confirming a Masking Policy.....	61
8.2.4 Enabling and Disabling a Masking Policy.....	61
8.2.5 Deleting a Masking Policy.....	62
8.3 Data Types for Masking.....	63
8.4 Security Notes.....	63
Chapter 9 Periodic Operations.....	64
9.1 Configuring and Monitoring the Log.....	64
9.2 Monitoring Disk Usage and Securing Free Space.....	64
9.2.1 Monitoring Disk Usage.....	64
9.2.2 Securing Free Disk Space.....	65

9.3 Automatically Closing Connections.....	65
9.4 Monitoring the Connection State of an Application.....	65
9.4.1 Using the View (pg_stat_activity).....	66
9.4.2 Using pgAdmin.....	67
9.5 Reorganizing Indexes.....	67
9.6 Monitoring Database Activity.....	68
9.6.1 Information that can be Collected.....	69
9.6.2 Collection Configuration.....	70
9.6.3 Information Reset.....	70
9.7 Scheduling of an Aggressive Freeze for Tuples (VACUUM FREEZE).....	71
9.7.1 Monitoring Trends in Transaction ID Usage.....	71
9.7.2 How to Schedule Aggressive Freeze for Tuples.....	72
9.7.3 Tuning the Allocation Time for Aggressive Freeze for Tuples.....	72
9.8 Database Monitoring with Amazon CloudWatch.....	73
9.8.1 Metrics Collection.....	73
9.8.2 Log Collection.....	74
9.8.3 Metrics Collected by the Sample Script.....	74
9.9 Monitoring Deferred SQL and Periodically Backing up statistics.....	75
9.10 Performance Tuning.....	79
9.10.1 Enhanced Query Plan Stability.....	79
9.10.1.1 Fixing the Height of a Btree Index.....	79
Chapter 10 Building a Cluster System Using Patroni.....	81
Chapter 11 Autoscale of Replica Servers through Amazon EC2 Auto Scaling Integration.....	82
11.1 Autoscale Setup.....	82
11.1.1 Building a Cluster System.....	82
11.1.1.1 Preparation.....	82
11.1.1.2 Building a Cluster System Using a Mirroring Controller.....	83
11.1.1.3 Acquiring Snapshots to Restore Data when Deploying Replica Servers.....	83
11.1.1.4 Settings for Name Resolution of the Standby Server, which will be the Replica Source Server.....	84
11.1.2 Building a Replica Server.....	84
11.1.2.1 Installation of Fujitsu Enterprise Postgres on EC2.....	84
11.1.2.2 Creation of Startup Script for Replica Server Deployment.....	85
11.1.2.3 Volume Mount Definition.....	87
11.1.2.4 Creation of AMI for Replica Server.....	87
11.1.3 Enable Autoscale.....	87
11.1.4 Preparation for Scale in.....	88
11.2 Operations for Autoscale.....	89
11.2.1 Customizing Scaling Policies.....	89
11.2.2 Scale in of Replica Servers.....	91
11.2.3 Switching/Failback of the Cascade Replication Source for the Replica Server.....	91
11.2.4 Updating an Existing AMI.....	92
11.2.5 Storage Scale up.....	92
11.2.6 Backup/Restore.....	95
Chapter 12 Streaming Replication Using WebAdmin.....	96
12.1 Creating a Standby Instance.....	96
12.2 Promoting a Standby Instance.....	97
12.3 Converting an Asynchronous Replication to Synchronous.....	97
12.4 Converting a Synchronous Replication to Asynchronous.....	98
12.5 Joining a Replication Cluster.....	98
Chapter 13 Installing and Operating the In-memory Feature.....	100
13.1 Installing Vertical Clustered Index (VCI).....	100
13.1.1 Evaluating whether to Install VCI.....	100
13.1.2 Estimating Resources.....	101
13.1.3 Setting up.....	101

13.1.3.1 Setting Parameters.....	101
13.1.3.2 Installing the Extensions.....	102
13.1.3.3 Creating a VCI.....	102
13.1.3.4 Confirming that the VCI has been Created.....	103
13.1.4 Data that can Use VCI.....	103
13.1.4.1 Relation Types.....	103
13.1.4.2 Data Types.....	104
13.2 Operating VCI.....	105
13.2.1 Commands that cannot be Used for VCI.....	105
13.2.2 Data Preload Feature.....	107
Chapter 14 Parallel Query.....	108
14.1 CPU Load Calculation.....	108
14.2 Increase of Workers during Runtime.....	108
14.3 Statistics View Displays the Action State.....	108
Chapter 15 High-Speed Data Load.....	110
15.1 Installing High-Speed Data Load.....	110
15.1.1 Deciding whether to Install.....	110
15.1.2 Estimating Resources.....	110
15.1.3 Setup.....	111
15.1.3.1 Setting Parameters.....	111
15.1.3.2 Installing the Extension.....	112
15.2 Using High-Speed Data Load.....	112
15.2.1 Loading Data.....	112
15.2.2 Checking Progress.....	113
15.2.3 Recovering from a Data Load that Ended Abnormally.....	114
15.3 Removing High-Speed Data Load.....	115
15.3.1 Removing the Extension.....	115
Chapter 16 Global Meta Cache.....	116
16.1 Usage.....	116
16.1.1 Deciding Whether to Enable the Global Meta Cache Feature.....	116
16.1.2 Estimating Memory for Global Meta Cache.....	116
16.1.3 How the GMC Memory Area Is Used.....	116
16.1.4 Enabling the Global Meta Cache Feature.....	116
16.1.5 Estimating Resources.....	117
16.2 Statistics.....	117
16.2.1 System View.....	117
Chapter 17 Local Meta Cache Limit.....	118
17.1 Usage.....	118
17.1.1 Deciding Whether to Enable the Local Meta Cache Limit Feature.....	118
17.1.2 How to Set Parameters for the Local Meta Cache Limit Feature.....	118
17.1.3 Cache Removal when Local Meta Cache Limit is Enabled.....	118
17.1.4 Performance Impact and Parameter Tuning of the Local Meta Cache Limit Feature.....	119
Chapter 18 Backup/Recovery Using the Copy Command.....	121
18.1 Configuration of the Copy Command.....	121
18.2 Backup Using the Copy Command.....	124
18.3 Recovery Using the Copy Command.....	125
18.4 Copy Command Interface.....	126
18.4.1 Copy Command for Backup.....	126
18.4.2 Copy Command for Recovery.....	128
Chapter 19 Actions when an Error Occurs.....	130
19.1 Recovering from Disk Failure (Hardware).....	131
19.1.1 Using WebAdmin.....	131
19.1.2 Using Server Command.....	132

19.2 Recovering from Data Corruption.....	136
19.2.1 Using WebAdmin.....	137
19.2.2 Using the pgx_rcvll Command.....	137
19.3 Recovering from an Incorrect User Operation.....	138
19.3.1 Using WebAdmin.....	138
19.3.2 Using the pgx_rcvll Command.....	139
19.4 Actions in Response to an Application Error.....	140
19.4.1 When using the view (pg_stat_activity).....	141
19.4.2 Using the ps Command.....	141
19.4.3 Using pgAdmin.....	142
19.5 Actions in Response to an Access Error.....	142
19.6 Actions in Response to Insufficient Space on the Data Storage Destination.....	143
19.6.1 Using a Tablespace.....	143
19.6.2 Replacing the Disk with a Larger Capacity Disk.....	143
19.6.2.1 Using WebAdmin.....	144
19.6.2.2 Using Server Commands.....	144
19.7 Actions in Response to Insufficient Space on the Backup Data Storage Destination.....	146
19.7.1 Temporarily Saving Backup Data.....	146
19.7.1.1 Using WebAdmin.....	146
19.7.1.2 Using Server Commands.....	146
19.7.2 Replacing the Disk with a Larger Capacity Disk.....	149
19.7.2.1 Using WebAdmin.....	149
19.7.2.2 Using Server Commands.....	150
19.8 Actions in Response to Insufficient Space on the Transaction Log Storage Destination.....	154
19.8.1 Replacing the Disk with a Larger Capacity Disk.....	154
19.8.1.1 Using WebAdmin.....	154
19.8.1.2 Using Server Commands.....	155
19.9 Errors in More Than One Storage Disk.....	157
19.10 Actions in Response to Instance Startup Failure.....	157
19.10.1 Errors in the Configuration File.....	157
19.10.2 Errors Caused by Power Failure or Mounting Issues.....	157
19.10.3 Other Errors.....	157
19.10.3.1 Using WebAdmin.....	158
19.10.3.2 Using Server Commands.....	158
19.11 Actions in Response to Failure to Stop an Instance.....	158
19.11.1 Using WebAdmin.....	158
19.11.2 Using Server Commands.....	158
19.11.2.1 Stopping the Instance Using the Fast Mode.....	159
19.11.2.2 Stopping the Instance Using the Immediate Mode.....	159
19.11.2.3 Forcibly Stopping the Server Process.....	159
19.12 Actions in Response to Failure to Create a Streaming Replication Standby Instance.....	160
19.13 Actions in Response to Error in a Distributed Transaction.....	160
19.14 I/O Errors Other than Disk Failure.....	161
19.14.1 Network Error with an External Disk.....	161
19.14.2 Errors Caused by Power Failure or Mounting Issues.....	162
19.15 Anomaly Detection and Resolution.....	162
19.15.1 Port Number and Backup Storage Path Anomalies.....	162
19.15.2 Mirroring Controller Anomalies.....	163
Appendix A Parameters.....	164
Appendix B System Administration Functions.....	174
B.1 WAL Mirroring Control Functions.....	174
B.2 Transparent Data Encryption Control Functions.....	174
B.2.1 pgx_open_keystore.....	174
B.2.2 pgx_set_master_key.....	175
B.2.3 pgx_declare_external_master_key.....	175
B.2.4 pgx_set_keystore_passphrase.....	176

B.3 Profile Management Functions and User Management Functions.....	176
B.3.1 Profile Management Functions.....	176
B.3.2 User Management Functions.....	177
B.4 Data Masking Control Functions.....	178
B.4.1 pgx_alter_confidential_policy.....	178
B.4.2 pgx_create_confidential_policy.....	184
B.4.3 pgx_drop_confidential_policy.....	187
B.4.4 pgx_enable_confidential_policy.....	188
B.4.5 pgx_update_confidential_values.....	189
B.5 VCI Data Load Control Function.....	190
B.6 High-Speed Data Load Control Functions.....	191
Appendix C System Catalogs.....	192
C.1 pgx_profile.....	192
C.2 pgx_user_profile.....	192
C.3 pgx_auth_password.....	193
C.4 pgx_password_history.....	193
Appendix D System Views.....	195
D.1 pgx_tablespaces.....	195
D.2 pgx_stat_lwlock.....	195
D.3 pgx_stat_latch.....	195
D.4 pgx_stat_walwriter.....	196
D.5 pgx_stat_sql.....	196
D.6 pgx_stat_gmc.....	197
D.7 pgx_stat_progress_loader.....	197
Appendix E Tables Used by Transparent Data Encryption.....	198
E.1 pgx_tde_master_key.....	198
Appendix F Tables Used by Data Masking.....	199
F.1 pgx_confidential_columns.....	199
F.2 pgx_confidential_policies.....	199
F.3 pgx_confidential_values.....	200
Appendix G Tables Used by Aggressive Freeze for Tuples.....	201
G.1 pgx_stat_freeze_results.....	201
Appendix H Tables Used by High-Speed Data Load.....	202
H.1 pgx_loader_state.....	202
Appendix I Starting and Stopping the Web Server Feature of WebAdmin.....	203
I.1 Starting the Web Server Feature of WebAdmin.....	203
I.2 Stopping the Web Server Feature of WebAdmin.....	203
Appendix J WebAdmin Wallet.....	205
J.1 Creating a Credential.....	205
J.2 Using a Credential.....	205
Appendix K WebAdmin Disallow User Inputs Containing Hazardous Characters.....	206
Appendix L Copy Command Samples that Use the Advanced Copy Feature of the ETERNUS Disk Array.....	207
Appendix M Collecting Failure Investigation Data.....	209
Index.....	210

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

Operation management using GUI tools

This involves managing operations using the WebAdmin and pgAdmin GUI tools.

- Management using WebAdmin

This removes the requirement for complex environment settings and operational design for backup and recovery that is usually required for running a database. It enables you to easily and reliably monitor the state of the database, create a streaming replication cluster, back up the database, and restore it even if you do not have expert knowledge of databases.

- Management using pgAdmin

When developing applications and maintaining the database, you can use pgAdmin to perform simple operations on database objects, such as:

- Rebuild indexes and update statistics
- Create, delete, and update database objects

In addition, from pgAdmin of Fujitsu Enterprise Postgres, you can use the expanded features provided by Fujitsu Enterprise Postgres on the PostgreSQL SQL commands.



See

Refer to pgAdmin Help for information on the expanded features of pgAdmin provided by Fujitsu Enterprise Postgres.

Operation management using commands

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

Points to consider when choosing an operation method

- You cannot combine WebAdmin and server commands to perform the following operations:
 - Use commands to operate an instance created using WebAdmin.
 - Use WebAdmin to recover a database backed up using commands.

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

- To operate an instance created using the `initdb` command in WebAdmin, the instance needs to be imported using WebAdmin.
- You can perform backup and restoration in pgAdmin, but the backup data obtained with WebAdmin and `pgx_dmpall` is not compatible with the backup data obtained with pgAdmin.

Features used in each phase

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

Operation		Operation with the GUI	Operation with commands
Setup	Creating an instance	WebAdmin is used. The server machine capacity, and the optimum parameter for operations using WebAdmin, are set automatically.	The configuration file is edited directly using the initdb command.
	Creating a standby instance	WebAdmin is used. WebAdmin performs a base backup of the source instance and creates a standby instance.	A standby instance is created using the pg_basebackup command.
	Changing the configuration files	WebAdmin is used.	The configuration file is edited directly.
Starting and stopping an instance		WebAdmin is used.	The pg_ctl command is used.
Creating a database		This is defined using pgAdmin of the GUI tool, or using the psql command or the application after specifying the DDL statement.	
Backing up the database		WebAdmin, or the pgx_dmpall command, is used.	It is recommended that the pgx_dmpall command be used. Recovery to the latest database can be performed.
Database recovery		WebAdmin is used.	To use the backup that was performed using the pgx_dmpall command, the pgx_rcvall command is used.
Monitoring	Database errors	The status in the WebAdmin window can be checked. (*1)	The messages that are output to the database server log are monitored (*1)
	Disk space	The status in the WebAdmin window can be checked. A warning will be displayed if the free space falls below 20%. (*1)	This is monitored using the df command of the operating system, for example. (*1)
	Connection status	This can be checked using pgAdmin of the GUI tool, or referencing pg_stat_activity of the standard statistics view from psql or the application.	

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

1.2 Starting WebAdmin

This section describes how to start and log in to WebAdmin.

About using WebAdmin

- It is recommended to use the following browsers with WebAdmin:

- Microsoft Edge

WebAdmin will work with other browsers, such as Firefox and Chrome, however, the look and feel may be slightly different.

- You must start the Web server feature of WebAdmin before using WebAdmin. Refer to "[Appendix I Starting and Stopping the Web Server Feature of WebAdmin](#)" for information on how to start the Web server feature of WebAdmin.

1.2.1 Logging in to WebAdmin

This section describes how to log in to WebAdmin.

Startup URL for WebAdmin

In the browser address bar, type the startup URL of the WebAdmin window in the following format:

```
http://hostNameOrIpAddress:portNumber/
```

- *hostNameOrIpAddress*: The host name or IP address of the server where WebAdmin is installed.
- *portNumber*: The port number of WebAdmin. The default port number is 27515.

Example

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

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

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

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

1.3 Starting pgAdmin

This section describes how to start pgAdmin, how to add an instance required for managing a database, and how to connect to and disconnect from the instance.

You can use pgAdmin on the Windows client.

About using pgAdmin

- You must start the instance to be connected to before using pgAdmin. Refer to "2.1 Starting and Stopping an Instance" for information on how to start an instance.
- When using pgAdmin4 with Microsoft Edge, enable network access by loopback in Microsoft Edge. Also, add Microsoft Edge to the loopback exclusion list.

1.3.1 Starting pgAdmin

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

Note that "<x>" indicates the product version. Also, the shortcuts displayed in the [Start] menu may be arranged differently depending on the operating system.

1.3.2 Adding an Instance

This section describes how to add an instance to be connected to.

If you use a link-local address with version 6 of the TCP/IP protocol, you may encounter the following error when registering the server with pgAdmin. Therefore, do not use link-local addresses.

```
unsupported format character "'" (0x22) at index 96
```

1. In the [Browser] pane, right-click [Servers], and then click [Register] >> [Server...].

2. In the [Register - Server] window, specify a value for each item.

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

- Name:** An empty text input field with a red warning icon to its left.
- Server group:** A dropdown menu showing 'Servers'.
- Background:** A checkbox that is unchecked.
- Foreground:** A checkbox that is unchecked.
- Connect now?:** A toggle switch that is turned on.
- Comments:** A large empty text area.
- Error message:** A red bar at the bottom of the dialog containing the text: 'Name' cannot be empty.
- Buttons:** 'Close', 'Reset', and 'Save' buttons are located at the bottom right.

[General] tab

- [Name]: Name of the instance to be managed

[Connection] tab

- [Host name/address]: Host name or IP address of the server where Fujitsu Enterprise Postgres is installed
- [Port]: Port number of the instance
- [Maintenance database]: Database name to which the client will connect
- [Username]: User name of the instance administrator
- [Password]: Password for the user name specified in [Username]

Point

.....
If you select [Save password], the Fujitsu Enterprise Postgres connection password is stored in the following location. Set the appropriate access permissions for the password file to protect it from unauthorized access.

- %APPDATA%\Roaming\pgAdmin\pgadmin4.db
-

1.3.3 Connecting/Disconnecting an Instance

This section describes how to connect pgAdmin to an instance, and how to disconnect it.

To connect to an instance created using WebAdmin, you must first configure the settings in the [Client authentication] window of WebAdmin to permit connection from pgAdmin.



See

.....

Refer to "Changing the settings" in the Installation and Setup Guide for Server for information on the [Client authentication] window of WebAdmin.

.....

Connecting to an instance

Starting pgAdmin does not connect it to any instance.

To connect to an instance, in the [Browser] pane, right-click the server, and then click [Connect Server].

If a password was not saved when the instance was added, enter a password in the password entry window that is displayed.

Disconnecting from an instance

To disconnect from an instance, in the [Browser] pane, right-click the server, and then click [Disconnect from Server].

1.3.4 Changing the pgAdmin Display Language

This section describes the procedure for changing the display language of pgAdmin.

1. Start pgAdmin.
2. Click [File] >> [Preferences].
3. Click [Miscellaneous] >> [User language] in the left pane of the [Preferences] dialog.
4. Select the language you want to display from the [User language] drop-down list displayed in the right pane, and click the [Save] button.
5. Refresh page by clicking the [Refresh] button.

1.4 Operations Using Commands

You can operate and manage the database using the following commands:

- Server commands

This group of commands includes commands for creating a database cluster and controlling the database. You can run these commands on the server where the database is operating.

To use these commands, you must configure the environment variables.



See

-
- Refer to "PostgreSQL Server Applications" under "Reference" in the PostgreSQL Documentation, or "Reference" for information on server commands.
 - Refer to "Configure the environment variables" in the procedure to create instances in "Using the initdb Command" in the Installation and Setup Guide for Server for information on configuring the environment variables.
-

- Client commands

This group of commands includes the psql command and commands for extracting the database cluster to a script file. These commands can be executed on the client that can connect to the database, or on the server on which the database is running.

To use these commands, you must configure the environment variables.



See

-
- Refer to "PostgreSQL Client Applications" under "Reference" in the PostgreSQL Documentation, or "Reference" for information on client commands.


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

1.5 Operating Environment of Fujitsu Enterprise Postgres

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

1.5.1 Operating Environment

The following figure shows the configuration of the Fujitsu Enterprise Postgres operating environment. The tables given below list the roles of the OS resources and Fujitsu Enterprise Postgres resources.



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

Table 1.1 OS resources

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

Table 1.2 Fujitsu Enterprise Postgres client resources

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

Table 1.3 Server resources of Fujitsu Enterprise Postgres

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

1.5.2 File Composition

Fujitsu Enterprise Postgres consists of the following files for controlling and storing the database. The table below shows the relationship between the number of such files and their location within a single instance.

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

Table 1.4 Number of files within a single instance and how to specify their location

File type	Required	Quantity	How to specify the location
Program files	Y	Multiple	Note that "<x>" indicates the product version. /opt/fsepv<x>server64
Database cluster	Y	1	Specify using WebAdmin or server commands.
Tablespace	Y	Multiple	Specify a space other than that under the database cluster, using pgAdmin or the DDL statement.
Backup	Y	Multiple	Specify using WebAdmin or server commands.
Core file	Y	Multiple	Specify using WebAdmin, server commands, or postgresql.conf.
Server certificate file (*1)	N	1	Specify using postgresql.conf.
Server private key file (*1)	N	1	Specify using postgresql.conf.
Master encryption key file (*1)	N	1	Specify the directory created as the key store using postgresql.conf.
Connection service file (*1)	N	1	Specify using environment variables.
Password file (*1)	N	1	Specify using environment variables.
CA certificate file (*1)	N	1	Specify using environment variables.

Y: Mandatory

N: Optional

*1: Set manually when using the applicable feature.



Information

If you are using antivirus software, refer to "1.5.3 Precautions when Using Antivirus Software".

1.5.3 Precautions when Using Antivirus Software

If anti-virus software is used, the instance might not start or not respond during operation, and proper operation of the database might not be guaranteed. To avoid such troubles, it is necessary to change the settings of the antivirus software so that virus scans are not executed on the active database.

Do scan exclusion for directories so that all files that comprise Fujitsu Enterprise Postgres are excluded from virus scanning. 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.

However, setting exclusions for scans poses a risk of missing malware that should be detected by antivirus software. To meet the security requirements necessary for system operation, implement risk countermeasures against malware threats across the entire system. Consider the following measures, for example:

- Introduce WAF, IPS, and IDS to prevent unauthorized access to the database server across the entire system.
- Disable unnecessary features, services, and protocols in the OS and network you use, and always obtain and apply the latest patches to prevent malware intrusion.
- Enable access log acquisition for the OS and database to detect data tampering or leakage caused by malware.
- Perform regular backups to ensure recovery from backups in case data is tampered with.
For backup methods of Fujitsu Enterprise Postgres, refer to "[Chapter 3 Backing Up the Database](#)".

1.6 Notes on Compatibility of Applications Used for Operations

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

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



See

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

1.7 Notes on Upgrading Database Instances

If you are upgrading a database instance that uses the extensions, follow the procedure below to upgrade according to the source product version.

When migrating from a database instance of FUJITSU Enterprise Postgres 10 or later

If you are using `pg_upgrade` to upgrade a database instance from FUJITSU Enterprise Postgres 10 or later that uses the extensions, follow the procedure below to upgrade.

It is strongly recommended to back up the database using `pg_dump` before performing `pg_upgrade`.

1. Create a database cluster at the destination
2. Stop the source and destination instances
3. Run the `pg_upgrade` command
4. Start the destination instance
5. Update the extensions used with `ALTER EXTENSION` to the latest version at the destination

```
ALTER EXTENSION extensionName UPDATE;
```

If you are using multiple extensions, update each extension separately.



See

For information about `pg_upgrade` and `ALTER EXTENSION`, refer to "Reference" in the PostgreSQL Documentation.

When migrating from a database instance of FUJITSU Enterprise Postgres 9.x

The following describes the tasks required to upgrade a database instance of FUJITSU Enterprise Postgres 9.x using the extension with `pg_upgrade`.

It is strongly recommended to back up the database using `pg_dump` before performing `pg_upgrade` or using `DROP EXTENSION`.

Before running `pg_upgrade`, remove the following extensions from all databases of the instance, except for "template0". If a column was created in the user table using the following extended data types, "DROP EXTENSION" will also drop that column. In such cases you might need to consider alternative upgrade methods rather than `pg_upgrade`. In that case, consider using `pg_dump` or `pg_restore`.

- pg_stat_statements
- pgx_io
- pgx_paging
- pgx_network
- pgx_network_err
- pgx_cpu
- pgx_memory
- pgx_swap
- pgx_disk
- pgx_process
- pgx_log
- oracle_compatible
- pg_dbms_stats
- pg_hint_plan

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

```
DROP EXTENSION extensionName;
```

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

```
CREATE EXTENSION extensionName;
```



See

For information about pg_upgrade, DROP EXTENSION and CREATE EXTENSION, refer to "Reference" in the PostgreSQL Documentation.

1.7.1 Additional Steps for upgrading to Fujitsu Enterprise Postgres with Vertical Clustered Index (VCI) Enabled

When upgrading FUJITSU Enterprise Postgres 11.0 or earlier instances that are using the VCI extension to FUJITSU Enterprise Postgres 12 or later using pg_upgrade, additional steps must be performed because of the incompatibility of the VCI extension between FUJITSU Enterprise Postgres 12 or later and FUJITSU Enterprise Postgres 11 or earlier.

Follow the procedure below in all databases in the FUJITSU Enterprise Postgres 11 or earlier instance, except "template0".

Before upgrading

1. Obtain the CREATE INDEX Definitions

Run the query below to list all the VCI indexes created in the database. Ensure that these indexes are re-created in the FUJITSU Enterprise Postgres 12 or later instance after pg_upgrade has finished.

```
SELECT nspname || '.' || relname AS index_relname,* FROM pg_class, pg_namespace
WHERE relnamespace = pg_namespace.oid AND relam IN (SELECT oid FROM pg_am WHERE amname='vci');
```

For each index_relname listed above, execute the commands below to obtain the CREATE INDEX definition (to use the same SQL syntax while re-creating the indexes on the FUJITSU Enterprise Postgres 12 or later instance).

```
SELECT pg_get_indexdef('indexName'::regclass);
```

2. Drop the VCI indexes and VCI extension along with all its dependencies.

To remove all the VCI indexes and VCI internal objects that are created in Fujitsu Enterprise Postgres, execute the commands below. VCI internal objects will be created in FUJITSU Enterprise Postgres 12 or later automatically when CREATE EXTENSION for VCI is executed.

```
DROP EXTENSION VCI CASCADE;
```

Information

To restore the VCI extension in the FUJITSU Enterprise Postgres 11 or earlier instance, execute CREATE EXTENSION.

After upgrading

Once the pg_upgrade operation is complete, for all databases except "template0", execute CREATE EXTENSION to create the VCI extension, and then execute CREATE INDEX for all the VCI indexes as required.

1.7.2 Additional Steps to upgrading to Fujitsu Enterprise Postgres with Transparent Data Encryption (TDE) Enabled

If you are using pg_upgrade to upgrade an instance of Fujitsu Enterprise Postgres that uses Transparent Data Encryption (TDE), there are steps you can take before upgrading.

If your old cluster operation was an HA cluster operation without database multiplexing operations, and you shared a single keystore file, see "[Before upgrading if you shared a keystore file](#)". For other operations, see "[Before upgrading](#)".

Before upgrading

Before upgrading, perform the following steps.

This is not necessary if you are using the old cluster keystore file location as the new cluster keystore file location. In that case, the old cluster cannot continue to be used. Enable automatic keystore opening for old and new clusters.

1. Copy Master Encryption Key

Copy the keystore file from the old cluster to the new cluster.

You do not need to use the pgx_set_master_key function to generate a new master encryption key on the new cluster; you must copy the keystore file from the old cluster.

As a database superuser, do the following:

```
$ mkdir <NEW-KEY-STORE>/
$ cp -p <OLD-KEY-STORE>/keystore.ks <NEW-KEY-STORE>/
```

NEW-KEY-STORE: The directory specified by the keystore_location parameter in postgresql.conf for the new cluster

OLD-KEY-STORE: The directory specified by the keystore_location parameter in postgresql.conf on the old cluster

Before upgrading if you shared a keystore file

If the primary and standby servers shared the same keystore file, copy the keystore file from the old environment and share it as the keystore file from the new environment.

For more secure management of keystore files, place them on a secure, isolated key management server or key management storage.

Enable automatic keystore opening for old and new clusters.

Chapter 2 Starting an Instance and Creating a Database

This chapter describes basic operations, from starting an instance to creating a database.

2.1 Starting and Stopping an Instance

This section describes how to start and stop an instance.

- [2.1.1 Using WebAdmin](#)
- [2.1.2 Using Server Commands](#)

Point


- To automatically start or stop an instance when the operating system on the database server is started or stopped, refer to "Configuring Automatic Start and Stop of an Instance" in the Installation and Setup Guide for Server and configure the settings.
- The collected statistics are initialized if an instance is stopped in the "Immediate" mode or if it is abnormally terminated. To prepare for such initialization of statistics, consider regular collection of the statistics by using the SELECT statement. Refer to "The Statistics Collector" in "Server Administration" in the PostgreSQL Documentation for information on the statistics.


2.1.1 Using WebAdmin

WebAdmin enables you to start or stop an instance and check its operating status.

Starting an instance


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


 is displayed when an instance is stopped.

To start a stopped instance, click .

Stopping an instance

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

 is displayed when an instance is active.



To stop an active instance, click .

Stop mode

Select the mode in which to stop the instance. The following describes the operations of the modes:






Stop mode	Connected clients	Backup being executed using the command
Smart mode (*1)	Waits for all connected clients to be disconnected.	Waits for backups being executed using the command to finish.
Fast mode	Rolls back all transactions being executed and forcibly disconnects clients.	Terminates backups being executed using the command.
Immediate mode	All server processes are terminated immediately. Crash recovery is executed the next time the instance is started.	
Kill process mode	Send SIGKILL to the process and abort all active transactions. This will lead to a crash-recovery run at the next restart.	

*1: When the processing to stop the instance in the Smart mode has started and you want to stop immediately, use the following procedure:


1. Restart the Web server feature of WebAdmin.
2. In the [Instances] tab, click .
3. In the [Instances] tab, click , and select the Immediate mode to stop the instance.

Checking the operating status of an instance

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

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

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

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.

Example

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

The message FATAL:the database system is starting up(11189)" may be output in the following cases. This is output by the pg_ctl command to see if the instance's processes have finished starting. Therefore, if there are no other applications, commands, or processes connecting to the database, ignore this message.

- If an application, command, or process connects to the database while the instance is starting up
- If the instance is started without specifying the -W option

If the -W option is specified, the command will return without waiting for the instance to finish starting. Therefore, you need to check the server log to see if the instance has started successfully.

Stopping an instance

Use the pg_ctl command to stop an instance.

Specify the following values in the pg_ctl command:

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

Example

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

Checking the operating status of an instance

Use the pg_ctl command to check the operating status of an instance.

Specify the following values in the pg_ctl command:

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

Example

When the instance is active:

```
> pg_ctl status -D /database/inst1
pg_ctl: server is running (PID: 1234)
```

When the instance is inactive:

```
> pg_ctl status -D /database/inst1
pg_ctl: no server running.
```

See

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

2.2 Creating a Database

This section explains how to create a database.

- [2.2.1 Using pgAdmin](#)
- [2.2.2 Using Client Commands](#)

2.2.1 Using pgAdmin

Follow the procedure below to define a database using pgAdmin.

1. In the pgAdmin window, right-click [Databases] in the [Browser] pane, and then click [Create] >> [Database...] to display a [Create - Database] window.
 - [General] tab
 - [Database]: Name of the database to be managed
3. Click [Save] to create the database.

2.2.2 Using Client Commands

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

An example of operations on the server is shown below.

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

```
> psql postgres
psql (<x>) (*1)
Type "help" for help.
```

*1: <x> indicates the PostgreSQL version on which this product is based.

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

```
postgres=# CREATE DATABASE db01;
CREATE DATABASE
```

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

```
postgres=# \l+
```

4. Disconnect from the postgres database.
Execute \q to terminate the psql command.

```
postgres=# \q
```

You can create a database using the createdb command.



See

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

Chapter 3 Backing Up the Database

This chapter describes how to back up the database.

Backup methods

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

- Backup using WebAdmin

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

WebAdmin is used for recovery.

- Backup using the `pgx_dmpall` command

Execute the `pgx_dmpall` command with a script to perform automatic backup. To back up data automatically, you must register the process in the automation software of the operating system. Follow the procedure given in the documentation for your operating system.

Additionally, you can also take backups periodically by using any external scheduler.

The `pgx_rcvall` command is used for recovery.

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.

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 18 Backup/Recovery Using the Copy Command](#)" for details.

Approximate backup time

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

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

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

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

When defining a tablespace

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.

When encrypting data stored in the database

There are several considerations for the backup of the keystore and backup of the database in case the data stored in the database is encrypted. Refer to the following for details:

- [5.6.4 Backing Up and Recovering the Keystore](#)
- [5.7 Backing Up and Restoring/Recovering the Database](#)

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.

The following backup methods are available for the features provided by Enterprise Postgres:

Category	Backup method	Backup target	Enterprise Postgres					
			Transparent Data Encryption	Policy-based Login Security	Data Masking	Confidentiality Management	Audit Log Feature	Database Multiplexing
Physical Backup	WebAdmin	Database cluster	Y	Y	Y	Y	Y	Y
	pgx_dmpall command		Y	Y	Y	Y	Y	Y
	Backup by continuous archiving		Y	Y	Y	Y	Y	Y
	File system level backup		Y	Y	Y	Y	Y	Y
Logical backup	pg_dumpall command	Database	N	Y	N	Y	N	Y
	pg_dump command		N	N	N	Y	N	Y
	COPY command		Table	N	N	N	N	N

Y: Can be used

N: Cannot be used

In a logical backup, the restore process recreates the indexes. For example, if the index size is large, such as for vector data handled by pgvector, the restoration time may be long. If you want to restore quickly, use a physical backup.

To export large amounts of data from multiple tables at high speed, use multiple executions to export data in parallel.

Some features are important to keep in mind when backing up. For information about the features, see the following:

Transparent Data Encryption : "[5.7 Backing Up and Restoring/Recovering the Database](#)"

Policy-based Login Security : "[7.6 Backup and Recovery](#)"

Confidentiality Management : "Backup/Restore" in the Security Operation Guide

Database Multiplexing : "Backup Operation" in the Cluster Operation Guide(Database Multiplexing)

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.

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.

When encrypting data stored in the database

If the data to be stored in the database is to be encrypted, it is necessary to enable the automatic opening of the keystore before doing so. Refer to "[5.6.3 Enabling Automatic Opening of the Keystore](#)" for details.

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.
archive_library	“(default)”	Specify the archive library. Specify [“(default)”].

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
```

Backup stores the data obtained during the backup and the backup data of the data obtained during previous backup.

If the data to be stored in the database is encrypted, refer to the following and back up the keystore:

- [5.6.4 Backing Up and Recovering the Keystore](#)

Backup status

Use the pgx_rcvall command to check the backup status.

Specify the following values in the pgx_rcvall command:

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

```
> pgx_rcvall -l -D /database/inst1
Date                Status             Dir
2022-03-01 13:30:40 COMPLETE          /backup/inst1/2022-03-01_13-30-40
```

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

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

See

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

Setting a restore point

In case you want to recover your database to a certain point in time, you can name this particular point in time, which is referred to as the restore point, by using the psql command.

By setting a restore point before executing an application, it becomes easy to identify up to which point in time the data will be reverted.

A restore point can be set to any point in time after a backup is executed. However, if a restore point is set before a backup is executed, the database cannot be recovered to that point in time. This is because restore points are recorded in the archive logs, and the archive logs are discarded when backups are executed.

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.

Example

The following example uses the psql command to connect to the database and execute the SQL statement to set a restore point.

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

```
postgres=# SELECT pg_create_restore_point('batch_20220303_1');
LOG:  restore point "batch_20220303_1" created at 0/20000E8
STATEMENT:  select pg_create_restore_point('batch_20220303_1');
pg_create_restore_point
-----
0/20000E8
(1 row)
```

.....

Refer to "[19.3.2 Using the pgx_rcvall Command](#)" for information on using a restore point to recover the database.



See

.....

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

.....

Chapter 4 Configuring Secure Communication Using Secure Sockets Layer

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

4.1 Configuring Communication Data Encryption

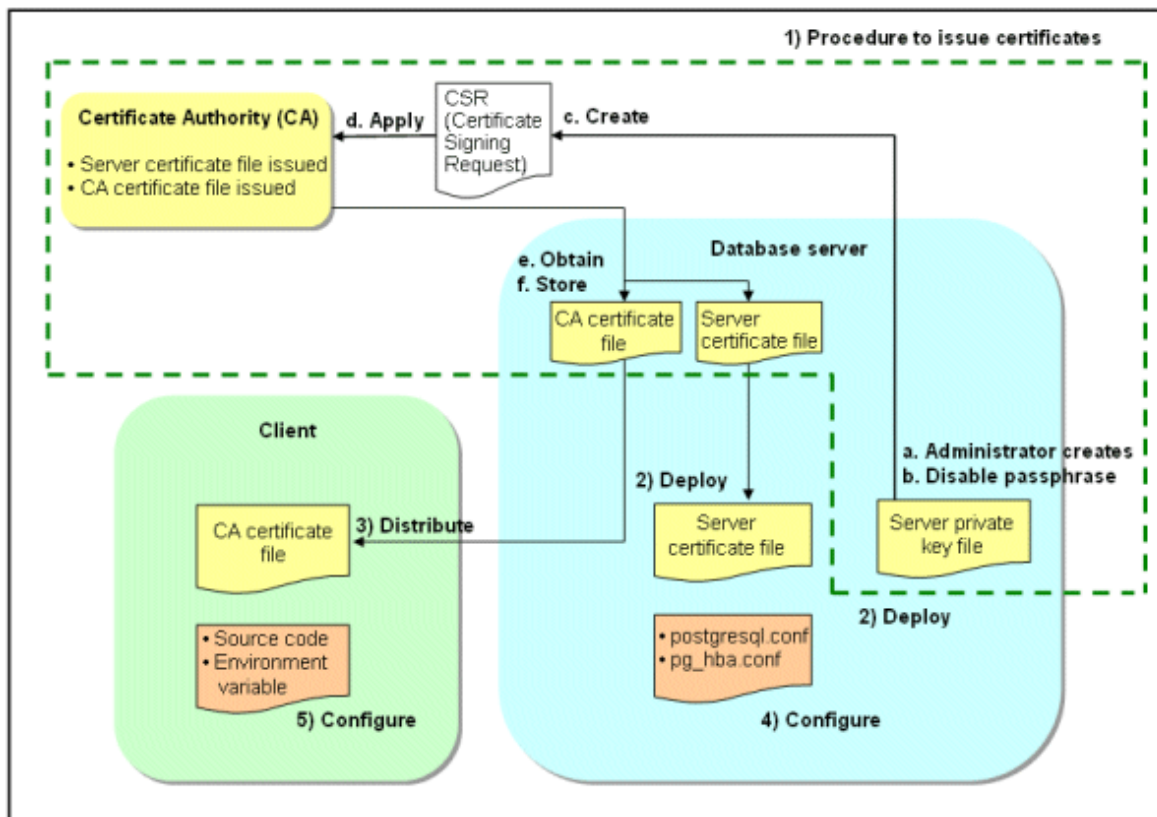
To encrypt communication data transferred between a client and a server, configure communication data encryption as described below. Communication data encryption not only protects the communication content, but it also guards against man-in-the-middle (MITM) attacks (for example, data and password theft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

Table 4.1 Configuration procedure

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

The following figure illustrates the environment for communication data encryption.

Figure 4.1 Environment for communication data encryption



4.1.1 Issuing a Certificate

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

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

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

- a. Create a server private key file
- b. Disable the passphrase for the server private key file
- c. Create a CSR (signing request for obtaining a server certificate) from the server private key file
- d. Apply to the certificate authority (CA) for a server certificate
- e. Obtain a server certificate file and a CA certificate file from the certificate authority (CA)
- f. Store the server certificate file and the CA certificate file

Note: If you lose or destroy the certificates, you will need to have them re-issued.

The above procedure enables you to prepare the following files:

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

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

Create a directory on the local disk of the database server and store the server certificate file and the server private key file in it.

Use the operating system features to set access privileges for the server certificate file and the server private key file so that only the database administrator has load privileges.

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

4.1.3 Distributing a CA Certificate File to the Client

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

4.1.4 Configuring the Operating Environment for the Database Server



See

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

4.1.5 Configuring the Operating Environment for the Client



See

.....
Refer to the following sections in the Application Development Guide for details, depending on your application development environment:

- "Settings for Encrypting Communication Data" under "Setup" in "JDBC Driver"
 - "Settings for Encrypting Communication Data" under "Setup" in "C Library (libpq)"
 - "Settings for Encrypting Communication Data" under "Setup" in "Embedded SQL in C"
-

4.1.6 Performing Database Multiplexing

When you perform communication that uses database multiplexing and a Secure Socket Layer server certificate, certificates with the same "Common Name" must be used. To ensure this, take one of the following actions:

- Create one server certificate, replicate it, and place a copy on each server used for database multiplexing.
- Create a server certificate with the same "Common Name" for each server used for database multiplexing.



See

.....
Refer to "Using the Application Connection Switch Feature" in the Application Development Guide for information on how to specify applications on the client.
.....

Chapter 5 Protecting Storage Data Using Transparent Data Encryption

This chapter describes how to encrypt data to be stored in the database.



See

.....
If you want to use an external key management system as the storage location for the encryption key, refer to "[Chapter 6 Using Transparent Data Encryption with Key Management Systems as Keystores](#)".
.....

5.1 Protecting Data Using Encryption

With PostgreSQL, data in a database is protected from access by unauthorized database users through the use of authentication and access controls. However, the OS file is not protected from attackers who bypass the database server's authentication and access controls.

With Fujitsu Enterprise Postgres, data inside the OS file is encrypted, so valuable information is protected even if the file or disk is stolen.

Data to be stored in a database is encrypted when it is written to the data file, and decrypted when it is read.

This is performed automatically by the instance, so the user and the application need not be aware of key management and encryption or decryption. This process is called TDE (Transparent Data Encryption).

The characteristics of TDE are described below.

Encryption mechanisms

Two-layer encryption key and the keystore

In each tablespace, there is a tablespace encryption key that encrypts and decrypts all the data within. The tablespace encryption key is encrypted by the master encryption key and saved.

Only one master encryption key exists in a database cluster. It is encrypted based on a passphrase specified by the user and stored in a keystore. Fujitsu Enterprise Postgres provides a file-based keystore. Attackers who do not know the passphrase cannot read the master encryption key from the keystore.

Strong encryption algorithms

TDE uses the Advanced Encryption Standard (AES) as its encryption algorithm. AES was adopted as a standard in 2002 by the United States Federal Government, and is used throughout the world.

Faster encryption and decryption based on hardware

TDE minimizes the overhead of encryption and decryption by using the AES-NI (Advanced Encryption Standard New Instructions) built into Intel(R) Xeon(R) processors since the 5600 series. This means that even in situations where previously the minimum encryption target was selected as a tradeoff between performance and security, it is now possible to encrypt all the data of an application.

Refer to the Intel Corporation's website for information on the list of processors equipped with AES-NI.

Zero overhead storage areas

Encryption does not change the size of data stored in tables, indexes, or WAL. There is, therefore, no need for additional estimates or disks.

Scope of encryption

All user data within the specified tablespace

The tablespace is the unit for specifying encryption. All tables, indexes, temporary tables, and temporary indexes created in the encrypted tablespace are encrypted. There is no need for the user to consider which tables and strings to encrypt.

Refer to "[5.4 Encrypting a Tablespace](#)" for details.

Backup data

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

Refer to "[5.7 Backing Up and Restoring/Recovering the Database](#)" after enabling or reconfiguring encryption to back up the database.

WAL and temporary files

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

Streaming replication support

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

The following are not encrypted:

- `pg_dump` and `pg_dumpall` output files
- Files output by the `COPY` command
- Notification event payloads that communicate using the `LISTEN` or `NOTIFY` command
- Checksum validation is not performed on encrypted tablespaces during backup and when using the `pg_checksum` utility.

5.2 Setting the Master Encryption Key

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

1. In the `keystore_location` parameter of `postgresql.conf`, specify the directory to store the keystore.

Specify a different location for each database cluster. 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

```
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 the `pgx_set_master_key` function to set the master encryption key.

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.

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.

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.

5.3 Opening the Keystore

To create encrypted tablespaces and access the encrypted data, you must first open the keystore. When you open the keystore, the master encryption key is loaded into the database server memory and becomes usable for encryption and decryption.

You need to open the keystore each time you start the instance. To open the keystore, the database superuser must execute the following SQL function.

```
SELECT pgx_open_keystore('passphrase');
```

The value "passphrase" is the passphrase specified during creation of the keystore.

Refer to "[B.2 Transparent Data Encryption Control Functions](#)" for information on the `pgx_open_keystore` function.

Note that, in the following cases, the passphrase must be entered when starting the instance, because the encrypted WAL must be decrypted for recovery. In this case, the above-mentioned `pgx_open_keystore` function cannot be executed.

- If performing crash recovery at the time of starting the instance
- If performing recovery using continuous archiving

For the above cases, specify the `--keystore-passphrase` option in the `pg_ctl` command, and then start the instance. This will display the prompt for the passphrase to be entered, as shown below.

```
> pg_ctl --keystore-passphrase start
Enter the passphrase:
The server is starting
>
```



When using an automatically opening keystore, you do not need to enter the passphrase and you can automatically open the keystore when the database server starts. Refer to "[5.6.3 Enabling Automatic Opening of the Keystore](#)" for details.

5.4 Encrypting a Tablespace

The keystore must be open before you can create an encrypted tablespace.

When creating a tablespace that will be encrypted, configure the encryption algorithm in the runtime parameters. For example, to create a tablespace with the name `secure_tablespace` using AES with a key length of 256 bits as the encryption algorithm, configure as shown below.

```
-- Specify the encryption algorithm for the tablespace to be created below
SET tablespace_encryption_algorithm = 'AES256';
CREATE TABLESPACE secure_tablespace LOCATION '/My/Data/Dir';
-- Specify that the tablespace to be created below is not to be encrypted
SET tablespace_encryption_algorithm = 'none';
```

Or

```
CREATE TABLESPACE secure_tablespace LOCATION '/My/Data/Dir' WITH (tablespace_encryption_algorithm = 'AES256' );
```

When the tablespace is empty, the encryption algorithm can be modified with the command below.

```
ALTER TABLESPACE secure_tablespace SET (tablespace_encryption_algorithm=AES256);
```

Trying to set the encryption algorithm for a non-empty tablespace causes an error.

You can use AES with a key length of 128 bits or 256 bits as the encryption algorithm. It is recommended that you use 256-bit AES. Refer to "[Appendix A Parameters](#)" for information on how to specify the runtime parameters.

If user provides both GUC and command line options while creating the tablespace, the preference is given to the command line option.

The `pg_default` and `pg_global` tablespaces cannot be encrypted.

Create tables and indexes in the encrypted tablespace that you created. Relations created in the encrypted tablespace are automatically encrypted.



Example

Example 1: Specifying an encrypted tablespace when creating it

```
CREATE TABLE my_table (...)  
    TABLESPACE secure_tablespace;
```

Example 2: Not explicitly specifying a tablespace when creating it and instead using the default tablespace

```
SET default_tablespace = 'secure_tablespace';  
CREATE TABLE my_table (...);
```

The process is the same for encrypting temporary tables and temporary indexes. In other words, either explicitly specify the `TABLESPACE` clause or list encrypted tablespaces in the `temp_tablespaces` parameter, and then execute `CREATE TEMPORARY TABLE` or `CREATE INDEX`.



Point

If an encrypted tablespace is specified in the `TABLESPACE` clause of the `CREATE DATABASE` statement, relations created in the database without explicitly specifying a tablespace will be encrypted. Furthermore, the system catalog will also be encrypted, so the source code of user-defined functions is also protected.

Example: Specifying a tablespace in a database definition statement

```
CREATE DATABASE DB01 TABLESPACE=SP01 ... ;
```

Part of the data is also stored in the system catalog - to encrypt this data as well, specify an encrypted tablespace as above and create a database.

An encrypted tablespace cannot be created from the window used for creating the pgAdmin tablespace, or from the query tool using a multi-command string. To create an encrypted tablespace, use psql console or a single command in the query tool such as that shown in the example below.

```
CREATE TABLESPACE tspace LOCATION '/some/dir/in/the/path' WITH  
(tablespace_encryption_algorithm='AES256')
```

5.5 Checking an Encrypted Tablespace

The `pgx_tablespaces` system view displays information about whether each tablespace has been encrypted, and about the encryption algorithm. Refer to "[D.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;
```



Example

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

spcname	spcencalgo
pg_default	none
pg_global	none
secure_tablespace	AES256

(3 rows)



See

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

5.6 Managing the Keystore

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

5.6.1 Changing the Master Encryption Key

Using the same encryption key for an extended period gives attackers an opportunity to decipher the encrypted data. It is recommended that you change the key at regular intervals, or whenever the key is exposed to risk.

Adhere to the industry's best practices for encryption algorithms and key management when considering how often the key should be changed. For example, the NIST in the United States has published "NIST Special Publication 800-57". The PCI DSS also refers to this publication. This publication recommends changing the master encryption key once a year.

To change the master encryption key, execute the `pgx_set_master_key` function, which is the same function used for configuring the key. Refer to "[5.2 Setting the Master Encryption Key](#)" for details.

After changing the master encryption key, you must immediately back up the keystore.

5.6.2 Changing the Keystore Passphrase

In security policies for organizations, it is usually a requirement that the passphrase be changed whenever a security administrator who knows the passphrase is removed from duties due to transfer or retirement. It is also recommended that the passphrase be changed if it is ever exposed to risks due to deception such as social engineering.

To change the keystore passphrase, execute the following SQL function as a superuser.

```
SELECT pgx_set_keystore_passphrase('oldPassphrase', 'newPassphrase');
```

After changing the passphrase, you must immediately back up the keystore.

Refer to "[B.2 Transparent Data Encryption Control Functions](#)" for information on the `pgx_set_keystore_passphrase` function.

5.6.3 Enabling Automatic Opening of the Keystore

When using an automatically opening keystore, you do not need to enter the passphrase and you can automatically open the keystore when the instance starts. Execute the `pgx_keystore` command to enable automatic opening of the keystore.

```
> pgx_keystore --enable-auto-open /key/store/location/keystore.ks
Enter the passphrase:
Automatic opening of the keystore is now enabled
>
```



Refer to "pgx_keystore" in the Reference for information on `pgx_keystore` command.

When automatic opening is enabled, an automatically opening keystore is created in the same directory as the original keystore. The file name of the automatically opening keystore is `keystore.aks`. The file `keystore.aks` is an obfuscated copy of the decrypted content of the `keystore.ks` file. As long as this file exists, there is no need to enter the passphrase to open the keystore when starting the instance.

Do not delete the original keystore file, `keystore.ks`. It is required for changing the master encryption key and the passphrase. When you change the master encryption key and the passphrase, `keystore.aks` is recreated from the original keystore file, `keystore.ks`.

Protect `keystore.ks`, `keystore.aks`, and the directory that stores the keystore so that only the user who starts the instance can access them.

Configure the permission of the files so that only the user who starts the instance can access the SQL functions and commands that create these files. Accordingly, manually configure the same permission mode if the files are restored.



```
# chown -R fsepuser:fsepuser /key/store/location
# chmod 700 /key/store/location
# chmod 600 /key/store/location/keystore.ks
# chmod 600 /key/store/location/keystore.aks
```

To use WebAdmin for backup and recovery, you must enable automatic opening of the keystore.

An automatically opening keystore will only open on the computer where it was created.

To disable automatic opening of the keystore, delete `keystore.aks`.

5.6.4 Backing Up and Recovering the Keystore

Back up the keystore at the following times in case it is corrupted or lost. Note that you must store the database and the keystore on separate data storage media. Storing both on the same data storage medium risks the danger of the encrypted data being deciphered if the medium is stolen. A passphrase is not required to open an automatically opening keystore, so store this type of keystore in a safe location.

- When the master encryption key is first configured
- When the master encryption key is changed
- When the database is backed up
- When the keystore passphrase is changed



Do not overwrite an old keystore when backing up a keystore. This is because during database recovery, you must restore the keystore to its state at the time of database backup. When the backup data of the database is no longer required, delete the corresponding keystore.

Example

- Back up the database and the keystore on March 1, 2022.

```
> pgx_dmpall -D /database/inst1
> cp -p /key/store/location/keystore.ks /keybackup/keystore_20220301.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 March 5, 2022.

```
> psql -c "SELECT pgx_set_master_key('passphrase')" postgres
> cp -p /key/store/location/keystore.ks /keybackup/keystore_20220305.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 March 5, 2022.

```
> cp -p /keybackup/keystore_20220305.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 March 2022.

```
> cp -p /keybackup/keystore_20220301.ks /key/store/location/keystore.ks
> pgx_rcvall -B /backup/inst1 -D /database/inst1 --keystore-passphrase
```

Specify the following in the `pgx_rcvall` command:

- Specify the data storage directory in the `-D` option. If the `-D` option is omitted, the value of the `PGDATA` environment variable is used by default.
 - Specify the backup data storage directory in the `-B` option.
 - The `--keystore-passphrase` option prompts you to enter the passphrase to open the keystore.
-

If you have restored the keystore, repeat the process of enabling automatic opening of the keystore. This ensures that the contents of the automatically opening keystore (`keystore.aks`) are identical to the contents of the restored keystore.

It is recommended that you do not back up the automatically opening keystore file, `keystore.aks`. If the database backup medium and the backup medium storing the automatically opening keystore are both stolen, the attacker will be able to read the data even without knowing the passphrase.

If the automatically opening keystore is corrupted or lost, you must again enable automatic opening. The `keystore.aks` file will be recreated from `keystore.ks` at this time.

See

Refer to "`pgx_rcvall`" and "`pgx_dmpall`" in the Reference for information on the `pgx_rcvall` and `pgx_dmpall` commands.

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

Refer to "[B.2 Transparent Data Encryption Control Functions](#)" for information on the `pgx_set_master_key` function.

Refer to "[5.6.3 Enabling Automatic Opening of the Keystore](#)" for information on how to enable automatic opening of the keystore.

5.7 Backing Up and Restoring/Recovering the Database

Fujitsu Enterprise Postgres enables you to use the five backup and recovery methods described below. Regardless of the method you use, you must back up the keystore at the same time.

Note that you must store the database and the keystore on separate data storage media. Storing both on the same data storage medium risks the danger of the encrypted data being deciphered if the medium is stolen.

Backup and recovery using WebAdmin

- Backup

WebAdmin backs up encrypted data.

Back up the key store after backing up the database.

- Recovery

Restore the keystore to its state at the time of database backup. Refer to "[5.6.4 Backing Up and Recovering the Keystore](#)" for details.

Enable automatic opening of the keystore in accordance with the procedure described in "[5.6.3 Enabling Automatic Opening of the Keystore](#)". Then, use WebAdmin to recover the database.

Backup and recovery using the `pgx_dmpall` and `pgx_rcvall` commands

- Backup

The `pgx_dmpall` command backs up the encrypted data.

Back up the key store after backing up the database.

- Recovery

Restore the keystore to its state at the time of the database backup.

Configure automatic opening of the key store as necessary.

If automatic opening of the keystore is not enabled, execute the `pgx_rcvall` command with the `--keystore-passphrase` option specified. This will display the prompt for the passphrase to be entered.



Example

- Back up the database and the keystore on March 1, 2022.

```
> pgx_dmpall -D /database/inst1
> cp -p /key/store/location/keystore.ks /keybackup/keystore_20220301.ks
```

Specify the following in the `pgx_dmpall` command:

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

- Recover the database and the keystore from the backup taken on March 1, 2022.

```
> cp -p /keybackup/keystore_20220301.ks /key/store/location/keystore.ks
> pgx_keystore --enable-auto-open /key/store/location/keystore.ks (Execute only when enabling
automatic opening)
> pgx_rcvall -B /backup/inst1 -D /database/inst1 --keystore-passphrase
```

Specify the following in the `pgx_rcvall` command:

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

- Specify the backup data storage directory in the -B option.
 - The --keystore-passphrase option prompts you to enter the passphrase to open the keystore.
-

Dump and restore using SQL

- Backup

The files output by the `pg_dump` and `pg_dumpall` commands are not encrypted. You should, therefore, encrypt the files using OpenSSL commands or other means before saving them, as described in "5.8 Importing and Exporting the Database" below.

Back up the key store after backing up the database.

- Restore

If the backup data has been encrypted using, for example Open SSL commands, decrypt that data.

The data generated by the `pg_dumpall` command includes a specification to encrypt tablespaces by default. For this reason, the `psql` command encrypts tablespaces during restoration.

File system level backup and restore

- Backup

Stop the instance and backup the data directory and the tablespace directory using the file copy command of the operating system. The files of encrypted tablespaces are backed up in the encrypted state.

Back up the key store after performing the backup.

- Restore

Restore the keystore to its state at the time of the database backup.

Stop the instance and restore the data directory and the tablespace directory using the file copy command of the operating system.

Continuous archiving and point-in-time recovery

- Backup

The `pg_basebackup` command backs up the encrypted data as is.

Back up the key store after performing the backup.

- Recovery

Restore the keystore to its state at the time of the database backup.

Configure automatic opening of the key store as necessary.

If automatic opening of the keystore is not enabled, execute the `pg_ctl` command to start the instance with the `--keystore-passphrase` option specified. This will display the prompt for the passphrase to be entered.



See

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

- pg_dumpall

If you have restored the keystore, repeat the process of enabling automatic opening of the keystore. This ensures that the contents of the automatically opening keystore (keystore.aks) are identical to the contents of the restored keystore.

Refer to "[5.6.3 Enabling Automatic Opening of the Keystore](#)" for information on how to enable automatic opening of the keystore.

5.8 Importing and Exporting the Database

The files output by the COPY TO command are not encrypted. Therefore, when transferring files to other systems, you should encrypt files using OpenSSL commands or other means and use scp or sftp to encrypt the data being transferred.

Use a safe method to delete obsolete plain text files.

You can use the following methods to safely delete files:

- shred command



Example

```
# Export the contents of the table my_table to a CSV file.
> psql -c "COPY my_table TO '/tmp/my_table.csv' (FORMAT CSV)" postgres

# Encrypt the exported file.
> openssl enc -e -aes256 -in my_table.csv -out my_table.csv.enc
(The user is prompted to enter the passphrase to be used for encryption)

# Safely delete plain text files.
> shred -u -x my_table.csv
(Transfer encrypted files to other systems)

# Decrypt the encrypted files on other systems.
> openssl enc -d -aes256 -in my_table.csv.enc -out my_table.csv
(The user is prompted to enter the passphrase to be used for decryption)
```

If you use COPY FROM to import data to tables and indexes in an encrypted tablespace, the imported data is automatically encrypted before being stored.

5.9 Encrypting Existing Data

You cannot encrypt existing unencrypted tablespaces. In addition, you cannot change encrypted tablespaces so that they do not encrypt.

As an alternative, transfer the tables and indexes to other tablespaces. You can use the following SQL commands for this.

```
ALTER TABLE table_name SET TABLESPACE new_tablespace;
ALTER INDEX index_name SET TABLESPACE new_tablespace;
ALTER DATABASE database_name SET TABLESPACE new_tablespace;
```



See

Refer to "SQL Commands" under "Reference" in the PostgreSQL Documentation for information on SQL commands.

5.10 Operations in Cluster Systems

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

5.10.1 HA Clusters that do not Use Database Multiplexing

Take the following points into account when using transparent data encryption in an HA cluster environment that does not use database multiplexing.

Placement and automatic opening of the keystore file

There are two alternatives for placing the keystore file:

- Sharing the keystore file
- Placing a copy of the keystore file

Sharing the keystore file

This involves using the same keystore file on the primary server and the standby server.

As the standby server is not active while the primary server is running, this file would not be accessed simultaneously, and therefore, it can be shared.

To manage the keystore file in a more secure manner, place it on the key management server or the key management storage isolated in a secure location.

Enable the automatic opening of the keystore on both the primary and standby servers.

Placing a copy of the keystore file

This involves placing a copy of the primary server keystore file on the standby server.

You can do this if you cannot prepare a shared server or disk device that can be accessed from both the primary and standby servers.

However, if you change the master encryption key and the passphrase on the primary server, you must copy the keystore file to the standby server again.

To manage the keystore file in a more secure manner, prepare the key management server or the key management storage isolated in a secure location for both the primary and standby servers, and place the keystore files there.

Enable the automatic opening of the keystore on both the primary and standby servers. Note that copying the automatically opening keystore file (keystore.aks) to the standby server does not enable the automatic opening of the keystore.



See

.....
Refer to the Cluster Operation Guide (PRIMECLUSTER) for information on building a cluster system environment for performing failover using the failover feature integrated with the cluster software.
.....

5.10.2 Database Multiplexing Mode

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

Placing the keystore file

Place a copy of the primary server keystore file on the standby server.

This is required as the keystore file cannot be shared, and both servers may need to access it simultaneously.



Point

.....
To manage the keystore file in a more secure manner, place it on the key management server or the key management storage isolated in a secure location. A keystore used by both the primary and standby servers can be managed on the same key management server or key management storage.
.....

.....
However, create different directories for the keystores to be used by the primary server and the standby server. Then copy the keystore for the primary server to the directory used on the standby server.
.....

Automatically opening the keystore

You must enable automatic opening of the keystore.

To do this, enable automatic opening of the keystore in all servers that make up database multiplexing. The settings for automatic opening of the keystore include information unique to each server, so simply copying the file does not enable it.

Changing the passphrase

Changes to the passphrase are reflected in all servers that make up database multiplexing, so no special operation is required.

Building and starting a standby server

Before using the `pg_basebackup` command or `pgx_rcvall` command to build a standby server, copy the keystore file from the primary server to the standby server. When using an automatically opening keystore, use the copied keystore file to enable automatic opening on the standby server.

Open the keystore each time you start the standby server. This step is necessary for decrypting and restoring encrypted WAL received from the primary server. To open the keystore, specify the `--keystore-passphrase` option in the `pg_ctl` command or `pgx_rcvall` command and enter the passphrase, or use an automatically opening keystore.

Changing the master encryption key and the passphrase

Change the master encryption key and the passphrase on the primary server. You need not copy the keystore from the primary server to the standby server. You need not even restart the standby server or reopen the keystore. Changes to the master encryption key and the passphrase are reflected in the keystore on the standby server.



See

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

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

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

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

5.11 Security-Related Notes

- Decrypted data is cached in the database server memory (shared buffer). As a result, unencrypted data is stored in a core file, which is a process memory dump. You should, therefore, safely delete the memory dump.
You can safely delete files by using the following command:
 - `shred` command
- Unencrypted data may be written from the database server memory to the operating system's swap area. To prevent leakage of information from the swap area, consider either disabling the use of swap area or encrypting the swap area using a full-disk encryption product.
- The content of the server log file is not encrypted. Therefore, in some cases the value of a constant specified in a SQL statement is output to the server log file. To prevent this, consider setting a parameter such as `log_min_error_statement`.
- When executing an SQL function that opens the keystore and modifies the master encryption key, ensure that the SQL statement containing the passphrase is not output to the server log file. To prevent this, consider setting a parameter such as `log_min_error_statement`. If you are executing this type of SQL function on a different computer from the database server, encrypt the communication between the client and the database server with SSL.
- Starting with FEP 10, logical replication is available, which allows non-backed up clusters to subscribe to databases where transparent data encryption is enabled. Logical replication does not need to have the same encryption strategy between publisher and subscriber.

In this scenario, if the user wants to encrypt the subscribed copy of data as well, then it is the user's responsibility to create encryption policies to the subscribed databases. By default, published encrypted tablespace data will not be encrypted in the subscriber side.

5.12 Tips for Installing Built Applications

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

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

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

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

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

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

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

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

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

Chapter 6 Using Transparent Data Encryption with Key Management Systems as Keystores

This chapter describes the operation of transparent data encryption when a key management system is used as a keystore.



Refer to "Key Management System Requirements" in the Installation and Setup Guide for Server for the key management system requirements that can be used with Fujitsu Enterprise Postgres.

6.1 Protecting Data Using Encryption

Refer to "5.1 Protecting Data Using Encryption". The following describes the differences from the transparent data encryption operation in the file-based keystore described in "5.1 Protecting Data Using Encryption".

Encryption mechanisms

Two-layer encryption key and the keystore

Each tablespace has a tablespace encryption key that encrypts/decrypts all data in it. Tablespace encryption keys are stored encrypted with the master encryption key.

Use an encryption key stored in a key management system as a common master encryption key for your database cluster. Fujitsu Enterprise Postgres refers to the key management system as a keystore for master encryption keys.

Type of key management system

Two types of key management systems are available:

- kmip

It is a key management system that can be used using a protocol called KMIP (Key Management Interoperability Protocol) standardized by OASIS (Organization for the Advancement of Structured Information Standards).

- custom

It is a key management system that cooperates using an adapter that converts the request format without adopting the KMIP protocol.

A sample of the adapter plugin required for this type is stored in the following location under the Fujitsu Enterprise Postgres installation directory.

- For Amazon Web Services (AWS)
`<Install directory>/share/aw-kms-plugin.sh.sample`
- For Microsoft Azure (Azure)
`<Install directory>/share/az-kms-plugin.sh.sample`

If you use the provided AWS sample plug-in as is, the following conditions apply.

Item	Contents
Available services	By using the AWS adapter, you can use encryption keys on the Key Management Service (KMS) provided by AWS. There are no regional restrictions as long as the region is supported by AWS KMS.
Available AWS KMS keys	The key spec for a KMS key (key spec) must be a symmetric cryptographic key. Asymmetric cryptographic keys cannot be used. Also, the KMS key usage (key usage) must be ENCRYPT_DECRYPT.
Required permissions	The user accessing AWS KMS must be permitted to perform the following operations for the KMS key to be used.

Item	Contens
	<ul style="list-style-type: none"> - Encrypt - Decrypt - DescribeKey
Key ID	<p>The following can be specified as the key ID.</p> <ul style="list-style-type: none"> - Key ARN
Dependent package	<p>Install AWS CLI on the Fujitsu Enterprise Postgres server. For details, refer to the AWS CLI manual. In addition, the following packages are required.</p> <ul style="list-style-type: none"> - jq <p>The plugin is executed by the OS user who starts the Fujitsu Enterprise Postgres server. You must set the PATH or modify the script file so that the OS user can execute the aws command and jq command.</p>
CLI configuration	<p>Configure the CLI so that aws commands executed by the OS user that starts the Fujitsu Enterprise Postgres server can access the AWS key management service without entering additional credentials.</p> <p>Configuring in this way allows keystores that use this plugin to be opened without entering a KMS secret.</p>

If you use the provided Azure sample plug-in as is, the following conditions apply.

Item	Contens
Available services	The Azure adapter allows you to use any key management service that is accessible through the Azure Key Vault API and that can use symmetric keys.
Available keys	Symmetric key is available.
Available algorithms	<p>The following algorithms are available for encryption/decryption operations.</p> <ul style="list-style-type: none"> - A256GCM
Key operations	<p>The user accessing the Azure key management service must be permitted to perform the following operations for the key to be used.</p> <ul style="list-style-type: none"> - encrypt - decrypt - get
Key ID	<p>The following can be specified as the key ID.</p> <ul style="list-style-type: none"> - Key object identifier
Dependent package	<p>Install Azure CLI on the Fujitsu Enterprise Postgres server. For details, refer to the Azure CLI manual. In addition, the following packages are required.</p> <ul style="list-style-type: none"> - jq <p>The plugin is executed by the OS user who starts the Fujitsu Enterprise Postgres server. You must set the PATH or modify the script file so that the OS user can execute the az command and jq command.</p>
Sign in	Sign in to Azure using a service principal. You need the application ID, tenant ID, and credentials to sign in. The available authentication methods are password authentication and certificate-based authentication.
Opening the Keystore	<p>To open the keystore, you must specify the following as the KMS secret:</p> <ul style="list-style-type: none"> - Password (for password authentication)

Item	Contents
	- Private key passphrase (for certificate-based authentication)



See

Refer to "Key Management System Requirements" in the Installation and Setup Guide for Server for the key management system requirements that can be used with Fujitsu Enterprise Postgres.

Sharing tablespace encryption keys

When using an adapter to link with a key management system, encryption and decryption of the tablespace encryption key using the master encryption key are performed on the key management system side.

Tablespace encryption keys can be shared within a database cluster so that you do not need to access the key management system each time you want to use the tablespace encryption key.

The cost of encryption/decryption using the master encryption key becomes an issue in the following cases:

- Multiple connections to the database access encrypted tablespaces
- Connections accessing encrypted tablespaces are repeated and connection pooling is disabled

Encryption key identifier

Key IDs are used as identifiers to identify encryption keys stored on the key management system.

Key ID

Information that identifies the encryption key on the key management system, and is unique within the key management system. The correspondence between the encryption key substance (byte string) and the key ID does not change throughout the life cycle of the encryption key.

The name of the identifier differs depending on each key management system, but in this feature, such information is called the key ID.

Changes to the key management system

After starting operation of the transparent data encryption function, the key management system to be used can be changed to another key management system.

6.2 Setting the Master Encryption Key

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

1. Load the `shared_preload_libraries` parameter in `postgresql.conf` with the library name "tde_kms"

```
shared_preload_libraries = 'tde_kms'
```

2. When using an adapter, register the adapter as a plug-in. Specify the directory where the plugin is stored in the `tde_kms.plugin_path` parameter in `postgresql.conf`. Store your plugins in this directory. If you want to use the samples, copy them into this directory. The plugin file requires execution privilege for the OS user that starts the Fujitsu Enterprise Postgres server.

```
tde_kms.plugin_path = '/home/fsepuser/plugin/'
```

3. To share the tablespace encryption key, set the `tde_kms.enable_shared_dek` parameter in `postgresql.conf` to "on".

```
tde_kms.enable_shared_dek = on
```

4. Set the `tde_kms.kms_conninfo_file` parameter in `postgresql.conf` to a file that contains key management system connection information. Refer to "[Appendix A Parameters](#)" for information.

Example for the key management system connection information file `kms_conninfo.conf`

```
tde_kms.kms_conninfo_file = 'kms_conninfo.conf'
```

Example of key management system connection information file

For type kmip

```
kmip mykmipsvr mykmipsvr.example.com 5696 cert sslcert=postgres.crt
sslkey=postgres.key sslrootcert=root.crt
```

For type custom (When using the AWS sample plugin)

```
custom mykms aw-kms-plugin.sh arg=--profile arg=user1
```

For type custom (When using the Azure sample plugin)

```
custom mykms az-kms-plugin.sh kms-secret-obf=password.ksc arg=--auth-method arg=password
arg=--user-id arg=ApplicationId arg=--tenant arg=TenantId arg=--algorithm arg=A256GCM
```

5. Execute a CREATE EXTENSION statement to install the extension.

```
CREATE EXTENSION tde_kms;
```

6. To enable transparent data encryption, call the `pgx_declare_external_master_key` function to declare the encryption key to use as the master encryption key. Specify a key ID as an identifier to identify the encryption key. Refer to "[B.2.3 pgx_declare_external_master_key](#)" for information on the `pgx_declare_external_master_key` function.

```
SELECT pgx_declare_external_master_key( kms_name => 'mykmipsvr', key_id =>
'a0eebc99-9c0b-0000-0000-000000000000', sslpassphrase => 'mykmippassphrase' );
```

6.3 Opening the Keystore

To create encrypted tablespaces and access the encrypted data, you must first open the keystore. When you open the keystore, the master encryption key is available 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( sslpassphrase => 'passphrase');
```

passphrase is the passphrase of the private key file for the client certificate.

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 `--kms-secret` 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 --kms-secret start
Enter secret:
```

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 "[6.6.2 Enabling Automatic Opening of the Keystore](#)" for details.

6.4 Encrypting a Tablespace

Refer to "[5.4 Encrypting a Tablespace](#)".

6.5 Checking an Encrypted Tablespace

Refer to ["5.5 Checking an Encrypted Tablespace"](#).

6.6 Managing the Keystore

Describes how to manage master encryption keys when a key management system is used.

6.6.1 Changing the Master Encryption Key

To change the master encryption key, run the `pgx_declare_external_master_key` function as you originally set it. Refer to ["6.2 Setting the Master Encryption Key"](#) for more information.

Also, by specifying a different key management system name than the key management system name specified during installation, you can change the key management system to be used. Refer to ["6.6.5 Changes to the Key Management System"](#) for details.

6.6.2 Enabling Automatic Opening of the Keystore

You can automatically open a keystore at instance startup without entering a passphrase by specifying all credentials, including those that should be kept secret, in the key management system connection information file. To enable automatic keystore opening, run the `pgx_keystore` command.

Example of storing obfuscated credentials in the file `sslkeypassphrase.ksc`

```
> pgx_keystore -s -o sslkeypassphrase.ksc
Enter secret:
```

Specify obfuscated credentials in the key management system connection information file.

```
kmip      mykmipsvr      kmip.example.com      5696      cert      sslcert=postgres.crt      sslkey=postgres.key
sslrootcert=root.crt      sslkeypassphrase-obf=sslkeypassphrase.ksc
```

The key management system connection information file is valid only on the computer on which it was created.

To disable automatic keystore opening, delete the file containing obfuscated credentials for the private key specified in `sslkeypassphrase-obf` and delete the `sslkeypassphrase-obf` option in the key management system connection information file.



See

.....

Refer to `"pgx_keystore"` in the Reference for information on `pgx_keystore` command.

Refer to ["Appendix A Parameters"](#) for information on the key management system connection information file.

.....

6.6.3 Changing Credentials for Key Management Systems

If the credentials for the key management service change, you must also change the credentials that Fujitsu Enterprise Postgres uses to connect to the key management system.

You can change the credentials used by Fujitsu Enterprise Postgres using the `pgx_open_keystore` function. The new credentials are used to connect to the new key management system.

If you have a streaming replication configuration, you must change credentials on all replicas.

Refer to ["B.2.1 pgx_open_keystore"](#) for information on the `pgx_open_keystore` function.

6.6.4 Verifying the Master Encryption Key

The view `pgx_tde_master_key` shows information about your master encryption key. For more information about columns, Refer to ["E.1 pgx_tde_master_key"](#) for more information about columns..

6.6.5 Changes to the Key Management System

If you need to change the linked key management system after installing the transparent data encryption function that uses the key management system, you can do so by following the procedure below.

After this step, data in Fujitsu Enterprise Postgres will be encrypted using encryption keys on the new key management system.

Defining a new key management system

If the new key management system is not listed in the key management system's connection information file, add a definition to that configuration file. Give the old key management system and the new key management system different key management system names.

Reload the configuration file for the changes to take effect.

Declaring the master encryption key to use

Use the `pgx_dexlare_external_master_key` function to declare a new encryption key to use. Specify the name you gave the new key management system as the key management system name. Other arguments required are the key ID to use on the new key management system and credentials. Upon successful completion, it will be encrypted using the encryption key on the new key management system.

Information

The master encryption key that encrypts the backup data before changing the key management system exists only on the old key management system. As such, you need the old key management system and the encryption keys residing there for any period of time when you might restore backup data from before the change. If you delete the encryption key on the old key management system, destroy the old key management system, cancel the key management service, etc., you will not be able to decrypt the backup data and you will not be able to use the data.

6.7 Backing Up and Restoring/Recovering the Database

Fujitsu Enterprise Postgres enables you to use the five backup and recovery methods described below.

Backup and recovery using WebAdmin

- Backup

WebAdmin backs up encrypted data.

- Recovery

If you recover to a point in time when the old master encryption key was in use, change to the most recent master encryption key immediately after recovery.

Enable automatic opening of the keystore in accordance with the procedure described in "[6.6.2 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.

- Recovery

If you recover to a point in time when the old master encryption key was in use, change to the most recent master encryption key immediately after recovery.

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 `--kms-secret` option specified. This will display the prompt for the passphrase to be entered.

Dump and restore using SQL

- Backup

The files output by the `pg_dump` and `pg_dumpall` commands are not encrypted. You should, therefore, encrypt the files using OpenSSL commands or other means before saving them, as described in "[5.8 Importing and Exporting the Database](#)" below.

- 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.

- Restore

Stop the instance and use the OS file copy command to restore the data storage directory or tablespace directory.

If you recover to a point in time when the old master encryption key was in use, change to the most recent master encryption key immediately after recovery.

Continuous archiving and point-in-time recovery

- Backup

The `pg_basebackup` command backs up the encrypted data as is.

- Recovery

If you recover to a point in time when the old master encryption key was in use, change to the most recent master encryption key immediately after recovery.

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 `--kms-secret` option specified. This will display the prompt for the passphrase to be entered.



See

.....

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

- Refer to "Reference" in the PostgreSQL Documentation for information on the following commands:

- `psql`
- `pg_dump`
- `pg_basebackup`

- Refer to the Reference for information on the following commands:

- `pgx_rcvall`
 - `pgx_dmpall`
 - `pg_dumpall`
-

6.8 Importing and Exporting the Database

Refer to "[5.8 Importing and Exporting the Database](#)".

6.9 Encrypting Existing Data

Refer to "[5.9 Encrypting Existing Data](#)".

6.10 Operations in Cluster Systems

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

6.10.1 HA Clusters that do not Use Database Multiplexing

Take the following points when using transparent data encryption with a key management system as a key store in an HA cluster environment that does not use database multiplexing.

Placement and automatic opening of the connection information file of the key management system

The file that describes the connection information of the key management system specified in the `tde_kms.kms_conninfo_file` parameter of the `postgresql.conf` file, and the files such as certificates that are referenced from that file can be shared by the primary server and the standby server. However, the obfuscated credential file used to enable automatic opening of the keystore must be created and placed on each server according to the instructions for enabling automatic opening.

If not shared, it must be possible to connect to the key management system used from the standby server with the same key management system name as the key management system name set on the primary server. Place the connection information file and files such as certificates referenced from the connection information file on the standby server. The obfuscated credential file used to enable automatic opening of the keystore must be created and placed on each server according to the instructions in enabling automatic opening.

Changing credentials for key management systems

If the credentials for the key management system have changed, use the `pgx_open_keystore` function on the primary server to change the credentials. Re-enable automatic opening of the keystore on the standby server.



See

.....
Refer to the Cluster Operation Guide (PRIMECLUSTER) for information on building a cluster system environment for performing failover operations linked to cluster software.
.....

6.10.2 Database Multiplexing Mode

Note the following when using transparent data encryption with a key management system as a key store in environments that use streaming replication, or database multiplexing with streaming replication.

Placement and automatic opening of the connection information file of the key management system

The file that describes the connection information of the key management system specified in the `tde_kms.kms_conninfo_file` parameter of the `postgresql.conf` file, and the files such as certificates that are referenced from that file can be shared by the primary server and the standby server. However, the obfuscated credential file used to enable automatic opening of the keystore must be created and placed on each server according to the instructions for enabling automatic opening.

If not shared, it must be possible to connect to the key management system to be used with the same key management system name as the key management system name set on the primary server. Place the connection information file and files such as certificates referenced from the connection information file on all servers that configure database multiplexing mode. The obfuscated credential file used to enable automatic opening of the keystore must be created and placed on each server according to the instructions in enabling automatic opening.

Changing credentials for key management systems

If the credentials for the key management system are changed, use the `pgx_open_keystore` function on all servers that configure database multiplexing to change the credentials.

Starting a standby server

Open the keystore when starting the standby server. This is required to decrypt and replay the encrypted WAL received from the primary server. To open a keystore, use the `pg_ctl` or `pgx_rcvall` command with `--kms-secret` and provide your credentials, or enable automatic opening of the keystore.

Changing the master encryption key

Change the master encryption key on the primary server. No need to restart the standby server or reopen the keystore. Changes to the master encryption key are also reflected 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 "High Availability, Load Balancing, and Replication" under "Server Administration" in the PostgreSQL Documentation for information on how to set up streaming replication.

6.11 Security-Related Notes

Refer to "[5.11 Security-Related Notes](#)".

6.12 Tips for Installing Built Applications

Refer to "[5.12 Tips for Installing Built Applications](#)".

6.13 Reference Information for Linking Key Management Systems Using Sample Plugins

Sample Plugin for AWS

See below for CLI configuration:

<https://docs.aws.amazon.com/cli/latest/userguide/cli-configure-files.html>

Sample Plugin for Azure

For key management services that support symmetric keys, see below.

<https://learn.microsoft.com/en-us/azure/key-vault/keys/about-keys#key-types-and-protection-methods>

For more information about Azure Key Management Services, see below.

<https://learn.microsoft.com/en-us/azure/security/fundamentals/key-management#azure-key-management-services>

For more information about service principals, see below.

<https://learn.microsoft.com/en-us/cli/azure/create-an-azure-service-principal-azure-cli#4-sign-in-using-a-service-principal>

Options that can be specified for AWS plugin

The following additional options can be specified for the plugin. These options can be specified as extra-args in the KMS connection information file.

`--config config-file`: Specify the path of the AWS CLI configuration file. If omitted, the default path of the AWS CLI will be used.

`--credentials credentials-file`: Specify the path to the AWS CLI authentication information file. If omitted, the default path of the AWS CLI will be used.

`--profile profile-name`: Specify the profile to use in the AWS CLI configuration file and credentials file. If omitted, the AWS CLI default profile will be used.

Options that can be specified for the Azure plugin

The following additional options can be specified for the plugin. These options can be specified as extra-args in the KMS connection information file.

- auth-method (password|cert): Specify the authentication method. (password: Password authentication, cert: Certificate-Based Authentication)
- user-id user-id: Specify the application ID.
- user-cert cert-file: For certificate-based authentication, specify the path to the certificate file.
- tenant tenant-id: Specify the tenant ID.
- algorithm algorithm: Specify the algorithm to be used.

Plugin Errors

If an error occurs in the operation of the plugin, a message will be output to the server log.

Verifying access to encryption keys

You can view information about the encryption key being used by executing the following command as the OS user running the Fujitsu Enterprise Postgres server.


- When using the sample plug-in for AWS
aws kms describe-key --key-id *Key ID*
- When using the sample plug-in for Azure
az keyvault key show --id *Key ID*

Notes on using the Azure sample plug-in

The service principal may need to sign in periodically. If the OS user who starts the Fujitsu Enterprise Postgres server is already signed in to Azure due to the environment settings, the credentials passed from Fujitsu Enterprise Postgres will not be used when using the plugin. They will only be used when periodic sign-in becomes necessary. If the KMS secret specified when opening the KMS connection information file or keystore is incorrect, the error will be detected when periodic sign-in becomes necessary, so check in advance that the specified information is correct.

Chapter 7 Policy-based Login Security

To apply a policy for login security to a user, define the policy as a profile and assign the profile to the user. The contents of the profile are saved as database objects.



The items that can be set for a profile are described below.

Managing Dormant Users

You can automatically lock users who have not been connected to the database for a long time.

This setting is for users with the LOGIN attribute.

Managing Policies When Using Password Authentication

You can set the following policies for users who use password authentication (password, md5, scram-sha-256).

- Set a password life time
- Restrict password reuse
- Lock accounts that have failed to login continuously
- Allow passwords to be set in encrypted form
- Set the gradual password rollover time

Gradual password rollover is when you change a password and then keep the old password in effect for a while.

This setting specifies the valid period.

This is useful, for example, if it is difficult to make a new password available system-wide instantly.

7.1 Advance Preparation

Ensure that the postgres database exists and that you can connect to it.

If no connections are possible, and if using a database other than the default postgres database, specify the name of the database to which you can connect in the `userprofile_database` parameter in the `postgresql.conf` file.

Refer to "[Parameters for the Policy-based Login Security](#)" for parameters.

7.2 Changing the Contents of the default Profile

When you create a database user with `CREATE ROLE`, the database user is assigned the default profile. When you create the default profile, all password configuration and authentication restrictions and the gradual password rollover feature are disabled, so change the parameter values to suit your policy. Change the value with the `pgx_alter_profile` function.

Users with `CREATEROLE` privilege can modify the contents of a profile.

Refer to "[B.3.1 Profile Management Functions](#)" for functions.

Refer to "7.7 Profile parameters" for more information on parameters.

[Example]

```
SELECT pgx_alter_profile('default',
'{
  "INACTIVE_USER_TIME": 30,
  "PASSWORD_LIFE_TIME": 50,
  "PASSWORD_GRACE_TIME": 10,
  "PASSWORD_REUSE_MAX": 5,
  "PASSWORD_LOCK_TIME": 0.5,
  "PASSWORD_ALLOW_HASHED": true,
  "PASSWORD_ROLLOVER_TIME": 0.125
}');
```

7.3 Creating and Assigning Profiles

If you want to apply a different policy than the default profile, create a new profile and assign it to the user.

The `pgx_create_profile` function creates a profile.

To assign a profile to a user, use the `pgx_assign_profile_to_user` function.

Users with `CREATEROLE` privilege can create and assign profiles.

Refer to "B.3.1 Profile Management Functions" and "B.3.2 User Management Functions" for functions.

7.4 Actions to be Taken in the Event of Deviation from Policy

When a user is locked

Locked in the following cases:

- Restriction by `INACTIVE_USER_TIME`
- Restriction by `FAILED_LOGIN_ATTEMPTS`

If you want to allow a locked user to login again, unlock the user. Unlocking is performed using the `pgx_unlock_user` function. Operations can be performed by users with `CREATEROLE` privilege.

Refer to "B.3.2 User Management Functions" for functions.

[Example]

```
SELECT pgx_unlock_user('user1');
```

To release a lock when a database administrator is locked, users other than the locked user with `CREATEROLE` privilege can do so. If the user does not exist, Fujitsu Enterprise Postgres must be started in single-user mode and the lock released.



Point

.....
You can set `PASSWORD_LOCK_TIME` to automatically release locks due to `FAILED_LOGIN_ATTEMPTS`.
.....

When the password expires

When a password life time is over, the password expires and cannot be used to connect to and operate on the database until the password is changed.

The password can be changed by the expired user itself, or by a user who has the `CREATEROLE` privilege and who is an administrator of that user (who has `ADMIN` privilege for that user).

7.5 Settings in Streaming Replication Configuration

For streaming replication, the policy is enabled on both the primary and standby servers.

Perform changing the contents of the profile, creating and assigning profiles, unlock, and change the password can only be done on the primary server.

This section describes how to setting streaming replication configuration.

1. Grant privilege to users for streaming replication

The standby server connects to the primary server and propagates any state changes it detects to the primary server. The used users for streaming replication (specified in `primary_conninfo`) is used to connect for state change. Grant privilege to this user. Allow the user to SET ROLE directly to the group role `pgx_update_profile_status`.

[Example]

```
# GRANT pgx_update_profile_status TO repluser WITH SET TRUE;
```

This action enables users for streaming replication to lock all other database users or unlock the lock state due to policy deviations. Also, membership in the `pgx_update_profile_status` group role should only be granted to users for streaming replication.

2. Add a record to `pg_hba.conf`

Add a record to accept connections from the standby server. Duplicate the record for streaming replication and specify the database name in the `userprofile_database` parameter as the destination database name. Authentication methods other than password authentication are recommended.

```
host replication repluser standbyServerAddress authenticationMethod # For Streaming
Replication
host postgres repluser standbyServerAddress authenticationMethod # For state change
propagation
```

Information

- If the primary and standby servers are disconnected, profile-based restrictions still apply on each server. However, the standby server itself cannot change its password or unlock it explicitly. After removing the cause of the disconnect, restore the connection to the primary server and business operations, and then change the password or explicitly unlock.
 - Connection information for reflecting status changes to the primary server or standby server is output as "User profile status sender" to `application_name` in the `pg_stat_replication` view. When monitoring the streaming replication status with the `pg_stat_replication` view, exclude the "User profile status sender" line from the monitoring target.
-

7.6 Backup and Recovery

Backup

Profile contents, user and profile assignment status, and password history can be backed up with physical backups and the `pg_dumpall` command.

When dumping with SQL statements using the `pg_dumpall` command, database role information is backed up at the same time it is dumped.

Information

When the `pg_dumpall` command is run without the `--no-role-passwords` option, the database user's password is included in the backup file in the form of an encrypted form. Therefore, when recovering the database user's password, temporarily set this parameter to true in the default profile of the restore destination. If the restore is successful, the contents of the default profile are restored to the contents at the time of the backup.

Recovery

If you recover using older backup data, you may be considered a dormant user at the time of recovery, or your password may have expired. In this case, unlock or change the password after recovery.

7.7 Profile parameters

Details of the parameters set in the profile are explained.

INACTIVE_USER_TIME

Number of days before auto-locking users who have not been connected to the database for a long time (users who cannot see their sessions)

[Supported values]

integer: A INTEGER value greater than or equal to 1

The unit is days. The maximum value is 24855 days.

DEFAULT: The value of the same parameter in the default profile

UNLIMITED: No auto lock

This setting is for users with the LOGIN attribute.

If a session to the database cannot be verified for at least INACTIVE_USER_TIME, the user is locked out and cannot log in.

Sessions are checked into the database every hour and the information is saved.

If the system shuts down before the save, it is assumed that there were no logins between the last save and the shutdown. As a result, the lock period might be shorter than the value specified in this parameter.

When using streaming replication, the session confirmation to the database is performed every hour on each server. The primary server makes a locking decision while the standby server sends session confirmation information to the upstream server. This can cause a time lag of several hours before the session confirmation information from the standby server is communicated to the primary server. This lag can result in locking even when connected to a standby server.

You can determine when each user's session was last seen by the system by looking at the value of the userprlastactivetime column in the pgx_user_profile system catalog.

If the device is locked by the setting of this parameter, use the pgx_unlock_user function to unlock the device. Automatic cancellation by PASSWORD_LOCK_TIME is not performed.

When a logged-in user changes to another role with SET ROLE, the new role is not considered logged in to the database. It is assumed that the old role is still logged in to the database.

PASSWORD_LIFE_TIME

Number of days the same password can be used for authentication

[Supported values]

numeric: A NUMERIC value greater than or equal to 0

The unit is days. Hours and seconds can be specified with decimal places (e.g. 4.5 is equivalent to "4 days and 12 hours"). Precision is 1 second. The maximum value is 24855 days.

DEFAULT: The value of the same parameter in the default profile

UNLIMITED: No life time (same password can be used indefinitely)

The password life time will be over after PASSWORD_LIFE_TIME days from the last time the password was updated for the target user. The timing of updating the profile is not the starting point. Therefore, if you specify an extremely short number of days (such as 1 day), it may already be past the life time at the time of renewal.

It is possible to specify when a password becomes invalid using the VALID UNTIL clause of CREATE ROLE or ALTER ROLE. If you specify both the VALID UNTIL clause and PASSWORD_LIFE_TIME, both values are valid. Note that in this case, you can login only if both constraints are met.

If there is no grace period, the password expires when the password life time is over. If you login using password authentication in this state, you will receive a "password expired" warning and you will not be able to execute commands other than changing your password. You can change the password to resume normal operations. Other than password authentication, you can connect to and work with the database.

For streaming replication, users will not be able to connect to the standby server after the password life time is over. There is no grace period for password expiration on standby servers. Password changes must be made on the primary server.

PASSWORD_GRACE_TIME

The number of days after a password life time is over before the password expires.

[Supported values]

numeric: A NUMERIC value greater than or equal to 0

The unit is days. Hours and seconds can be specified with decimal places (e.g. 4.5 is equivalent to "4 days and 12 hours"). Precision is 1 second. The maximum value is 24855 days.

DEFAULT: The value of the same parameter in the default profile

UNLIMITED: Indefinite period

The password expiration grace period is `PASSWORD_GRACE_TIME` days from the first login time after the password life time is over. You can login with your current password during the grace period, but a warning prompts you to change your password. You can operate normally except that a warning is displayed. If specified as **UNLIMITED**, the grace period is infinite and you will be warned to change your password at every login. Once transitioned to the grace period, even if you change the profile value after that, you cannot go back to before the grace period, and you cannot change the life time. If 0 is specified, there is no grace period and the password expires at the first login after the password life time is over. In this case, you can perform normal operations again by changing the password.

Because the standby server does not have a grace period before the password expires, users cannot connect to the standby server after the password life time is over. Password changes must be made on the primary server.

PASSWORD_REUSE_TIME

Number of days the same password cannot be reused

The password cannot be reused by the same user for this period from the time the password is updated.

This parameter must be set in combination with `PASSWORD_REUSE_MAX`.

[Supported values]

numeric: A NUMERIC value greater than or equal to 0

The unit is days. Hours and seconds can be specified with decimal places (e.g. 4.5 is equivalent to "4 days and 12 hours"). Precision is 1 second. The maximum value is 24855 days.

DEFAULT: The value of the same parameter in the default profile

UNLIMITED: Not reusable (However, if `PASSWORD_REUSE_MAX` is also **UNLIMITED**, the password can be reused without restriction)

PASSWORD_REUSE_MAX

Number of password changes required before the same password can be reused

This parameter must be set in combination with `PASSWORD_REUSE_TIME`.

[Supported values]

integer: An INTEGER value greater than or equal to 0

DEFAULT: The value of the same parameter in the default profile

UNLIMITED: Not reusable (However, if `PASSWORD_REUSE_TIME` is also **UNLIMITED**, the password can be reused without restriction)

Both `PASSWORD_REUSE_TIME` and `PASSWORD_REUSE_MAX` constraints must be met to reuse passwords.

For example, if you specify `PASSWORD_REUSE_TIME = 30`, `PASSWORD_REUSE_MAX = 10`, a certain password can be reused if 30 days have passed since the update time and the password has been updated 10 times or more. When password update fails due to these parameters, the `SQLSTATE` will be "22023: invalid_parameter_value".

Also, if one parameter has a value and the other parameter specifies `UNLIMITED`, the password cannot be reused. If you want to use only one condition for reuse judgment, you need to set the value of the unused parameter to 0. However, if `UNLIMITED` is specified for both, these parameters are ignored and passwords can be reused without restriction.

When changing a password, if the password is specified in `MD5` or `SCRAM` encrypted format, you cannot check that the password is being reused.

PASSWORD_ALLOW_HASHED

Whether to allow passwords to be specified in `MD5` or `SCRAM` encrypted form when changing passwords Allow if true.

[Supported values]

boolean: true or false

DEFAULT: The value of the same parameter in the default profile

If true, there are no restrictions on how to set or change passwords.

However, note the following when changing passwords in encrypted form:

- Unable to check for password reuse (`PASSWORD_REUSE_TIME`, `PASSWORD_REUSE_MAX`)
- Password rollover disabled (`PASSWORD_ROLLOVER_TIME`)
- Inability to check password complexity using the extension `passwordcheck`

If false, the password can only be set or changed by specifying the password in clear text in the `PASSWORD` clause of a `CREATE ROLE` or `ALTER ROLE` statement. You cannot change the password using the `psql` command `\password` meta-command.

In this case, all password checks that cannot be performed in encrypted form are possible.



Note

Passwords specified in `CREATE ROLE` or `ALTER ROLE` statements may, depending on configuration, be logged in the `psql` command history and the server log.

FAILED_LOGIN_ATTEMPTS

Number of consecutive failed login attempts allowed by the user

[Supported values]

integer: An `INTEGER` value greater than 0

DEFAULT: The value of the same parameter in the default profile

`UNLIMITED`: Indefinite period (Can fail any number of times)

If password authentication fails consecutively for the number of times specified by this parameter, the user is locked and cannot login.

The number of failed login attempts is counted separately on each server.

PASSWORD_LOCK_TIME

Number of days after a user is locked due to consecutive login failures before the user is unlocked

[Supported values]

numeric: A `NUMERIC` value greater than or equal to 0

The unit is days. Hours and seconds can be specified with decimal places (e.g. 4.5 is equivalent to "4 days and 12 hours"). Precision is 1 second. The maximum value is 24855 days.

DEFAULT: The value of the same parameter in the default profile

UNLIMITED: Locked indefinitely

If set to UNLIMITED, the user will be locked indefinitely and will not be automatically unlocked. Unlocking requires an explicit `pgx_unlock_user` function call by a user with CREATEROLE privilege.

PASSWORD_ROLLOVER_TIME

Number of days after password change before old password expires

[Supported values]

numeric: A NUMERIC value greater than or equal to 0

The unit is days. Hours and seconds can be specified with decimal places (e.g. 4.5 is equivalent to "4 days and 12 hours"). Precision is 1 second. The maximum value is 60 days. The minimum value is 0.0416 days (approximately 1 hour).

DEFAULT: The value of the same parameter in the default profile

From the time the user changes the password until PASSWORD_ROLLOVER_TIME has elapsed, the user can login using the old password.

Specify a value equal to or less than the lesser of the PASSWORD_GRACE_TIME (except when 0 is specified) or PASSWORD_LIFE_TIME. If a larger value is specified, the smaller of these parameters is used. For example, if PASSWORD_GRACE_TIME is 10 and PASSWORD_LIFE_TIME is 50, specify 10 or less for this parameter. If 15 is specified in this parameter, 10 is assumed to be specified.

If 0 is specified, the user cannot login with the old password. The default profile has an initial value of 0.

To expire the combinable period immediately before PASSWORD_ROLLOVER_TIME expires, a user or a user with CREATEROLE privilege can call the `pgx_make_password_rollover_expire` function.

If you execute the ALTER ROLE statement with PASSWORD NULL, password authentication is disabled. Therefore, you cannot log in with the old password.

In either of the following cases, the password rollover is disabled and the previous password expires immediately:

- When changing the password, the password is specified in MD5 or SCRAM encrypted format.
 - The password method has changed since the last time the password was set.
- The relevant parameters are:

- `password_encryption`
- `scram_iterations`

7.8 Worker Processes

This feature allows up to two background worker processes to reside. This information is displayed in the statistics view `pg_stat_activity` when it needs to be processed by a worker process. The worker processes added by this feature are those with column `backend_type` values of "user profile status writer" and "user profile status sender".

These processes may wait on the wait event "UserProFileWorkerMain" of type "Activity".

Chapter 8 Data Masking

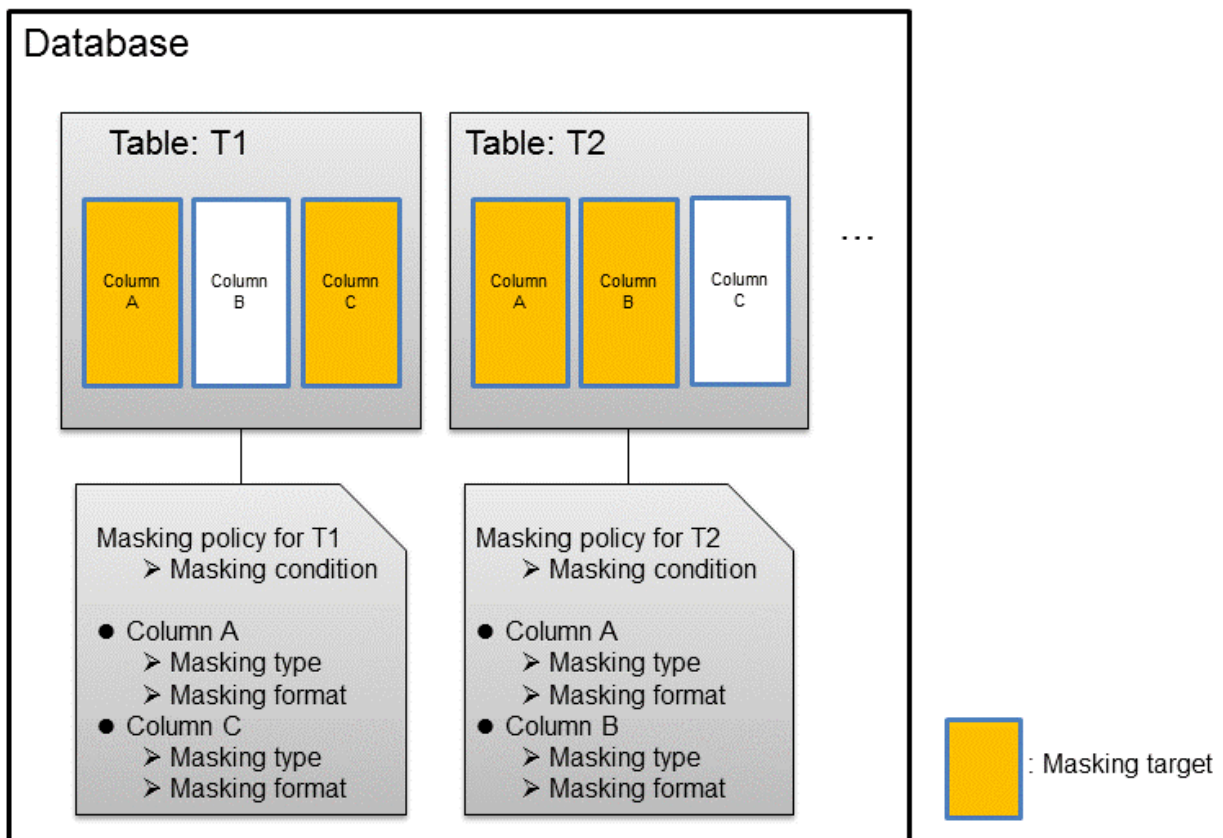
Data masking is a feature that can change the returned data for queries generated by applications, so that it can be referenced by users. For example, for a query of employee data, digits except the last four digits of an eight-digit employee number can be changed to "*" so that it can be used for reference.

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.

8.1 Masking Policy

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

Figure 8.1 Masking policy



Information

When a masking policy is defined, the search performance for the corresponding table may deteriorate.

8.1.1 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

Information

- 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.

8.1.2 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.

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.

The following masking types can be specified and selected depending on the masking target data type. When masking the data that includes multibyte characters, do not specify partial masking for masking type. The result may not be as expected.

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.

8.1.3 Masking Condition

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

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

Information

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

8.1.4 Masking Format

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

Full masking

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




See

Refer to "8.3 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 <code>function_parameters</code>
Numeric type	<p><i>'replacementCharacter, startPosition, endPosition'</i></p> <ul style="list-style-type: none"> - <i>replacementCharacter</i>: Specify the number to display. Specify a value from 0 to 9. - <i>startPosition</i>: 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>. <p> Example</p> <p>Specify as below to change the values from the 1st to 5th digits to 9.</p> <pre>function_parameters := '9, 1, 5'</pre> <p>In this example, if the original data is "123456789", it will be changed to "999996789".</p>
Character type	<p><i>'inputFormat, outputFormat, replacementCharacter, startPosition, endPosition'</i></p> <ul style="list-style-type: none"> - <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. - <i>endPosition</i>: Specify the position of "V" as an end position of masking. When working out the end position, do not include positions of "F". For example, to specify the position of the 11th "V" from the left, specify 11. Specify a positive integer that is greater than <i>startPosition</i>. <p> Example</p> <p>Specify as below to mask a telephone number other than the first three digits using *.</p> <pre>function_parameters := 'VVVFVVVVFVVVV, VVV-VVVV-VVVV, *, 4, 11'</pre> <p>In this example, if the original data is "012-3156-7890", it will be changed to "012-****-****".</p>
Date/timestamp type	'MDYHMS'

Category	Method of specifying function_parameters
	<ul style="list-style-type: none"> - 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. <p> Example</p> <p>Specify as below to mask hour, minute, and second and display 00:00:00.</p> <pre>function_parameters := 'MDYh0m0s0'</pre> <p>In this example, if the original data is "2022-02-10 10:10:10", it will be changed to "2022-02-10 00:00:00".</p>

 **See**

- Refer to "B.4.2 pgx_create_confidential_policy" for information on function_parameters.
- Refer to "8.3 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.

When column data type is `character(n)` or `char(n)` and if the string length after change exceeds `n`, the extra characters will be truncated and only characters up to the `n`th character will be displayed.

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

 **Example**

Specify as below to change all three characters starting from `b` to `X`.

```
regexp_pattern := 'b..'
regexp_replacement:= 'X'
regexp_flags := 'g'
```

In this example, if the original data is "foobarbaz", it will be changed to "fooXX".



See

- Refer to "POSIX Regular Expressions" in the PostgreSQL Documentation and check pattern, replacement, and flags for information on the values that can be specified for `regex_pattern`, `regex_replacement`, and `regex_flags`.
- Refer to "8.3 Data Types for Masking" for information on the data types for which masking can be performed.

8.2 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.
 - Be sure to add "pgx_datamasking" to the beginning of the parameter "shared_preload_libraries".
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
```



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. Only database superusers can configure masking policies.

The masking policy setting is explained in the following steps using an example.

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

8.2.1 Creating a Masking Policy

Create a masking policy. 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.

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

1. Create a masking policy

Execute the `pgx_create_confidential_policy` system management function to create a masking policy.

The following values are configured in this example.

- Masking target: Numeric type c1
- Masking type: FULL
- Masking condition: '1=1'

```
postgres=# select pgx_create_confidential_policy(table_name := 't1', policy_name := 'p1',
expression := '1=1', column_name := 'c1', function_type := 'FULL');
pgx_create_confidential_policy
-----
t
(1 row)
```

2. Confirm the displayed data

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

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



See

- Refer to "[B.4.2 pgx_create_confidential_policy](#)" for information on the `pgx_create_confidential_policy` system management function.

8.2.2 Changing a Masking Policy

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

1. Change a masking policy

Execute the `pgx_alter_confidential_policy` system management function to change a masking policy.

The following values are changed in this example.

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

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

2. Confirm the displayed data
Confirm that the masking target data has been correctly changed.

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



See

- Refer to "B.4.1 [pgx_alter_confidential_policy](#)" for information on the `pgx_alter_confidential_policy` system management function.

8.2.3 Confirming a Masking Policy

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

1. 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 |
function_parameters | regexp_pattern | regexp_replacement | regexp_flags |
column_description
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----
 public      | t1         | p1          | c1          | FULL         |
 public      | t1         | p1          | c2          | PARTIAL      | VVVVVVVVVVVVVVV, VVV-VVVV-
 VVVV, *, 4, 11 |
(2 row)
```

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

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



See

- Refer to "F.1 [pgx_confidential_columns](#)" for information on the `pgx_confidential_columns` table.
- Refer to "F.2 [pgx_confidential_policies](#)" for information on the `pgx_confidential_policies` table.

8.2.4 Enabling and Disabling a Masking Policy

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

1. Disable a masking policy
Execute the `pgx_enable_confidential_policy` system management function to disable a masking policy.

```
postgres=# select pgx_enable_confidential_policy(table_name := 't1', policy_name := 'p1',
enable := 'f');
```

```
pgx_enable_confidential_policy
-----
t
(1 row)
```

2. Confirm the displayed data

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

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

3. Enable a masking policy

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

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

4. Confirm the displayed data

Confirm that the masking target data has been correctly changed.

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



See

- Refer to "[B.4.4 pgx_enable_confidential_policy](#)" for information on the `pgx_enable_confidential_policy` system management function.

8.2.5 Deleting a Masking Policy

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

1. Delete a masking policy

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

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

2. Confirm the displayed data

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

```
postgres=# select * from t1;
 c1 |      c2
-----+-----
```

1		012-3456-7890
2		012-3456-7891
3		012-3456-7892
(3 row)		



.....
- Refer to "[B.4.3 pgx_drop_confidential_policy](#)" for information on the `gx_drop_confidential_policy` function.
.....

8.3 Data Types for Masking

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

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

Category	Data type	Masking type		
		Full masking	Partial masking	Regular expression masking
Numeric type	<code>smallint</code>	Y	Y	N
	<code>integer</code>	Y	Y	N
	<code>bigint</code>	Y	Y	N
	<code>decimal</code>	Y	Y	N
	<code>numeric</code>	Y	Y	N
	<code>float</code>	Y	Y	N
	<code>real</code>	Y	Y	N
	<code>double precision</code>	Y	Y	N
Character type	<code>character varying(n)</code>	Y	Y	Y
	<code>varchar(n)</code>	Y	Y	Y
	<code>character(n)</code>	Y	Y	Y
	<code>char(n)</code>	Y	Y	Y
Date/timestamp type	<code>date</code>	Y	Y	N
	<code>timestamp</code>	Y	Y	N

8.4 Security Notes

- Starting with FEP 10, logical replication is available, which allows non-backed up clusters to subscribe to databases where data masking policies are enabled. Logical replication allows publisher and subscriber databases to have their own or the same data masking policies.

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

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

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

Chapter 9 Periodic Operations

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

Information

Among routine operations, those that can be implemented with scripts can be automated using cron, while those that can be implemented as SQL statements or stored procedures can be automated using pg_cron.

9.1 Configuring and Monitoring the Log

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

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

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

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

See

- Refer to "Error Reporting and Logging" under "Server Administration" in the PostgreSQL Documentation for information on logs.
- Refer to "Configuring Parameters" in the Installation and Setup Guide for Server for information on log settings when operating with WebAdmin.

9.2 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

9.2.1 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%

9.2.2 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 "9.5 Reorganizing Indexes" for details.
- To secure space on the backup data storage destination disk:
Execute backup using WebAdmin or the pgx_dmpall command.

9.3 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. If a value other than 0 is specified, 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.	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.



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

9.4 Monitoring the Connection State of an Application

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

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



See

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

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

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

- [9.4.1 Using the View \(pg_stat_activity\)](#)
- [9.4.2 Using pgAdmin](#)

9.4.1 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 backend_type = 'client backend' and state='idle in
transaction' and current_timestamp > cast(query_start + interval '60 minutes' as timestamp);
-[ RECORD 1 ]-----+-----
datid          | 13003
datname        | db01
pid            | 4638
leader_pid     |
usesysid       | 10
username       | fsep
application_name | apl01
client_addr    | 192.33.44.15
client_hostname |
client_port    | 27500
backend_start  | 2022-02-24 09:09:21.730641+09
xact_start     | 2022-02-24 09:09:23.858727+09
query_start    | 2022-02-24 09:09:23.858727+09
state_change   | 2022-02-24 09:09:23.858834+09
wait_event_type | Client
wait_event     | ClientRead
state          | idle in transaction
backend_xid    |
backend_xmin   |
query_id       |
query         | begin;
backend_type   | client backend
```



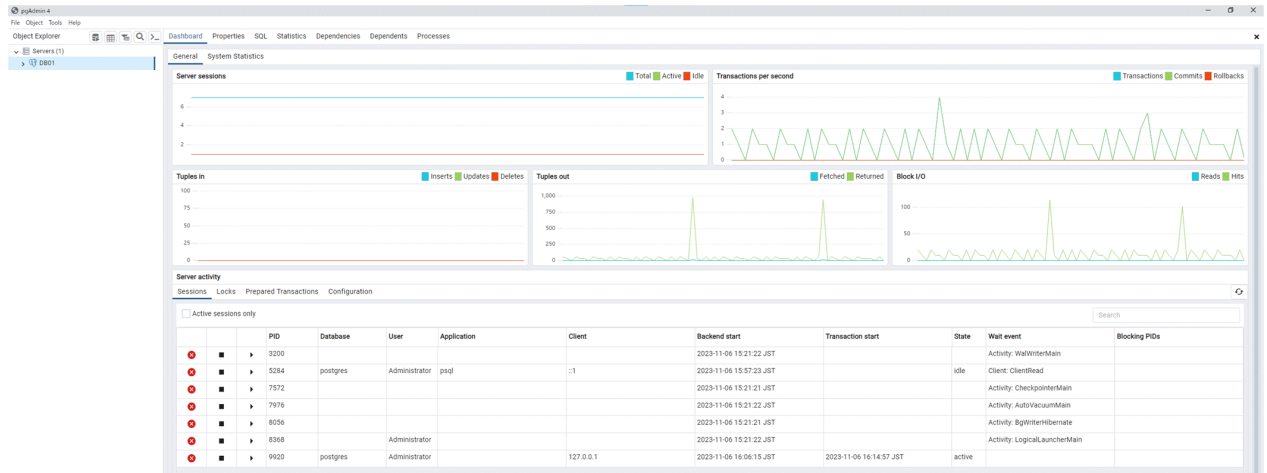
See

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

9.4.2 Using pgAdmin

This section describes the procedure for monitoring connections using [Server Status] in pgAdmin.

1. In the [Browser] pane, click the database server for monitoring.
2. In the [Dashboard] tab, identify client connections that have been in the waiting state for an extended period.



9.5 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. 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.

Check the disk usage of the data storage destination using the method described in "9.2 Monitoring Disk Usage and Securing Free Space".



See

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



Point

Typically, reorganize indexes once a month at a suitable time such as when conducting database maintenance. Use SQL statements to check index usage. If this usage is increasing on a daily basis, adjust the frequency of recreating the index as compared to the free disk space.

The following example shows the SQL statements and the output.

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

[SQL statements]

```
SELECT
  nspname AS schema_name,
  relname AS index_name,
  round(100 * pg_relation_size(indexrelid) / pg_relation_size(indrelid)) / 100 AS index_ratio,
  pg_size_pretty(pg_relation_size(indexrelid)) AS index_size,
  pg_size_pretty(pg_relation_size(indrelid)) AS table_size
FROM pg_index I
```

```

LEFT JOIN pg_class C ON (C.oid = I.indexrelid)
LEFT JOIN pg_namespace N ON (N.oid = C.relnamespace)
WHERE
  C.relkind = 'i' AND
  pg_relation_size(indrelid) > 0
ORDER BY pg_relation_size(indexrelid) DESC, index_ratio DESC;

```

[Output]

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



See

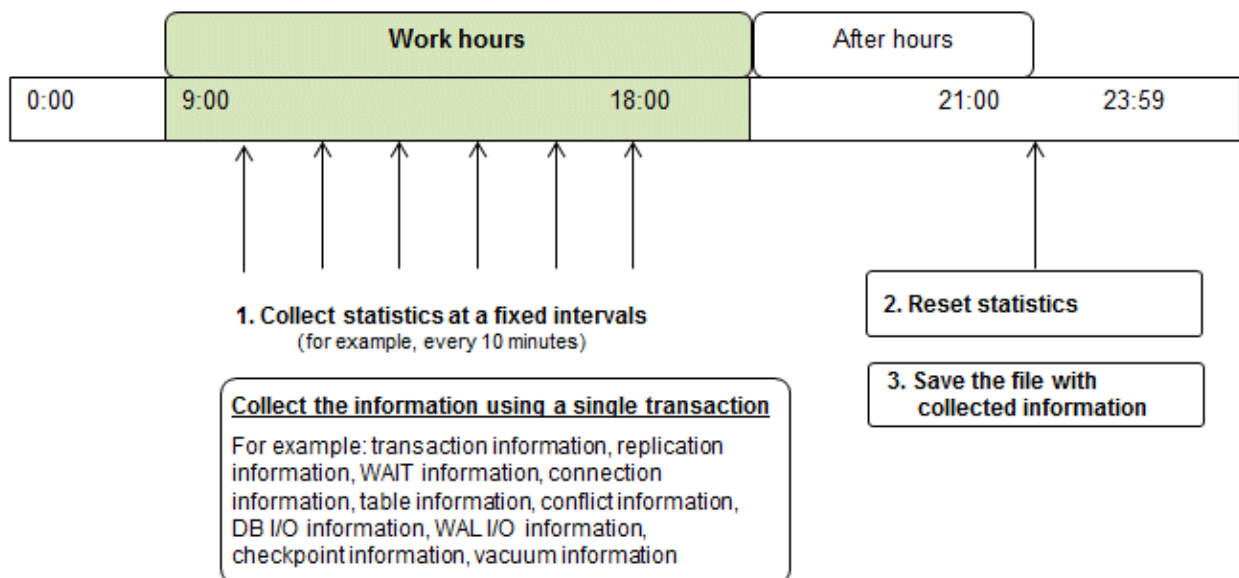
Refer to "Notes on Application Compatibility" in the Application Development Guide for information on maintaining application compatibility.

9.6 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 9.1 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 "9.6.1 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 "9.6.3 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

9.6.1 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 9.1 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 "D.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 "D.3 pgx_stat_latch" for details.
pgx_stat_walwriter	Displays statistics related to WAL writing, in a single line. Refer to "D.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 "D.5 pgx_stat_sql" for details.
pgx_stat_gmc	Displays statistics related to Global Meta Cache hit ration and used memory size. Refer to "D.6 pgx_stat_gmc" for detail. Also refer to Chapter 16 Global Meta Cache for information on the Global Meta Cache.

9.6.2 Collection Configuration

The procedure for configuring collection depends on the information content.

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

Information common to PostgreSQL



See

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

Information added by Fujitsu Enterprise Postgres

Information added by Fujitsu Enterprise Postgres is collected by default.

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

View name	Parameter
pgx_stat_lwlock	track_waits (*1)
pgx_stat_latch	
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.

9.6.3 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

Function	Type of return value	Description
		<ul style="list-style-type: none"> - If <code>pg_stat_reset_shared('walwriter')</code> is called: All values displayed in <code>pgx_stat_walwriter</code> - If <code>pg_stat_reset_shared('sql')</code> is called: All values displayed in <code>pgx_stat_sql</code> - If <code>pg_stat_reset_shared('gmc')</code> is called: All values except size column in <code>pgx_stat_gmc</code>



See

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

9.7 Scheduling of an Aggressive Freeze for Tuples (VACUUM FREEZE)

An architectural limitation of PostgreSQL is that it requires transaction ID reclamation via tuple freezing to avoid the problem of disruptive transaction wraparound. For more information about transaction wraparound issues, refer to "Preventing Transaction ID Wraparound Failures" in the PostgreSQL Documentation.

Tuple freezing works with autovacuum and `VACUUM FREEZE`(aggressive freeze for tuples). However, on a system that consumes transactions quickly, autovacuum alone might be enough to run into wraparound problems. This is because autovacuum is slow to reduce frequency of collisions with applications and increase system load. Such systems must perform an aggressive tuple freeze at the appropriate time.

Fujitsu Enterprise Postgres provides scripts to perform efficient aggressive freeze for tuples.

This section describes how to monitor whether the freezing process is on time in a running system, and how to estimate the running time of the script.



Point

If you insert or update a large amount of data at once, the next freeze process will take a long time, so be sure to perform aggressive freeze for tuples after such operations.

9.7.1 Monitoring Trends in Transaction ID Usage


Execute the following SQL to monitor the trend in transaction ID usage.

```
# SELECT max(age(datfrozenxid)) FROM pg_database WHERE dataallowconn = true;
```

If autovacuum does not sufficiently reduce transaction ID usage, schedule aggressive freeze for tuples.

Determine whether the amount of waste reduced by automatic vacuum is sufficient

Looking at the entire interval that includes several autovacuum freeze processes, the transaction ID usage will be graphed like a sawtooth, sandwiching the `autovacuum_freeze_max_age` setting value, depending on the start and end of the autovacuum operation. If the usage peaks are gradually getting larger, the reduction may be insufficient. In particular, if it no longer falls below `autovacuum_freeze_max_age`, the reduction is definitely insufficient.



9.7.2 How to Schedule Aggressive Freeze for Tuples

1. Determine a time period with low business load when aggressive freeze for tuples will have minimal impact.
2. Copy the script that implements aggressive freeze for tuples and set the connection information, multiplicity, and timeout value parameters.

For streaming replication configuration, specify multiple connection destinations in PGHOST and PGPORT, and specify "read-write" or "primary" in PGTARGETSESSIONATTRS to connect to the primary server.

The scripts are stored in the following location under the Fujitsu Enterprise Postgres installation directory:

```
<Install directory>/share/execute_freeze.sh.sample
```

3. Register the script from step 2 in the job scheduler.
Specify the execution timing according to the planned time period.

Example: Use cron to run a script called execute_freeze.sh at 2:05 AM

```
5 2 * * 0 sh /path/to/execute_freeze.sh
```

When a schedule review is required

Even after aggressive freeze for tuples is scheduled, continue to monitor the trend in transaction ID usage. If you determine that the freezing process is not keeping up, the schedule must be reviewed. To collect the statistical information required for the review, enable the extended functionality for aggressive freeze for tuples using the following procedure.

1. Modify the shared_preload_libraries parameter in postgresql.conf

```
shared_preload_libraries = 'pgx_stat_vacuum_freeze'
```

2. Execute the following SQL

```
SELECT datname FROM pg_database WHERE datallowconn = true;
```

Execute the following CREATE EXTENSION for all databases to be output and define the extension.

```
CREATE EXTENSION pgx_stat_vacuum_freeze;
```

Monitor the trends in transaction usage, and if you determine that the freezing process is not keeping up, refer to "[9.7.3 Tuning the Allocation Time for Aggressive Freeze for Tuples](#)" to calculate the time required for freezing tuples and add more allocation time.

9.7.3 Tuning the Allocation Time for Aggressive Freeze for Tuples

The interval between autovacuum runs is treated as one interval, and aggressive freeze for tuples is scheduled to be completed before the next autovacuum run. For example, if your environment is one week where freezing by autovacuum occurs, determine the required time within that interval.

Calculate the time required for aggressive freeze for tuples as follows, based on the time it took to freeze transaction IDs by past autovacuum stored in the [pgx_stat_freeze_results](#) table for each database you are using.

In this example, the interval for autovacuum freezing is set to one week. Perform this for each database you operate and calculate the total.

```
# SELECT sum(total_elapsed_time) - sum(sleep_times) as total_time FROM pgx_stat_freeze_results WHERE
measurement_start_time >= '2024/10/21 00:00:00' AND measurement_end_time < '2024/10/28 00:00:00' AND
autovacuum = true;
```

If you are unable to secure sufficient time when the workload is low, increase the degree of parallelism. You can specify the degree of parallelism in the script.

9.8 Database Monitoring with Amazon CloudWatch

Amazon CloudWatch is a monitoring tool provided by Amazon Web Services (hereinafter referred to as AWS). Fujitsu Enterprise Postgres works with Amazon CloudWatch to collect database-related metrics and logs. The collected metrics can be used to monitor the status of the database and create alert notifications when thresholds are exceeded. In addition, by analyzing the collected logs on CloudWatch, it is possible to detect abnormalities in logins and operations.

9.8.1 Metrics Collection

Metrics collection uses psql commands to retrieve metrics and the AWS CLI to forward metrics to Amazon CloudWatch.

1. Preparation

Refer to the AWS documentation, install the AWS CLI, and configure it to transfer metrics to Amazon CloudWatch.

2. Create a metrics collection script

Sample scripts for collecting metrics are stored in the following location under the Fujitsu Enterprise Postgres installation directory. Create a script for collecting metrics by referring to the sample script. For the metrics collected by the sample script, refer to "[9.8.3 Metrics Collected by the Sample Script](#)".

```
<Install directory>/share/putting-metrics-to-monitoring-tool.sh.sample
```

Edit the following connection information defined in the sample script.

- Amazon CloudWatch connection destination to store metrics

Define the Amazon CloudWatch Namespace to collect metrics in "NAMESPACE".

```
NAMESPACE="YOUR Cloud Watch NAMESPACE"
```

- Database connection destination for collecting metrics

Edit the variables below to define the PostgreSQL information for retrieving metrics.

PGHOST: Specify the host name to connect to.

PGPORT: Specify the port to connect to.

PGDATABASE: Specify the database name to connect to.

PGUSER: Specify the PostgreSQL user to use for connection. To view all statistical information for all sessions, specify a superuser or a role with the built-in role `pg_read_all_stats` privileges.

3. Setting environment variables

To run the metrics collection script, you must set the path where the Fujitsu Enterprise Postgres server feature and AWS CLI are installed in the environment variable `PATH`.

The following is an example of how to set environment variables when the AWS CLI is installed in `/awscli/bin`.

If you want to set environment variables within a script, write them at the beginning of the script before the environment variable for the database connection destination.

```
export PATH=/opt/fsepv<x>server64/bin:/awscli/bin:$PATH
```

4. Define the database connection credentials

If necessary, define the password and the storage location of the SSL communication certificate in an environment variable or configuration file.

5. Periodically run metrics collection scripts

Use crontab or similar to periodically run the metrics collection script as the OS user who can transfer metrics using the aws cloudwatch put-metric-data command.

Below is an example of a crontab definition for collecting metrics every 10 minutes.

```
$ cat /var/spool/cron/<Running User>  
*/10 * * * * /<Script placement destination>/putting-metrics-to-monitoring-tool.sh
```

Edit the script to run periodically using the "crontab -e" command.

When running with crontab, you can use rsyslog to output the script's standard output to a log file and detect whether the script is running normally or abnormally.

The metrics collection interval should be determined based on the purpose of collection. The following are examples of collection intervals.

- Long intervals, such as once every 10 minutes or once an hour, are recommended when monitoring to check operational trends.
- Short intervals, such as once a minute, are recommended when monitoring to quickly respond to increases in load, such as a sudden increase in the number of connections.

If the collection interval is short, the load may increase because the number of queries to the database increases accordingly. Consider the business load when deciding on the execution interval and the queries to be executed for metric collection.

9.8.2 Log Collection

You can collect logs using the CloudWatch agent. For instructions on how to collect logs, see the Amazon CloudWatch manual.

9.8.3 Metrics Collected by the Sample Script

In the sample script, metrics items are created from the following perspectives to regularly check whether the database cluster is operating properly.

- Is the database capacity sufficient?
- Are your SQL performance and connections adequate?
- Are each node operating normally?

Table 9.2 List of metrics collected by the sample script

Metrics name	Metrics
pg_stat_bgwriter_*	Maps displayed in statistics collection
pg_stat_database_*	Maps displayed in statistics collection
pg_stat_database_conflicts_*	Maps displayed in statistics collection
pg_stat_archiver_*	Maps displayed in statistics collection
pg_stat_activity_*	Maps displayed in statistics collection
pg_stat_replication_*	Maps displayed in statistics collection
pg_replication_slots_*	Mapping to the pg_replication_slots system catalog
pg_locks_*	Mapping to the pg_locks system catalog
pg_capacity_connection_*	Connection metrics (e.g. txns running for 1 hour)
pg_capacity_schema_*	Schema disk space metrics
pg_capacity_tblspace_*	Tablespace disk space metrics

Metrics name	Metrics
pg_capacity_tblvacuum_*	Metrics for tables that haven't been vacuumed in days
pg_capacity_longtx_*	Number of transactions running for more than 5 minutes
pg_capacity_datfrozenxid_*	Trends in transaction ID usage (Used for scheduling aggressive freeze for tuples (VACUUM FREEZE))
pg_performance_locking_detail_*	Blocked process details
pg_performance_locking_*	Number of blocked processes
pg_stat_user_tables_*	Vital statistics from pg_stat_user_tables
pg_statio_user_tables_*	Vital statistics from pg_statio_user_tables
pg_stat_statements_*	Statistics for SQL statements executed by the server
pg_capacity_tblbloat_*	Table fetch bloat

9.9 Monitoring Deferred SQL and Periodically Backing up statistics

After production starts, SQL execution can be delayed due to a variety of factors, including unexpected system load and data bias. You can detect delays early by monitoring for delays that affect operations. Monitoring can also reduce the time required for recovery tasks, such as determining the cause and resolving delays.

One cause of deferred SQL is a change in statistics. Some delays caused by statistics can be resolved by restoring the statistics before the delay occurred. For this reason, back up your statistics periodically to guard against delays.

Operational workflow

The workflow required to monitor delayed SQL and back up statistics is shown below.

Corresponding time	Work item	Work details
During environment construction	(1) Enabling the pg_dbms_stats extension	Enable the pg_dbms_stats extension with the Create Extension statement.
	(2) Logging settings for delayed SQL monitoring	Set the log_min_duration_statement parameter etc. in postgresql.conf.
During system development	(3) Creating and reviewing a script for backing up statistics (*1)	Provide a script to back up statistics according to the statistics and Vacuum operation.
During operation	(4) Periodic backup of statistics	Execute the script created in (3) to back up the statistics.
	(5) Monitoring delayed SQL and adjusting the allowable delay time	Monitor the occurrence of delayed SQL and adjust the allowable delay time to find the optimal settings for operation.
When delayed SQL occurs	(6) Investigating the cause of delayed SQL	If a delayed SQL occurs, we will investigate the cause of the SQL delay and consider countermeasures.
	(7) Restoring statistics (*2)	Delayed SQL is eliminated by restoring and fixing the backed up statistics.

*1: Although it is possible to perform the work during operation, please consider the impact on operation and perform the work before starting operation.

*2: If it is determined that the cause is not statistics, take appropriate measures according to the cause.

The specific work content is shown below.

(1) Enabling the pg_dbms_stats extension

To back up and restore statistics, you need the export and import functions of the pg_dbms_stats extension, so enable the pg_dbms_stats extension.

For instructions on how to enable it, refer to "Setting Up pg_dbms_stats" in the Installation and Setup Guide for Server.

(2) Logging settings for delayed SQL monitoring

Database logs are used to monitor delayed SQL.

Make the following settings in postgresql.conf. Note that the date and time are important when investigating the cause of delayed SQL, set the format and file name to include the date and time.

Parameter	Specified value
logging_collector	on
log_line_prefix	Specify the following format and add information required for investigation, such as the time and executed application. [%t]%u %d %p[%l]
log_filename	Since you will need to check past execution records, set it so that you can see the date and time. postgresql.%Y-%m-%d
log_min_duration_statement	300000



Point

.....
If the value specified for the log_min_duration_statement parameter is small, even SQL that does not need to be considered delayed will be output to the log, and if it is too large, delayed SQL cannot be detected. Therefore, if the setting value cannot be estimated, set it small at the start of operation, and adjust it to a larger value as necessary while monitoring during operation.
.....

(3) Creating and reviewing a script for backing up statistics

To back up statistics, use the export function of the pg_dbms_stats extension.

The backup is used not only for restoration but also for identifying the cause of delayed SQL, so the time the backup was performed is important. Therefore, be sure to create a script like the one below so that the backup time is known.

We also provide a sample script that uses the export function to export on a database basis. This allows users to respond more flexibly by making modifications themselves, such as making changes on a schema basis, so copy the script to the execution current before using it.

The sample script is stored below. Also, "<x>" indicates the product version.

```
/opt/fsepv<x>server64/share/doc/extension/export_plain_stats-<x>.sql.sample
```

The following is an example of a file created when it is copied to the current execution directory as "export_plain_stats-<x>.sql".

```
expprt PATH=/opt/fsepv<x>server64/bin:$PATH
export CURRENTDIR=/pg_dbms_stats/backup

pushd "${CURRENTDIR}"

# make file name
EXECDATE=`date '+%Y%m%d%H%M%S'`
FILENAME=pg_dbms_stats_back.${EXECDATE}.dmp

# export
psql -d database1 -f export_plain_stats-<x>.sql

# rename dump file
```

```
mv export_stats.dmp ${FILENAME}
popd
```

(4) Periodic backup of statistics

Execute the script created in "[\(3\) Creating and reviewing a script for backing up statistics](#)" to back up the statistics information on a regular basis.

Backup execution interval

Perform backup at the following timings according to the statistics update method.

When statistics is automatically updated (autovacuum is enabled)

Statistics are updated when autovacuum is run.

Therefore, use the OS's cron command or schedule function of the task manager to periodically execute the export function and back up the statistics to a file.

Therefore, use the OS's cron command or a scheduler such as task manager to periodically execute the script created in "[\(3\) Creating and reviewing a script for backing up statistics](#)" and back up the statistics information to a file.

Obtain the shortest autovacuum execution interval from the log, and back up the statistics at an interval shorter than the obtained autovacuum execution interval. autovacuum logs can be obtained by setting the `log_autovacuum_min_duration` parameter to 0.

When users update statistics

If the user controls the updating of statistics using the ANALYZE command, etc., back up the statistics at the same time as updating the statistics.

Backup retention period

The backed up statistics will be used as investigation material to determine the cause of delayed SQL if a performance problem occurs. Therefore, keep it for the period of time expected to be required from the occurrence of the operational problem to its resolution.

Point

- The capacity required for backup mainly depends on the number of objects (schemas, tables, etc.) contained in the database. Estimate the required capacity from the backup file when checking the operation of the script and secure this space.
- To restore backed up statistics, the names of objects such as schemas and tables in the backup source must match. If you change the name of an object during operation, change the backup file name or storage location so that you can check whether it matches the object definition when restoring statistics.

(5) Monitoring delayed SQL and adjusting the allowable delay time

Monitors SQL that is delayed beyond the time set in the `log_min_duration_statement` parameter.

The settings in this section will be output in the following format.

```
[2024-04-30 10:20:11 JST]user1 postgres 3414[1]LOG: duration: 301001.541 ms statement: SELECT
pg_sleep(301)
[2024-04-30 10:26:12 JST]user1 postgres 3414[1]LOG: duration: 302002.321 ms statement: SELECT
pg_sleep(302)
```

When detected SQL is deemed to be delayed more than expected

Refer to "[\(6\) Investigating the cause of delayed SQL](#)" to determine the cause and take appropriate measures.

When SQL with execution time that does not affect operation is detected

The current setting just outputs redundant logs, so make major changes to the `log_min_duration_statement` parameter and then reload to enable it.

(6) Investigating the cause of delayed SQL

There are many factors that can cause SQL delays, so you must first identify the cause.

Here we will show you how to investigate SQL delays caused by changes in statistics.

Note that in the following cases, delays caused by statistics are unlikely, so recommend investigating from a different perspective (I/O, system load, etc.).

- When the SQL statement does not contain a search condition
- When the SQL statement uses the `pg_hint_plan` extension

1. Check the delayed SQL from the server log and identify the schema used by that SQL.
2. In a database cluster (such as a development environment) separate from the production environment, create an environment (hereafter referred to as the reproduction environment) with the same database name and table definition as the environment where the problem occurred. The reproduction environment also requires the `pg_dbms_stats` extension, so enable the `pg_dbms_stats` extension.
3. Use the `pg_dump` command from the production environment to obtain the target table definition. Since the data is not needed at this time, use the `-s` option of the `pg_dump` command from the production environment and use the extracted table definition.

The following is an example of output to the file `ddl_schema.dmp` when the schema to be investigated is `schema_1` in database `database1`.

```
pg_dump -d database1 -s schema_1 > ddl_schema.dmp
```

Using `ddl_schema.dmp`, create a research schema named `schema_1` in the database named `database1` in the reproduction environment.

```
psql -d database1 -f ddl_schema.dmp
```

4. In order to identify the cause of the SQL delay, use the backed up statistics to check whether there is a cause of the delay in the statistics. Import the statistics into the schema in the reproduction environment using the `pg_dbms_stats import` function (`dbms_stats.import_schema_stats`). The following is an example of restoring statistics for schema `schema_1` from the backup file "`pg_dbms_stats_back.20240422011223.dmp`".

```
psql -d database1 -c
"SELECT dbms_stats.import_schema_stats('schema_1', 'pg_dbms_stats_back.20240422011223.dmp')"
```

Then, execute the `EXPLAIN` command for the deferred SQL to obtain the execution result (query plan).

5. Restore the statistics by separately importing the statistics that were backed up during the previous time period when deferred SQL was running stably. For example, if the target SQL is executed daily as part of regular batch processing, use the data from the same time the previous day. For online processing, use the data from a previous time period when a similar business application is expected to be executed. The following is an example of restoring the statistics from the backup file "`pg_dbms_stats_back.20240421010001.dmp`" from one day ago, assuming the SQL is used in batch processing.

```
psql -d database1 -c
"SELECT dbms_stats.import_schema_stats('schema_1', 'pg_dbms_stats_back.20240421010001.dmp')"
```

It likewise obtains the result of executing the `EXPLAIN` command on deferred SQL (query plan).

6. Compare the execution results (query plans) of 4. and 5., and take the following action based on the results.
 - If the query plans are the same or the expected execution time of the query plan in 5. is longer
This may be a problem other than statistics (system load, index imbalance, I/O, etc.), so investigate from a different perspective.
 - Other than the above
The problem is likely that the optimal query plan was not selected due to updated statistics.
Refer to "[\(7\) Restoring statistics](#)" and import the statistics from the period when the system was operating stably into the production environment to restore the system.

(7) Restoring statistics

Restore the production environment statistics to the state they were in before the delay occurred. The procedure is the same as for verification, but it is as follows:

```
psql -d database1 -c
"SELECT dbms_stats.import_schema_stats('schema_1', 'pg_dbms_stats_back.20240422000000')"
```

9.10 Performance Tuning

9.10.1 Enhanced Query Plan Stability

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

If you know that the statistics will not change significantly, you can control the selection of the same query plan by directly specifying which query plan to select in `pg_hint_plan`, or by fixing the statistics referenced by the planner in `pg_dbms_stats`.

Also, `pg_dbms_stats` from Fujitsu Enterprise Postgres can fix statistics (height of Btree indexes) that OSS `pg_dbms_stats` cannot account for. Refer to "9.10.1.1 Fixing the Height of a Btree Index".



See

- For information about setting up `pg_hint_plan` and `pg_dbms_stats`, refer to "`pg_hint_plan`" and "`pg_dbms_stats`" in the Installation and Setup Guide for Server.
- For basic usage of `pg_hint_plan` and `pg_dbms_stats`, refer to below.
 - Control execution plans with `pg_hint_plan`
<https://www.postgresql.fastware.com/postgresql-insider-tun-hint-plan>
 - Control execution plans by fixing statistics with `pg_dbms_stats`
<https://www.postgresql.fastware.com/postgresql-insider-tun-dbms-stt>
- For more information about `pg_hint_plan` and `pg_dbms_stats`, refer to the OSS web page.

9.10.1.1 Fixing the Height of a Btree Index

When PostgreSQL creates a query plan, it also references the height of the Btree index to be used as part of the statistics. However, OSS `pg_dbms_stats` does not include the height of the Btree index in its fixed statistics. The height of a Btree index is also fixed in `pg_dbms_stats` for Fujitsu Enterprise Postgres.

Fixing the height of a Btree index covers Fujitsu Enterprise Postgres 16 and later, so to maintain compatibility with earlier versions, specify something like:

- PostgreSQL 9.2 and earlier compatibility

```
pg_dbms_stats.use_tree_height = off
```

- Fujitsu Enterprise Postgres 15 SP2 and earlier compatibility

```
pg_dbms_stats.lock_tree_height = off
```

See the table below for details on each parameter. If you want the parameter value to be valid for all sessions, specify it in `postgresql.conf` and then reload.

Parameter	Specified value	Default
pg_dbms_stats.use_tree_height	Specify whether to include the Btree index height in cost calculation and plan generation when statistics are fixed. on: Includes Btree index height in cost calculations and plan generation. off: Do not include Btree index height in cost calculation or plan generation.	on
pg_dbms_stats.lock_tree_height	Specify whether to also fix the height of the Btree index when you fix the statistics. Note that this is valid only when the specified value of pg_dbms_stats.use_tree_height is "on". on: Fixes the height of the Btree index at the height when locked. off: The height of the Btree index is not fixed and varies depending on the amount of data in the index.	on

When you want to import statistics from your legacy environment

The format of files handled by the import/export functions of this version of pg_dbms_stats is different from the format of pg_dbms_stats in OSS or Fujitsu Enterprise Postgres 15SP2 or earlier. Therefore, statistics exported from pg_dbms_stats for OSS or Fujitsu Enterprise Postgres 15SP2 or earlier should be imported using the dbms_stats.import_<obj>_stats_no_tree_height function.

How to specify the dbms_stats.import_<obj>_stats_no_tree_height function

```
dbms_stats.import_<obj>_stats_no_tree_height('Absolute path of export file')
```

<obj> indicates a database (currently connected), schema, table, or column.

Import using the dbms_stats.import_<obj>_stats function provided by OSS is not possible.

If you use the wrong import function, an error message containing the following HINT will be output. Check the combination of input file and import function type, and use the correct import function.

```
HINT: The import function may be incorrectly combined with the format of the exported data. Please check the documentation for the relationship between the import function and the available data.
```

Chapter 10 Building a Cluster System Using Patroni

Fujitsu Enterprise Postgres can build a cluster system in cooperation with Patroni and etcd.

For information on how to build and operate a cluster system, refer to the Patroni and etcd documentation.

For information on how to install Patroni and etcd, refer to the following.

- Patroni: "Patroni" in the Installation and Setup Guide for Server
- etcd: "etcd" in the Installation and Setup Guide for Server Assistant

System configuration

Regarding the etcd cluster configuration, a major choice is whether to place Patroni and etcd on different servers or not.

Placing etcd on the same server as Patroni can save on the number of servers, but for example, in the following cases, it is recommended to place Patroni and etcd on different servers to avoid trouble.

- Assuming operation to dynamically change the number of replica nodes

Placing etcd on all nodes might be excessive. However, placing etcd on only some nodes makes the system configuration a bit more complex.

- The number of replica nodes is large

This is to avoid resource contention between etcd and the database server.

- Share the same DCS (etcd) with other system.


Regarding the Patroni cluster configuration, various configurations are possible depending on the system requirements, so please decide on your own. However, please note that depending on the configuration method, a single server failure could cause Patroni to demote the primary server to a standby server and halt update operations.

How to configure PostgreSQL

To enable Patroni to manage the configuration of PostgreSQL's streaming replication, all PostgreSQL data directories and settings are described in `patroni.yml`. In particular, do not directly edit `postgresql.conf`.

Chapter 11 Autoscale of Replica Servers through Amazon EC2 Auto Scaling Integration

By using Amazon EC2 Auto Scaling (Auto Scaling), you can scale out/scale in a group of replica servers that are cascaded and replicated from standby servers in a multiplexed database operation, depending on the database load.



11.1 Autoscale Setup

Describes the autoscale setup.

11.1.1 Building a Cluster System

Describes how to build a cluster system that will be the replication source for autoscaling.

11.1.1.1 Preparation

Refer to the AWS documentation to install the AWS Command Line Interface (AWS CLI) and configure the necessary settings for the following operations.

Operations	Example IAM policies to assign
Take an EBS snapshot	"ec2:DescribeInstances" "ec2:CreateTags" "ec2:CreateSnapshot"
Store WAL archives in S3	"s3:PutObject" "s3:AbortMultipartUpload" "s3:ListBucket"
Update the launch template for snapshot restoration	ec2:CreateLaunchTemplateVersion ec2:DescribeLaunchTemplates ec2:DescribeLaunchTemplateVersions
Edit the Route53 Private Hosted Zone to enable updating replica server connections	"route53:ChangeResourceRecordSets" "route53:GetChange"
Enable storage expansion via SSM	AmazonSSMManagedInstanceCore
Enable metric collection via the CloudWatch Agent	CloudWatchAgentServerPolicy

Point

To use AWS commands in the script that takes snapshots, you need to set the path where AWS CLI is installed in the PATH environment variable. Specify the PATH environment variable either as a global environment variable or in the sample script.

11.1.1.2 Building a Cluster System Using a Mirroring Controller

Refer to the Cluster Operation Guide(Database Multiplexing) to build a streaming replication configuration on Amazon EC2 (EC2).

To quickly copy data during scale out, we will use the Amazon Elastic Block Store (EBS) restore feature. Therefore, mount EBS to the EC2 instance where Fujitsu Enterprise Postgres server features are installed, and specify the folder where EBS is mounted as the storage location for the database cluster with initdb. Similarly, create tablespaces by specifying a folder where EBS is mounted.

11.1.1.3 Acquiring Snapshots to Restore Data when Deploying Replica Servers

When scaling out, a snapshot of the database cluster's data area is taken to restore data when deploying replica servers.

Creating a script to take snapshots

Create a script to take snapshots of the EBS volume storing data, in order to restore data when scaling out. On the primary server, use OS Cron or similar tools to regularly execute the created script and take snapshots.

Sample scripts for taking snapshots are stored under the Fujitsu Enterprise Postgres installation directory below.

```
<Install directory>/share/get-storage-snapshot-for-autoscaling.sh.sample
```

The sample script takes a snapshot between the `pg_backup_start()` and `pg_backup_stop()` functions. EBS processes the data at the time the snapshot command is executed in the background, allowing read/write access to the volume immediately after the snapshot command is executed.

The sample script waits for the background snapshot process to complete, ensuring the snapshot finishes successfully. Since EBS snapshots capture incremental data since the last snapshot, the wait time varies depending on the amount of data changed since the previous snapshot. The first EBS snapshot captures all data stored on EBS. Therefore, it is recommended to take the snapshot after building the streaming replication cluster and storing the necessary data before starting database operations.

Additionally, when restoring data to a replica server from a snapshot, the WAL archives created after the snapshot was taken are applied. The time required to apply the WAL varies depending on the amount of database changes since the snapshot was taken. Therefore, determine the snapshot interval based on the startup time settings for the replica server.

1. Edit the EBS volume ID to take a snapshot.

Specify the ID, device name to be specified in the launch template and mount path of the EBS volume to be snapshotted.

Specify all EBS volumes that store data and tablespaces.

```
# Specify the mount path of the EBS storing the data.
declare -A volumes
volumes=(
  [ "vol-xxxxxxxxxxxxxxxx" ]="/dev/sdf:/data"
  [ "vol-yyyyyyyyyyyyyy" ]="/dev/sdg:/tablespace"
)
```

Information

You can check the ID of an EBS volume mounted to a specific instance with a command like the following.

```
$ aws ec2 describe-volumes --filters Name=attachment.instance-id,Values=<your-instance-id> --
query "Volumes[*].{ID:VolumeId}" --output text
```

2. Editing PostgreSQL database connection information

PostgreSQL connection information is specified to execute `pg_backup_start()` and `pg_backup_stop()` functions before and after creating a snapshot. If necessary, add and define environment variables related to authentication, such as passwords and SSL communication certificates.

```
# PostgreSQL database connection information
export PGHOST=localhost
export PGPORT=27500
export PGDATABASE=postgres
export PGUSER=fsepuser
```

3. Edit installation path

Specify the installation path for Fujitsu Enterprise Postgres and AWS CLI. "<x>" indicates the product version.

```
export FEP_INSTALL_DIR=/opt/fsepv<x>server64
export AWSCLI_BIN_DIR=/usr/local/bin
```

4. When updating the definition of the launch template for the snapshot restore destination, specify the following variables.

Replace the snapshot ID to restore to the EBS volume defined in the specified launch template with the ID of the most recent snapshot. Utilizing the most recent snapshot reduces the time required to apply the WAL to the data restored to replica servers deployed in a scale out configuration.

```
# Define if you want to update the snapshot restore source for the volume of the launch template.
export LAUNCH_TEMPLATE_NAME=your-launch-template-name
export LAUNTH_TEMPLATE_SOURCE_VERSION=1
```

Editing postgresql.conf

When restoring EBS from a snapshot during scale out, we restore to the latest data using archived WALs stored in Amazon Simple Storage Service (S3). Therefore, we enable `archive_mode` in `postgresql.conf` and archive WALs to S3.

The following is an example of archiving WALs to the S3 bucket "primary-wal-bucket".

```
archive_mode = on
archive_command = 'aws s3 cp %p s3://primary-wal-bucket/%f'
```

11.1.1.4 Settings for Name Resolution of the Standby Server, which will be the Replica Source Server

Use Amazon Route 53 private hosted zones to resolve names from replica servers to standby servers.

Private hosted zones allow you to configure DNS records for resources within your Virtual Private Cloud (VPC). You can set an arbitrary domain name (e.g., `standby-server.com`) and associate it with the private IP address of the standby server's EC2 instance.

While EC2 instances are assigned a private IP DNS name upon creation, by simply rewriting the Route 53 DNS record, it becomes easier to switch the replication connection destination for multiple replica servers when replacing standby servers or switching between primary and standby.

11.1.2 Building a Replica Server

Describes the procedure for creating an Amazon Machine Image (AMI) that will be the source for replica server instances.

Set up the necessary startup scripts for AMI deployment on the instance where Fujitsu Enterprise Postgres server is installed.

11.1.2.1 Installation of Fujitsu Enterprise Postgres on EC2

Install the Fujitsu Enterprise Postgres server on any EC2 instance and create an OS user who will be the instance administrator.

Database initialization and volume attachment are not required at this time. These will be performed when deploying the replica server.

11.1.2.2 Creation of Startup Script for Replica Server Deployment

Create a launch script that runs when replica servers are deployed by Auto Scaling. Specify the created script to run in the launch template.

During script execution, heartbeats are continuously sent to extend the Auto Scaling group's lifecycle hooks. During this period, connections from the load balancer are disabled. Once data restoration completes and replication catches up on the standby server, the lifecycle hooks conclude, and connections from the load balancer become possible again.

The script to be created will perform the following processes:

- Modify postgresql.conf (define replication settings, restore_command, etc.)
- Start the database
- Wait until replication catch-up is complete, then transition to a state where connections from the load balancer are accepted

Preparation

1. Refer to the AWS documentation to install the AWS Command Line Interface (AWS CLI) and configure it as needed for the following steps.

Operations	Example IAM policies to assign
Manipulate lifecycle hooks until database startup	"autoscaling:RecordLifecycleActionHeartbeat" "autoscaling:DescribeAutoScalingInstances" "autoscaling:CompleteLifecycleAction"
Retrieve WAL archives stored in S3	"s3:PutObject" "s3:GetObject" "s3:ListBucket"
Apply tags to mounted EBS volumes	"ec2:DescribeInstances" "ec2:DescribeVolumes" "ec2:DescribeVolumeStatus" "ec2:CreateTags"
Collect metrics via the CloudWatch Agent	CloudWatchAgentServerPolicy
Perform remote operations via SSM during storage expansion and scale in	AmazonSSMManagedInstanceCore

2. Refer to the WS documentation to install the SSM Agent.

The SSM Agent is required to stop the database before instances are deleted during scale in.

Creating a startup script

Sample script for database startup is stored under the Fujitsu Enterprise Postgres installation directory below.

```
<Install directory>/share/start-replica-database-for-autoscaling.sh.sample
```

1. Edit the EBS volume ID for taking snapshots.

Specify the ID and mount path of the EBS volume from which the snapshot was taken.

Specify all EBS volumes that store data and tablespaces. Specify the same values as the script that creates snapshots placed on the primary server.

```
# Specify the mount path of the EBS storing the data.
declare -A volumes
volumes=(
  ["vol-xxxxxxxxxxxxxxxx"]="/dev/sdf:/data"
  ["vol-yyyyyyyyyyyyyy"]="/dev/sdg:/tablespace"
)
```

Information

You can check the ID of an EBS volume mounted to a specific instance with a command like the following.

```
$ aws ec2 describe-volumes --filters Name=attachment.instance-id,Values=<your-instance-id> --query "Volumes[*].{ID:VolumeId}" --output tex
```

2. Editing the OS user for the instance administrator

Specify the OS user for the instance administrator. The database will be started using the specified user.

```
export DB_OWNER=fsepuser
```

3. Editing the hostname of the standby server for replication source

Specify the hostname of the standby server that will be the replication source and connection string for connection (primary_conninfo value in postgresql.conf).

For an example of how to set the hostname of the standby server, refer to ["11.1.1.4 Settings for Name Resolution of the Standby Server, which will be the Replica Source Server"](#).

```
export REPLICATION_HOST=standby-server.com
export PRIMARY_CONNINFO="host=${REPLICATION_SOURCE_DNS}"
```

4. Editing PostgreSQL database connection information

Specify PostgreSQL connection information to confirm database startup. If necessary, add and define environment variables related to authentication, such as passwords and SSL communication certificates. Also, define the PGDATA variable for database startup.

```
# PostgreSQL database connection information
export PGDATA=/data
export PGHOST=localhost
export PGPORT=27500
export PGDATABASE=postgres
export PGUSER=${DB_OWNER}
```

5. Editing the S3 bucket storing WAL archives

Specify the name of the S3 bucket that stores WAL archives used for data recovery when the database starts up and the value of restore_command in postgresql.conf.

```
# S3 bucket name storing WAL archives used for data recovery
export ARCHIVE_S3_BUCKET =primary-wal-bucket
export RESTORE_COMMAND="${AWS_CLI_BIN}/aws s3 cp s3://${ARCHIVE_S3_BUCKET}/%f %p"
```

6. Editing installation path

Specify the installation path for Fujitsu Enterprise Postgres and AWS CLI. "<x>" indicates the product version.

```
export FEP_INSTALL_DIR=/opt/fsepv<x>server64
export AWSCLI_BIN_DIR=/usr/local/bin
```

7. Editing postgresql.conf

The following values are inserted into postgresql.conf when the script runs. If different values are specified for the same parameters in postgresql.auto.conf, or if you wish to add parameters, modify the values below.

```
# Recreate postgresql.conf for the replica
postgresql_conf="${PGDATA}/postgresql.conf"
cat << EOF >> "${postgresql_conf}"
# Replication Settings
hot_standby = on
primary_conninfo = '${PRIMARY_CONNINFO}'
```

```
restore_command = '${RESTORE_COMMAND}'
EOF
```

11.1.2.3 Volume Mount Definition

When EC2 instances of replica servers are deployed by scaling out, `/etc/fstab` is edited to mount the attached EBS volume as a data storage area.

1. Obtaining UUID

On the primary server where the snapshot is to be taken, obtain the UUID to be defined in `/etc/fstab` using the `lsblk` or `blkid` command.

```
# lsblk -o NAME,UUID,MOUNTPOINTS
NAME                UUID                                MOUNTPOINTS
nvme1n1             11111111-1111-11aa-aaaa-aaaaaaaaaaaa /database/data
nvme2n1             22222222-2222-22bb-bbbb-bbbbbbbbbbbb /database/tblspc1
```

2. Editing `/etc/fstab`

On the EC2 instance where you are setting up the replica server, define the acquired UUID in `/etc/fstab`.

```
# vi /etc/fstab
UUID=11111111-1111-11aa-aaaa-aaaaaaaaaaaa /database/data    xfs  defaults,nofail 0 2
UUID=22222222-2222-22bb-bbbb-bbbbbbbbbbbb /database/tblspc1 xfs  defaults,nofail 0 2
```

3. Creating a folder to mount the volume

Create a folder to mount the volume and change the owner of the folder to the OS user of the database instance administrator.

```
# mkdir -p /database/data
# mkdir -p /database/tblspc1
# chown fsepuser:fsepuser /database/data /database/tblspc1
```

11.1.2.4 Creation of AMI for Replica Server

Create an AMI for the replica server you set up.

In the AWS Management Console, select the EC2 instance where you set up the replica server from the EC2 instance list and create an AMI.

The AMI created here will be used when defining scale out with Auto Scaling.

11.1.3 Enable Autoscale

Create an Auto Scaling launch template and an Auto Scaling group when deploying an AMI. When creating an Auto Scaling group, integrate it with Elastic Load Balancing to distribute queries.

Creating an Auto Scaling launch template

1. From the "My AMIs" tab displayed in the EC2 console, select the AMI created with "[11.1.2.4 Creation of AMI for Replica Server](#)".
2. Under Storage and Volume, specify the EBS volume to mount to EC2. Specify a snapshot to restore data from when launching the instance.
 - Size: Specify a size equal to or larger than the primary server being replicated.
 - Device Name: Specify the device name of the volumes defined in the script that acquires the snapshot.
 - Snapshot: Specify the ID of an existing snapshot.
3. Enable the "Allow tags in metadata" option to read EC2 instance tags in the startup script and obtain the Auto Scaling Group name.
4. Define the script to be executed when the instance starts in "User data". Here, define the execution of the database startup script created in "[11.1.2.2 Creation of Startup Script for Replica Server Deployment](#)".

The following is an example definition when the database startup script is defined under `/home/fsepuser`.

```
#!/bin/bash

/home/fsepuser/start-replica-database-for-autoscaling.sh
```

Creating an Auto Scaling group

1. Select a "launch template" that defines the AMI and launch script for the replica instances.
2. Select the instance launch options.
 - Select "Override launch template" for the instance type and define a mixed instances policy. This allows you to update instances by replacing only the root volume when patching replica servers, omitting the copying of data storage areas.
 - Specify the VPC and Availability Zones where the scaled-out instances will be placed as instance launch options. It is recommended to specify multiple Availability Zones.
3. Create a load balancer. Create and attach a Network Load Balancer to load balance TCP traffic.
4. Configure group size and scaling policies.
5. Add notification settings for instance launches.
6. Add tags to the group.

11.1.4 Preparation for Scale in

Describes how to define the processing to execute before instances are deleted during scale in.

The example implements processing to stop the database before instances are deleted. You can implement processing to execute before the targeted instances are deleted during scale in, such as notifying the application.

1. Create a Lambda function to stop the database on the console

Define a Lambda function to be called by the lifecycle hook when an instance stops.

The following is an example of the function code that stops the database of the instance to be stopped.

```
import boto3

ssm_client = boto3.client('ssm')
autoscaling = boto3.client('autoscaling')

def lambda_handler(event, context):

    # Specify the data directory as a variable
    data_directory = '/database/inst1'

    # Get EC2 instance ID
    instance_id = event['EC2InstanceId']
    # Get AutoScaling information
    hook_name = event['LifecycleHookName']
    asg_name = event['AutoScalingGroupName']
    token = event['LifecycleActionToken']

    # Execute SSM command
    ssm_client.send_command(
        InstanceIds=[instance_id],
        DocumentName='AWS-RunShellScript',
        Parameters={
            'commands': [
                f'pg_ctl stop -D {data_directory}'
            ]
        }
    )

    # Restart the termination process
```

```

autoscaling.complete_lifecycle_action(
    LifecycleHookName=hook_name,
    AutoScalingGroupName=asg_name,
    LifecycleActionToken=token,
    LifecycleActionResult='CONTINUE'
)

return {
    'statusCode': 200,
    'body': 'Command sent successfully'
}

```

2. Define Auto Scaling lifecycle hooks

Define a lifecycle hook that invokes a Lambda function when Auto Scaling performs a scale in.

```

$ aws autoscaling put-lifecycle-hook \
  --lifecycle-hook-name <Life cycle hook name> \
  --auto-scaling-group-name <auto scaling group name> \
  --lifecycle-transition autoscaling:EC2_INSTANCE_TERMINATING \
  --notification-target-arn arn:aws:lambda:ap-northeast-1:<id>:function:<lambda function name> \
  --default-result CONTINUE

```

3. Set the delay in load balancer deregistration time

Set the load balancer deregistration delay as needed. The following command specifies the time until the load balancer stops connection traffic for EC2 instances targeted for scale in.

```

$ aws elbv2 modify-target-group-attributes --target-group-arn <Target group ARN> --attributes
Key=deregistration_delay.timeout_seconds,Value=300

```

11.2 Operations for Autoscale

Describes the operational procedures after auto scale out is configured.

11.2.1 Customizing Scaling Policies

Auto Scaling's scaling policy, by default, determines scale out based on CPU utilization.

Describes the steps to auto-scale out using metrics other than CPU utilization (memory utilization, number of database connections).

1. Deploy EC2 using the replica server AMI.
2. Install the CloudWatch agent on the deployed EC2.

```

$ curl "https://amazoncloudwatch-agent.s3.amazonaws.com/redhat/amd64/latest/amazon-cloudwatch-agent.rpm" -o "amazon-cloudwatch-agent.rpm"
$ sudo rpm -Uvh amazon-cloudwatch-agent.rpm
$ rm amazon-cloudwatch-agent.rpm

```

3. Enable CloudWatch metrics collection

Below is an example of how to get memory usage and the number of database connections.

Example) When collecting memory usage

Use the CloudWatch Agent to collect metrics. Below is an example of editing the Agent configuration file to enable collection of the "mem_used_percent" metric.

```

# Create a Json definition to collect memory usage
$ cat << 'EOF' > /opt/aws/amazon-cloudwatch-agent/etc/amazon-cloudwatch-agent.json
{
  "metrics": {
    "namespace": "Custom/FEP",
    "append_dimensions": {

```

```

    "AutoScalingGroupName": "${aws:AutoScalingGroupName}",
    "InstanceId": "${aws:InstanceId}"
  },
  "metrics_collected": {
    "mem": {
      "measurement": [
        "mem_used_percent"
      ],
      "metrics_collection_interval": 60
    },
  }
}
EOF

# Enable the Cloud Watch agent
$ systemctl enable amazon-cloudwatch-agent

```

Example) When getting the number of database connections

Use the aws command to send database connection counts to CloudWatch. Below is an example of defining database connection counts using the psql command and sending them to CloudWatch.

```

# Create a script to send the number of connections to CloudWatch
$ cat << 'EOF' > /home/fsepuser/send_fep_conn.sh
#!/bin/bash

TOKEN=$(curl -s -X PUT "http://169.254.169.254/latest/api/token" -H "X-aws-ec2-metadata-token-ttl-seconds: 21600")
INSTANCE_ID=$(curl -sH "X-aws-ec2-metadata-token: $TOKEN" http://169.254.169.254/latest/meta-data/instance-id)
ASG_NAME=$(curl -sH "X-aws-ec2-metadata-token: $TOKEN" http://169.254.169.254/latest/meta-data/tags/instance/aws:autoscaling:groupName)

CONN=$(sudo -u fsepuser psql -U fsepuser -d postgres -t -c "select count(*) from pg_stat_activity;" | xargs)
aws cloudwatch put-metric-data \
  --namespace Custom/FEP \
  --metric-name DBConnections \
  --dimensions AutoScalingGroupName="$ASG_NAME",InstanceId="$INSTANCE_ID" \
  --value "$CONN"
EOF

# Register the script with Cron for regular execution
cat <<EOF > /etc/cron.d/send_fep_conn
* * * * * root /bin/bash /home/fsepuser/send_fep_conn.sh
EOF
chmod 644 /etc/cron.d/send_fep_conn

# Enable Cron
$ systemctl enable crond

```

4. Create an AMI on the EC2 instance you set up.

5. Set up autoscale using the AMI you created.

If you are creating a new one, refer to "[11.1.3 Enable Autoscale](#)".

If you are updating an existing AMI that is auto scaling, refer to "[11.2.4 Updating an Existing AMI](#)".

6. Create a CloudWatch alert in the AWS console.

Set the following to define when to scale out by issuing an alarm:

- Metrics to monitor

- Threshold for issuing an alert
 - Duration for which the threshold is exceeded
7. Modify the Auto Scaling group's scaling policy in the AWS console.
Update the scaling policy by specifying the alert rule you created.

11.2.2 Scale in of Replica Servers

Describes the procedure for scale in a scaled out instance.

Preparation as shown in "11.1.4 Preparation for Scale in" is required in advance.

Execute scale in

Execute scale in of the replica server with the following command.

```
$ aws autoscaling update-auto-scaling-group --auto-scaling-group-name <AutoScalingGroupName> --
desired-capacity <NewDesiredCapacity>
```

11.2.3 Switching/Failback of the Cascade Replication Source for the Replica Server

Describes how to switch the replication source to continue operating the replica server, for example, when the standby server of the replication source stops abnormally.

Creation of a sample script for promoting to primary server

Prepare 2 scripts for switching the replication target of the replica server.

- Switching script to execute when the standby server is disconnected and the primary server becomes the replication source for the replica server
- Script to be executed when the standby server is restored and reintegrated, to reconfigure it as the replication source.

1. Creating the switchover script

Sample scripts are stored in the Fujitsu Enterprise Postgres installation directory. Switching the replication destination for the replica server is achieved by modifying the Route 53 DNS server records for each script.

- A script that is called and executed before the standby server is disconnected

```
<Install directory>/share/switch-wal-source-to-primary-for-autoscaling.sh.sample
```

- A script that is called and executed after the standby server is incorporated

```
<Install directory>/share/return-wal-source-back-to-standby-for-autoscaling.sh.sample
```

2. Edit the connection information

Edit the following connection information defined in the sample script.

```
# Specify the ID of the Route53 Private Hosted Zone used for standby server name resolution
HOSTED_ZONE_ID="YOUREHOSTEDZONEID"
# Specifies how long the DNS resolver keeps information about records
# A shorter value allows faster record updates but increases query volume and cost. A longer
value has the opposite effect
# Set this value based on the time required for replica servers to switch their connection targets
TTL=10
# Specifies the DNS server record for the standby server used as the connection target for
replica servers
WAL_SOURCE_RECORD="wal-source.sample.internal"
```

3. Specify for Mirroring Controller state transition command

Specify the following in the server definition files for the primary and standby servers.

```
$ vi workdir/MC/server1.conf
db_instance = '/data'
pre_detach_command = '/home/fsepuser/workdir/switch-wal-source-to-primary-for-autoscaling.sh'
post_attach_command = '/home/fsepuser/workdir/return-wal-source-back-to-standby-for-autoscaling.sh'
```

Operations during failback

After the standby server is restored, the hostname managed by ROUTE 53 is switched back by the script executed during failback. At this point, the connections already established between the primary server and replica server remain. To switch back the replication connection destination, restart the replica servers sequentially at any time.

The sample script for restarting replica servers managed by AutoScaling is stored below.

```
<Install directory>/share/rolling_restart_fep_replica-for-autoscaling.sh.sample
```

Edit the following in the sample script and run the script. "<x>" indicates the product version.

```
REGION="your-region"

# AutoScaling name managing replica servers
ASG_NAME="your-read-replica-asg"

# Specify the installation path of Fujitsu Enterprise Postgres and the path to the AWS CLI command
FEP_INSTALL_DIR="/opt/fsepv<x>server64"
AWS_CLI_BIN="/usr/local/bin"

# Specify the path to the data storage area
PGDATA="/data"

# Specify the OS user for the database instance administrator
DB_OWNER='fsepuser'
```

11.2.4 Updating an Existing AMI

Describes the procedure for applying Fujitsu Enterprise Postgres and OS patches to replica servers.

After changing the AMI for the replica server, edit the launch template and Auto Scaling group respectively to apply the new AMI. When performing instance updates in the Auto Scaling group, it is recommended to specify the "Replace root volume" strategy for the instance update strategy. This prevents EBS replacement from occurring and eliminates the need for snapshot restoration time.

1. Edit the replica server AMI.

Deploy an EC2 instance using the AMI of the replica server currently in use.

Edit the script on the deployed EC2 instance, and once modifications are complete, create a new AMI.

2. Edit the launch template.

Create a new template or replace the AMI specified in an existing template with the new AMI. In the latter case, a new version of the launch template will be added.

3. Edit the Auto Scaling group.

Specify the new template or new version of the launch template.

There are two types of instance update strategies: "rolling strategy" and "root volume replacement strategy". In the former case, EBS replacement occurs and data restoration begins, so it is recommended to use the latter strategy.

11.2.5 Storage Scale up

Describes how to monitor storage utilization and scale up storage capacity when utilization exceeds a threshold.

1. Install the CloudWatch agent on each EC2 instance.
2. Create and execute a CloudWatch metric collection script on each EC2 instance.

Example)

```
# Create a Json definition to collect disk usage in addition to memory usage
$ cat << 'EOF' > /opt/aws/amazon-cloudwatch-agent/etc/amazon-cloudwatch-agent.json
{
  "metrics": {
    "namespace": "Custom/FEP",
    "append_dimensions": {
      "AutoScalingGroupName": "${aws:AutoScalingGroupName}",
      "InstanceId": "${aws:InstanceId}"
    },
    "metrics_collected": {
      "mem": {
        "measurement": [
          "mem_used_percent"
        ],
        "metrics_collection_interval": 60
      },
      "disk": {
        "measurement": [
          "used_percent"
        ],
        "resources": [
          "/database/inst1",
          "/database/tblspc1"
        ]
      }
    }
  }
}
EOF

# Enable the Cloud Watch agent
$ systemctl enable amazon-cloudwatch-agent
```

Creating a CloudWatch alert in the AWS console

Describes the following steps to automatically scale up storage when its utilization exceeds a threshold:

- Expanding the file system on the EC2 instance where EBS is mounted using SystemManager
- Invoke a Lambda function via a Simple Notification Service (SNS) topic
- Execute storage scaling using the Lambda function (create Python script, grant privileges via IAM role)
- Create a CloudWatch alert rule

The sample script here assumes the following:

- The EBS file system uses xfs or ext4
- When targeting EBS from a streaming replication cluster built with Mirroring Controller for expansion, apply the following tags. These tags identify the target EBS during storage scale up execution.
 - Target EBS:
 - Key=MountPoint, Value=<EBS mount path>
 - Primary server and Standby server EC2:
 - Key=DBRole, Value=<any value>

- Key=ASG_NAME, Value=<Auto Scaling group name that manages the replica server at the replication destination>

1. Defining EC2 file system expansion via System Manager

In the AWS Management Console, create a System Manager Command Document. The Command Document receives the mount path of the expanded EBS volume from the Lambda function, identifies the file system, and expands it.

The sample script for the SSM Command Document is stored below. Paste the definition example content into the Command Document editor.

```
<Install directory>/share/expand-ebs-volume-for-autoscaling.yaml.sample
```

2. Defining notifications via an SNS topic

In the AWS Management Console, create an SNS topic. Select Standard as the type.

3. Defining scale up using Lambda functions

Create an IAM role to attach to the Lambda function.

Operations	Example IAM policies to assign
Expand the EBS volume	"ec2:DescribeInstances" "ec2:DescribeVolumes" "ec2:DescribeVolumesModifications" "ec2:ModifyVolume"
Execute the command in SystemManager to expand the file system	"ssm:SendCommand"

1. Create a Lambda function with the following specifications:

- Function creation option: Select "Create from scratch"
- Runtime: Python 3.9 or later
- Execution role: Select "Use another role" and specify the IAM role created for the Lambda function

2. Configure the following in the created Lambda function:

- Define the Python script to expand the EBS volume (using the sample script)
- Define the environment variables used by the script
- Define the trigger to execute the Lambda function

3. Define the Python script to perform storage expansion in the "Code Source" editor. Paste the contents of the sample script stored below into the editor.

```
<Install directory>/share/sns-triggered-lambda-template-for-autoscaling.py.sample
```

4. Edit the "Environment Variables" under "Settings" and define the following environment variables.

Key	Value
ASG_NAME	Specify the autoscaling group name.
INCREASE_GIB	Specify the storage capacity increase amount numerically. Specify either this parameter or TARGET_SIZE_GIB.
TARGET_SIZE_GIB	Specify the storage capacity as a numeric value. Specify either this parameter or INCREASE_GIB.
SSM_DOCUMENT_NAME	Specify the name of the created SSM Command Document.
MOUNT_PATH	Specify the mount path of the EBS to be expanded.

Key	Value
DB_ROLE	Specify the tag names defined for the primary server and replica server.

5. Specify the trigger to execute the Lambda function from "Add trigger".

- Select source: Choose "SNS"
- SNS topic: Select the created SNS topic

4. Create an alert rule

In the AWS Management Console, define thresholds and create an alert rule for the storage you want to monitor.

Select "Create alarm" in CloudWatch.

- Specify metrics and conditions
 - Metric: Define the expression for the metric you want to monitor
 - Threshold: Specify the threshold at which alerts will be notified for the selected metric
- Configuring Actions

In the Notifications section, select the following:

- Alarm state trigger: "Alarm state"
- Send notifications to the following SNS topic: "Select an existing SNS topic"
- Notification destination: Select the SNS topic specified as the trigger for the Lambda function

When targeting multiple EBS volumes (e.g., tablespaces), create separate SNS topics, Lambda functions, and CloudWatch alert rules for each EBS mount path.

11.2.6 Backup/Restore

Describes how to acquire backups during auto scale out execution.

Auto scale out restores data to replica servers using EBS snapshots and WAL archives. If WAL archives after the snapshot acquisition point are insufficient, the replica server cannot catch up to the latest state of the replication source, causing the database startup to fail.

When using pgBackRest, if you acquire full backups exceeding the retention generation count, older WAL is deleted. Therefore, if the retention generation count for full backups is set to 1, WAL archives prior to the full backup are deleted. This can result in a point where WAL after the snapshot acquisition time is insufficient. Consequently, to prevent WAL archives after the snapshot used for restoration during auto scale out from being deleted, you must carefully consider the timing of snapshot and full backup acquisition.

For example, if you specify 'count' for the retention-full-type option in pgBackRest, set retention-full to a value of 2 or higher. This retains multiple generations of full backups and the WAL archives corresponding to the retention period of the full backups.

By taking an EBS snapshot immediately after acquiring a full backup, the WAL archives from the time of the previous full backup will still be retained even after the next full backup is acquired. This allows the use of WAL archives created after the latest snapshot.

This approach prevents the deletion of WAL archives created after the latest snapshot, avoiding situations where insufficient WAL during snapshot restoration prevents data recovery.

Chapter 12 Streaming Replication Using WebAdmin

This chapter describes streaming replication using WebAdmin.

A standby instance for streaming replication can be created from a standalone instance, a master instance, or another standby instance. The standby instance connects to the master instance and uses WAL records to replicate data. The standby instance can be used for read-only operations.

WebAdmin cannot back up standby instances.

Additionally, streaming replication can be set to asynchronous or synchronous mode.


Point

- 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].

For characters that cannot be specified in WebAdmin, refer to “[Appendix K WebAdmin Disallow User Inputs Containing Hazardous Characters](#)”.

12.1 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.
 - [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 J WebAdmin Wallet](#)" for information on creating wallet entries.
 - [Instance name]: Name of the standby database instance to create.


The name must meet the conditions below:

 - Maximum of 16 characters
 - The first character must be an ASCII alphabetic character
 - The other characters must be ASCII alphanumeric characters
 - [Instance port]: Port number of the standby database instance.
 - [Host IP address]: The IP address of the server machine where the standby instance is to be created. This information is needed to configure the standby instance to be connected to the master.
 - [Data storage path]: Directory where the database data will be stored
 - [Backup storage path]: Directory where the database backup will be stored
 - [Transaction log path]: Directory where the transaction log will be stored
 - [Encoding]: Database encoding system
 - [Replication mode]: Replication mode of the standby instance to be created ("Asynchronous" or "Synchronous")

- [Application name]: The reference name of the standby instance used to identify it to the master instance.

The name must meet the conditions below:

- Maximum of 16 characters
- The first character must be an ASCII alphabetic character
- The other characters must be ASCII alphanumeric characters

4. Click  to create the standby instance.


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
ERROR: pgx_rcvall: backup of the database has not yet been performed, or an incorrect backup storage
directory was specified
```


12.2 Promoting a Standby Instance

Streaming replication between a master and standby instance can be discontinued using WebAdmin.

Follow the procedure below to promote a standby instance to a standalone instance, thereby discontinuing the streaming replication.

1. In the [Instances] tab, select the standby instance that needs to be promoted.
2. Click .
3. Click [Yes] from the confirmation dialog box.

The standby instance will be promoted and will become a standalone instance, which is not part of a streaming replication cluster.

Once the standby instance is promoted to become a standalone instance, the backup storage status will be "Error". This is because no backups are available when the instance is newly promoted to a standalone instance. The status will be reset if a new backup is performed by clicking [Solution] or .




12.3 Converting an Asynchronous Replication to Synchronous

You can convert an asynchronous standby instance to a synchronous standby instance.

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.

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.



See

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.

12.4 Converting a Synchronous Replication to Asynchronous

Follow the procedure below to convert a Synchronous standby instance to Asynchronous.

1. In the [Instances] tab, select the master instance of the relevant cluster.
2. Click
3. In the [Streaming replication] section, edit the value for [Synchronous standby names].
 - Remove the "Application name" of the standby instance you want to be in Asynchronous mode.
4. Click
5. Select the master instance and click
6. Select the standby instance. [Instance type] will now show the updated status.



See

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.

12.5 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 J WebAdmin Wallet](#)" for information on creating wallet entries. Replication credential (user name and password) should not contain hazardous characters. Refer to "[Appendix K WebAdmin Disallow User Inputs Containing Hazardous Characters](#)".
6. Enter [Host IP address].

Specify the IP address of the node where the standby instance was created.
7. Click to open the [Join replication cluster] dialog box.

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

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

Chapter 13 Installing and Operating the In-memory Feature

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

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

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

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

This feature can only be used in Advanced Edition.

13.1 Installing Vertical Clustered Index (VCI)

This section describes the installation of VCI.

1. [Evaluating whether to Install VCI](#)
2. [Estimating Resources](#)
3. [Setting up](#)

13.1.1 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

Point

.....
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 "13.1.4 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.
-

13.1.2 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.

13.1.3 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](#)

13.1.3.1 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

Parameter name	Setting value	Description	Required
vci.max_parallel_degree	Maximum number of background workers used for parallel scan	Maximum number of background workers used for parallel scan. Add this value to max_worker_processes.	N

Example

```
shared_preload_libraries = 'vci, pg_prewarm'
session_preload_libraries = 'vci, pg_prewarm'
reserve_buffer_ratio = 20
vci.control_max_workers = 8
vci.max_parallel_degree = 4
max_worker_processes = 18 # Example: If the initial value was 6, 6 + 8 + 4 = 18
```

Information

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

See

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

13.1.3.2 Installing the Extensions

Execute CREATE EXTENSION to install the VCI and pg_prewarm extensions. Both extensions need to be installed for each database. Also, extensions can only be installed on the public schema by a superuser.

- Installing VCI

```
db01=# CREATE EXTENSION vci;
```

- Installing pg_prewarm

```
db01=# CREATE EXTENSION pg_prewarm;
```

Some operations cannot be performed for VCI extensions. Refer to "[13.2.1 Commands that cannot be Used for VCI](#)" for details.

13.1.3.3 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);
```

Point

- Some table types cannot be specified on the ON clause of CREATE INDEX. Refer to "[13.1.4.1 Relation Types](#)" for details.
- Some data types cannot be specified on the column specification of CREATE INDEX. Refer to "[13.1.4.2 Data Types](#)" for details.
- Some operations cannot be performed for VCI. Refer to "[13.2.1 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.

13.1.3.4 Confirming that the VCI has been Created

Execute the SELECT statement to reference the `pg_indexes` catalog, and confirm that the VCI was created for the target columns.

Example

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

13.1.4 Data that can Use VCI

This section describes on which relation types and for which data types VCIs can be created.

13.1.4.1 Relation Types

VCIs cannot be created on some relation types.

The ON clause of CREATE INDEX described in "13.1.3.3 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

13.1.4.2 Data Types

VCI's cannot be created for some data types.

The column specification of CREATE INDEX described in "13.1.3.3 Creating a VCI" cannot specify a column with data type on which VCI's 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	N
	tsquery	N
UUID	uuid	Y
XML	xml	N
JSON	json	N
	jsonb	N
Range	int4range	N
	int8range	N
	numrange	N
	tsrange	N
	tstzrange	N
	daterange	N
Object identifier	oid	N
	regproc	N
	regprocedure	N
	regoper	N
	regoperator	N
	regclass	N
	regtype	N
	regconfig	N
	regdictionary	N
pg_lsn type	pg_lsn	N
Array type	-	N
User-defined type (Basic type, enumerated type, composite type, and range type)	-	N

13.2 Operating VCI

This section describes how to operate VCI.

13.2.1 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, except when the public schema is specified. This operation is not required for VCI.

- Operations that cannot be performed on relations containing a VCI

Command	Subcommand	Description
ALTER INDEX	SET	The subcommands on the left cannot be performed if a VCI is specified.
	SET TABLESPACE	
	ALL IN TABLESPACE	If the operation is required, delete the VCI using DROP INDEX, and re-create it using CREATE INDEX after completing the operation.
ALTER OPERATOR CLASS	RENAME TO	The subcommands on the left cannot be performed if a VCI is specified.
	OWNER TO	
	SET SCHEMA	This operation is not supported in VCI.
ALTER OPERATOR FAMILY	ADD	
	DROP	
	RENAME TO	
	OWNER TO	
	SET SCHEMA	
ALTER TABLE	ALL IN TABLESPACE <i>name</i> [OWNED BY <i>roleName</i>] SET TABLESPACE <i>newTablespace</i>	A tablespace that contains a VCI cannot be specified. If the operation is required, delete the VCI using DROP INDEX, and re-create it using CREATE INDEX after completing the operation.
	DROP [COLUMN] [IF EXISTS] <i>colName</i> [RESTRICT CASCADE]	A column that contains a VCI cannot be specified. If the operation is required, delete the VCI using DROP INDEX, and re-create it using CREATE INDEX after completing the operation.
	ALTER [COLUMN] <i>colName</i> [SET DATA] TYPE <i>dataType</i> [COLLATE <i>collation</i>] [USING <i>expr</i>]	
	CLUSTER ON <i>indexName</i>	A VCI cannot be specified.
	REPLICA IDENTITY {DEFAULT USING INDEX <i>indexName</i> FULL NOTHING}	This operation is not supported in VCI.
CLUSTER	-	A table that contains a VCI and VCI cannot be specified. If the operation is required, delete the VCI using DROP INDEX, and re-create it using CREATE INDEX after completing the operation.
CREATE INDEX	UNIQUE	The subcommands on the left cannot be performed if a VCI is specified.

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

Client application command

- Operations that cannot be performed on relations containing a VCI

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

13.2.2 Data Preload Feature

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

```
db01=# SELECT pgx_prewarm_vci('idx_vci');
```



See

Refer to "[B.5 VCI Data Load Control Function](#)" for information on `pgx_prewarm_vci`.

Chapter 14 Parallel Query

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

- CPU load calculation
- Increase of workers during runtime
- Statistics view displays the action state

14.1 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 checks the available CPU resources when determining the number of workers for a parallel query.

When checking the available CPU resources, the system checks the CPU usage (whether there are sufficient free system resources) at the time the SQL statement is executed, and a parallel query is only planned if there is enough free CPU usage for two or more cores on the server.

For example, if a server has four cores and the total free space from core 1 to core 4 is 200% or more, a parallel query will be planned, but if it is less than 200%, a parallel query will not be planned.

14.2 Increase of Workers during Runtime

Parallel query can cause CPU overload because workers cannot be added, for example, because there are not enough workers or no workers at the start of the query.

This Fujitsu Enterprise Postgres enhancement allows systems to allocate additional workers during query execution (if there are workers available at the time). The ability to increase workers during runtime is available only with parallel query. This improves query performance.

14.3 Statistics View Displays the Action State

You can check the action state by using the existing statistics collector. Monitoring the statistics enables you to check the action state of parallel processing through, for example, changes in the system load.

The following table describes the column added to the statistics views:

pg_stat_all_tables

Column	Data type	Description
parallel_query	bigint	Number of parallel query initialized in the table.

The parallel query column has also been added to the following views:

- pg_stat_sys_tables
- pg_stat_user_tables
- pg_stat_xact_all_tables
- pg_stat_xact_sys_tables
- pg_stat_xact_user_tables



See

Refer to "The Statistics Collector" in "Server Administration" in the PostgreSQL Documentation for information on the statistics collector and views.

Values will be set for the pgx_stat_sql view in Fujitsu Enterprise Postgres if parallel query is run.



See

.....
Refer to "D.5 pgx_stat_sql" for information on pgx_stat_sql view.
.....

Chapter 15 High-Speed Data Load

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

This feature is available only in the Advanced Edition.

Information

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.

15.1 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](#)

15.1.1 Deciding whether to Install

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

Database server memory

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

See

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

Frequency of checkpoints

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

Example

```
LOG:  checkpoints are occurring too frequently (19 seconds apart)
HINT:  Consider increasing the configuration parameter "max_wal_size".
```

15.1.2 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.

Information

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 "[15.1.1 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.

15.1.3 Setup

This section describes how to set up high-speed data load.

Setup flow

1. [Setting Parameters](#)
2. [Installing the Extension](#)

15.1.3.1 Setting Parameters

Set the parameters required for high-speed data load in `postgresql.conf`. After that, start or restart the instance.

The table below lists the `postgresql.conf` parameters that must be changed, and the values that must be added to their current values. After editing `postgresql.conf`, start or restart the instance.

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

As shown in the following example, 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.

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
```

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

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

 See

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

15.1.3.2 Installing the Extension

Execute CREATE EXTENSION to install the high-speed data load extension. The extension needs to be installed on each database. Also, extensions can only be installed on the public schema by a superuser.

 Example

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

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

15.2 Using High-Speed Data Load

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

15.2.1 Loading Data

To load data from a file into a Fujitsu Enterprise Postgres table, execute the pgx_loader command in load mode. The data inserted using this feature is dumped as a COPY command by the pg_dump command and the pg_dumpall command.

 Example

The example below loads the file /path/to/data.csv (2000 records) into table tbl using a degree of parallelism of 3.

```
$ pgx_loader load -j 3 -c "COPY tbl FROM '/path/to/data.csv' WITH CSV"
LOAD 2000
```

 Point

If an external file contains data that violates the format or constraints, the data load may fail partway through, resulting in delays for routine tasks such as nightly batch processing. Therefore, it is recommended to remove the invalid data before executing the data load.



- 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.

15.2.2 Checking Progress

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



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

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

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

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

Refer to "D.7 `pgx_stat_progress_loader`" for information on the `pgx_stat_progress_loader` view.



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

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

		1000														
9313		222		testdb		333		COPY FROM		FILE		0		0		63000
		1000														
9314		222		testdb		333		COPY FROM		FILE		0		0		63000
		1000														

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

15.2.3 Recovering from a Data Load that Ended Abnormally

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

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

Checking for in-doubt transactions

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

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

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

Example

The example below retrieves the global transaction identifier (gid) of in-doubt transactions performed by the database role myrole and that used table tbl. The retrieved global transaction identifiers pgx_loader:9589 and pgx_loader:9590 identify in-doubt transactions.

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

2. Refer to the pg_prepared_xacts system view.

Check if the in-doubt transactions retrieved above exist.

Example

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

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



See

Refer to "[H.1 pgx_loader_state](#)" for information on the `pgx_loader_state` table.

Resolving in-doubt transactions

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

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



Example

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

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



Point

The recovery mode of the `pgx_loader` command only resolves transactions prepared by high-speed data load. For transactions prepared by an application using distributed transactions other than this feature, follow the procedure described in "[19.13 Actions in Response to Error in a Distributed Transaction](#)".

15.3 Removing High-Speed Data Load

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

15.3.1 Removing the Extension

Execute `DROP EXTENSION` to remove the high-speed data load 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.

The extension needs to be removed on each database. Also, extensions can only be removed on the public schema by a superuser.



Example

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

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

Chapter 16 Global Meta Cache

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

This feature can only be used in Advanced Edition.

16.1 Usage

Describes how to use the Global Meta Cache feature.

16.1.1 Deciding Whether to Enable the Global Meta Cache Feature

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

In particular, consider using Global Meta Cache if the total size of the meta cache for each process exceeds the amount of installed memory, or takes up a large portion of that memory, thereby squeezing memory allocations to the database cache or the Operating system file cache. Using Global Meta Cache may increase the time it takes to execute a single SQL to reference a meta cache on shared memory, but you can expect a greater benefit from being able to allocate more memory, such as for the database cache.

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

16.1.2 Estimating Memory for Global Meta Cache

To enable the Global Meta Cache feature, the `pgx_global_metacache` parameter must specify an upper limit on the size of the shared memory (Hereinafter, the GMC area) dedicated to Global Meta Cache. Ideally, this upper limit should be the size estimated in "[Appendix A Parameters](#)". Values lower than this can still work, but refer to "[16.1.3 How the GMC Memory Area Is Used](#)" on using the GMC area to understand the disadvantages.

16.1.3 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.

16.1.4 Enabling the Global Meta Cache Feature

To Enable the Global Meta Cache feature edit the `postgresql.conf` file and set the `pgx_global_metacache` parameter. Restarting the instance after editing the `postgresql.conf` file is required. Refer to "[Appendix A Parameters](#)" for information on the parameters.

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

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

Example

Here is an example postgresql.conf configuration:

```
pgx_global_metacache = 800 MB
```

Wait Events

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

[GMC Feature Wait Events]

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

Point

If GMCSweep is happened frequently, increase the pgx_global_metacache setting.

See

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

16.1.5 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.

16.2 Statistics

Describes the statistics for the Global Meta Cache feature.

16.2.1 System View

You can check the cache hit ratio and size of the GMC area in the system view pgx_stat_gmc. Refer to "D.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 "9.6 Monitoring Database Activity" in the Operations Guide for information on the statistics.

Chapter 17 Local Meta Cache Limit

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

This feature is available only in the Advanced Edition.

17.1 Usage

Describes how to use the Local Meta Cache Limit feature.

17.1.1 Deciding Whether to Enable the Local Meta Cache Limit Feature

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

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

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

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

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

17.1.2 How to Set Parameters for the Local Meta Cache Limit Feature

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

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



Example

Here is an example postgresql.conf configuration:

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

17.1.3 Cache Removal when Local Meta Cache Limit is Enabled

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

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

In PostgreSQL, in order to acquire memory at high speed, a memory block of a certain size is acquired from the OS, and a small memory is cut out from the block and used. The memory for the metacache is cut out in the same way. Therefore, it is possible to return the memory block to the OS by destroying all the meta caches scattered throughout the memory block. When this happens, the next SQL execution will

be slowed down due to the re-creation of the meta cache. Therefore, upper limit of feature should be set to a value larger than the size of the meta cache used by at least one transaction.

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

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

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

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

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

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

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

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

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

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

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

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

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

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



Example

Here is an example postgresql.conf configuration:


```
pgx_cache_hit_log_interval= 30min
```

Chapter 18 Backup/Recovery Using the Copy Command

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

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

This appendix describes backup/recovery using the copy command.



Point

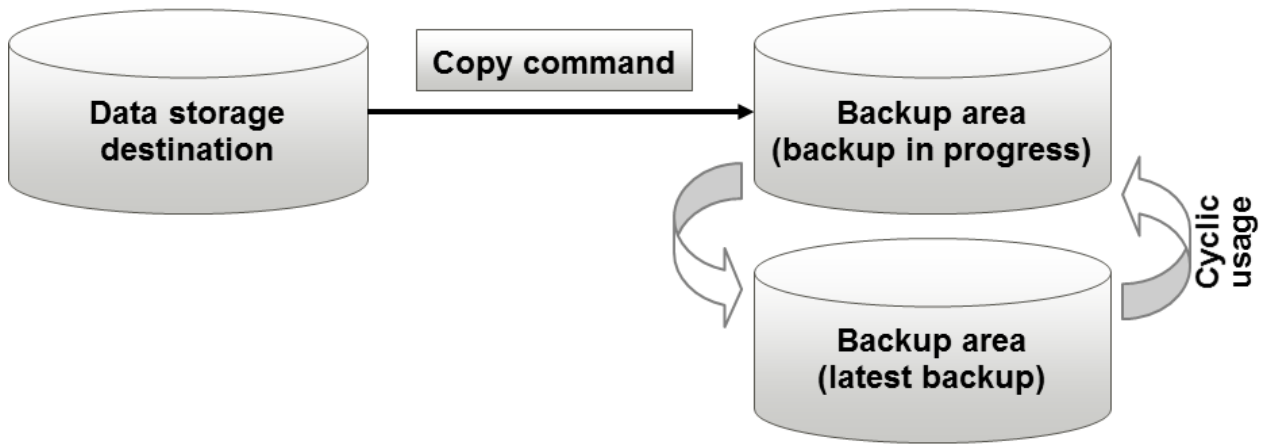
- By using the high-speed copy feature of the storage device to copy the data storage destination, the processing time for backup of large databases can be greatly reduced.
- It is also possible to back up only some tablespaces using the copy command. However, database resources not backed up using the copy command are still backed up to the backup data storage destination.
- 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.

18.1 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.




 **Information**

The backup data storage destination cannot be used as these backup areas used by the copy command.

Backup using the backup information file

The copy command must determine the backup destination on each backup, as it is necessary to cycle through the backup areas. Backup can be automated by using the backup information file, which contains information about the backup destination.



Information

The backup information file is prepared in the backup data storage destination by the `pgx_dmpall` command, and contains information that can be read or updated by the `copy` command. This file is managed by associating it with the latest backup successfully completed by the `pgx_dmpall` command, so the latest backup information relating to the `copy` command registered by the user can be retrieved. Additionally, the content of the backup information file can be displayed using the `pgx_rcvall` command.

Configuration of the copy command for backup

The `pgx_dmpall` command calls the `copy` command for backup after execution for the three modes below. It is therefore necessary for the `copy` command for backup to implement the required processing for each of the modes. The user can use any method to hand over backup information between modes within the `copy` command, such as creating temporary files.

- 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.

Configuration of the copy command for recovery

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

- restore mode

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

Point

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

Example

Using a bash script

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

```
;;
esac
```

Point

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

The sample is stored in the directory below:

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

- A sample script that uses OPC (an advanced copy feature of Fujitsu Storage ETERNUS disk array) is supplied. Refer to "[Appendix L Copy Command Samples that Use the Advanced Copy Feature of the ETERNUS Disk Array](#)" for details.

The samples cannot be used on SLES 15.

18.2 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 "[18.4 Copy Command Interface](#)" for details.

Performing backup

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

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

Example

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

Point

- To exclude up the database cluster from backup using the copy command, specify the --exclude-copy-cluster option.
- To back up only some tablespaces using the copy command, use the -P option specifying the full path of the file created in step 1 of preparation for backup.

See

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

Checking backup status

Use the pgx_rcvall command to check the backup status.

Execute the pgx_rcvall command with the -l option specified to output backup data information. If backup was performed using the copy command, the resources backed up using the copy command will also be output.

Example

```
$ pgx_rcvall -l -D /database/inst1  
Date           Status   Dir                                     Resources backed up by the copy command  
2022-03-01 13:30:40 COMPLETE /backup/inst1/2022-03-01_13-30-40 pg_data,dbspace,indexspace
```

18.3 Recovery Using the Copy Command

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

Determining the backup area of the latest backup

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

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

Example

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

Perform recovery

Execute the pgx_rcvall command with the -Y option specifying the full path of the copy command for recovery created in step 3 of the preparation for backup described in "18.2 Backup Using the Copy Command".

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

Example

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

Point

If the latest backup was performed using the copy command, the `pgx_rcvall` command automatically recognizes which database resources were backed up using the copy command, or whether resources were backed up to the backup data storage destination. Therefore, recovery can be performed by simply executing the `pgx_rcvall` command specifying the copy command for recovery.

See

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

18.4 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.

18.4.1 Copy Command for Backup

Feature

User exit (for the copy command) called from the `pgx_dmpall` command.

Format

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

If the operation mode is "prepare"

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

If the operation mode is "backup"

```
copyCommandName backup
```

If the operation mode is "finalize"

```
copyCommandName finalize 'pathOfBackupInfoFile'
```

Argument

- Operation mode

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

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

Point

If using high-speed copy of the storage device (which performs high-speed retrieval of snapshots and copies data to different physical areas), it is possible to invoke the snapshot process in backup mode, and the copy process in finalize mode.

- Full path of the backup information file

Full path of the backup information file of the latest backup, enclosed in single quotation marks. If a backup has not been performed, specify '.'.

- Full path of the backup target list file

Full path of the file containing the resources to be backed up using the copy command, enclosed in single quotation marks. One of the following is described in each resource name.

Resource	Description
Database cluster	pg_data
Tablespace	Tablespace name

Example

To back up the database cluster and the tablespaces dbspace and indexspace using the copy command, the file should contain the following:

```
pg_data
dbspace
indexspace
```

Information

The encoding of resource names output to the backup target list file by the pgx_dmpall command is the encoding used when this command connects to the database with auto specified for the client_encoding parameter, and is dependent on the locale at the time of command execution.

The number of arguments vary depending on operation mode. The argument of each operation mode is as follows.

Operation mode	First argument	Second argument	Third argument
prepare	Operation mode	Backup information file path name	Backup target list file path name
backup		None	None
finalize		Backup information file path name	

Additionally, the access permissions for the backup information file and backup target list file are different depending on the operation mode. The access permissions of each operation mode are as follows.

Operation mode	Backup information file	Backup target list file
prepare	Can be viewed by the instance administrator only	Can be viewed by the instance administrator only
backup	-	-

Operation mode	Backup information file	Backup target list file
finalize	Can be viewed and updated by the instance administrator only	-

Return value

Return value	Description
0	Normal end The <code>pgx_dmpall</code> command continues processing.
Other than 0	Abnormal end The <code>pgx_dmpall</code> command terminates in error.

Description

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

18.4.2 Copy Command for Recovery

Feature

User exit (for the copy command) called from the `pgx_rcvall` command.

Format

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

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

Argument

- Operation mode

Mode	Description
restore	Performs restore.

- Full path of the backup information file

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

- Full path of the backup target list file

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

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

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

Return value

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

Description

- The copy command operates with the privileges of the operating system user who executed the pgx_rcvall command. Therefore, grant copy command execution privileges to users who will execute the pgx_rcvall command. Additionally, have the copy command change users as necessary.
- The copy command is called once only in restore mode.
- To copy multiple resources simultaneously, have the copy command copy them in parallel.
- The backup information file and backup target list file cannot be deleted. Additionally, the privileges cannot be changed.
- Standard output and standard error output of the copy command are output to the terminal where the pgx_rcvall command was executed.
- If the copy command becomes unresponsive, the pgx_rcvall command will also become unresponsive. If the status of the copy command is deemed to be unresponsive by the operating system, use an operating system command to forcibly stop it.
- Output the copy command execution trace and the result to a temporary file, so that if it terminates in error, the cause can be investigated at a later time.
- Do not execute Fujitsu Enterprise Postgres commands or PostgreSQL applications in the copy command.
- There may be files and directories not required for recovery using the archive log included in the backup, such as postmaster.pid, pg_wal/*subdirectory* and pg_replslot in the database cluster. If such unnecessary files and directories exist, have the copy command delete them after the restore.

Chapter 19 Actions when an Error Occurs

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

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

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

Even if a disk is not defective, the same input-output error messages, as those generated when the disk is defective, may be output. The recovery actions differ for these error messages.

Check the status of the disk, and select one of the following actions:

- If the disk is defective

Refer to "[19.1 Recovering from Disk Failure \(Hardware\)](#)", and take actions accordingly.

- If the disk is not defective

Refer to "[19.14 I/O Errors Other than Disk Failure](#)", and take actions accordingly.

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

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

Determining the cause of an error

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



See

Refer to "Configuring Parameters" in the Installation and Setup Guide for Server for information on server logs.

Approximate recovery time

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

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

- Data storage destination or transaction log storage destination

```
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

19.1 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:

- [19.1.1 Using WebAdmin](#)
- [19.1.2 Using Server Command](#)



Point

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

19.1.1 Using WebAdmin

Recover a failed disk using WebAdmin.

In a streaming replication configuration, the master instance can be recovered, but the standby instance cannot. If disk failure occurs on a standby instance, it may be necessary to delete and re-create the instance. Also, recovering the master instance stops streaming replication to and from all standby instances. In such an event, the standby instances can be promoted to standalone instances or can be deleted and re-created

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

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

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

1. Stop applications

Stop applications that are using the database.

2. Stop the instance

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

3. Recover the failed disk

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

4. Create a tablespace directory

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

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

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

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

6. Recover the database cluster

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

7. Run recovery

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

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

8. Resume applications

Resume applications that are using the database.

Point

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

If this happens, refer to "[19.10.3 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.

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.

19.1.2 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 "[19.11 Actions in Response to Failure to Stop an Instance](#)".

3. Recover the failed disk

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

4. Create a storage destination directory

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

Example

To create a data storage destination directory:

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



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

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

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

6. Recover the database cluster

Recover the database cluster using the backup data.

Specify the following in the `pgx_rcvall` command:

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

Example

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

If recovery fails

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

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



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

```
FATAL: the database system is starting up
```

7. Start the instance

Start the instance.

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

8. Resume applications

Resume applications that are using the database.

If failure occurred on the backup data storage disk

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

There are two methods of taking action:

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

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

No	Step	Instance stopped	
		No	Yes
1	Confirm that transaction log mirroring has stopped	Y	N
2	Stop output of archive logs	Y	N
3	Stop applications	N	Y
4	Stop the instance	N	Y
5	Recover the failed disk	Y	Y
6	Create a backup data storage destination directory	Y	Y
7	Resume output of archive logs	Y	N
8	Resume transaction log mirroring	Y	N
9	Start the instance	N	Y
10	Run backup	Y	Y
11	Resume applications	N	Y

Y: Required

N: Not required

The procedure is as follows:

If an instance has not been stopped

1. Confirm that transaction log mirroring has stopped

Use the following SQL function to confirm that transaction log mirroring has stopped.

```
postgres=# SELECT pgx_is_wal_multiplexing_paused();
pgx_is_wal_multiplexing_paused
-----
t
(1 row)
```

If transaction log mirroring has not stopped, then stop it using the following SQL function.

```
postgres=# SELECT pgx_pause_wal_multiplexing();
LOG:  multiplexing of transaction log files has been stopped
pgx_pause_wal_multiplexing
-----
(1 row)
```

2. Stop output of archive logs

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

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

- Changing archive_command

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

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

- Reload the configuration file

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

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

3. Recover the failed disk

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

4. Create a backup data storage destination

Create a backup data storage destination.

Example

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

Refer to "[3.2.2 Using Server Commands](#)" for information on how to create a backup data storage destination.

5. Resume output of archive logs

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

6. Resume transaction log mirroring

Execute the `pgx_resume_wal_multiplexing` SQL function.

Example

```
SELECT pgx_resume_wal_multiplexing()
```

7. Run backup

Use the `pgx_dmpall` command to back up the database cluster.

Specify the following value in the `pgx_dmpall` command:

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

Example

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

If an instance has been stopped

1. Stop applications

Stop applications that are using the database.

2. Stop the instance

Stop the instance. Refer to "[2.1.2 Using Server Commands](#)" for details.

If the instance fails to stop, refer to "[19.11 Actions in Response to Failure to Stop an Instance](#)".

3. Recover the failed disk

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

4. Create a backup data storage destination

Create a backup data storage destination.

Example

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

Refer to "[3.2.2 Using Server Commands](#)" for details.

5. Start the instance

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

6. Run backup

Use the `pgx_dmpall` command to back up the database cluster.

Specify the following value in the `pgx_dmpall` command:

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

Example

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

7. Resume applications

Resume applications that are using the database.



See

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

19.2 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:

- [19.2.1 Using WebAdmin](#)
- [19.2.2 Using the `pgx_rcvall` Command](#)



Point

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

19.2.1 Using WebAdmin

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

Refer to "[19.1.1 Using WebAdmin](#)" for details.

19.2.2 Using the `pgx_rcvall` Command

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

Follow the procedure below to recover the data storage disk.

1. Stop applications

Stop applications that are using the database.

2. Stop the instance

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

If the instance fails to stop, refer to "[19.11 Actions in Response to Failure to Stop an Instance](#)".

3. Confirm the backup date and time

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

Specify the following values in the `pgx_rcvall` command:

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

Example

```
> pgx_rcvall -D /database/inst1 -B /backup/inst1 -l
Date                Status             Dir
2022-03-20 10:00:00 COMPLETE          /backup/inst1/2022-03-20_10-00-00
```

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

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

5. Recover the database cluster

Use the `pgx_rcvall` command to recover the database cluster.

Specify the following values in the `pgx_rcvall` command:

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

Example

In the following examples, "March 20, 2022 10:00:00" is specified as the recovery time.

```
> pgx_rcvall -D /database/inst1 -B /backup/inst1 -e '2022-03-20 10:00:00'
```

If recovery fails

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.

Information

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

```
FATAL: the database system is starting up
```

6. Start the instance

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

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

7. Resume applications

Resume applications that are using the database.

See

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

19.3 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:

- [19.3.1 Using WebAdmin](#)
- [19.3.2 Using the pgx_rcvall Command](#)

Point

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

19.3.1 Using WebAdmin

Recover data to a backup point using WebAdmin.

In a streaming replication configuration, the master instance can be recovered, but the standby instance cannot. If disk failure occurs on a standby instance, it may be necessary to delete and re-create the instance. Also, recovering the master instance stops streaming replication to and from all standby instances. In such an event, the standby instances can be promoted to standalone instances or can be deleted and re-created.

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

1. Stop applications

Stop applications that are using the database.

2. Stop the instance


Stop the instance. Refer to ["2.1.1 Using WebAdmin"](#) for information on how to stop an instance.

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

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

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

4. Recover the database cluster

Log in to WebAdmin, and in the [Instances] tab, select the instance to be recovered and click .

5. Recover to the backup point

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

Recovery is performed. An instance is automatically started when recovery is successful.

6. Resume database operations

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

19.3.2 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 ["19.11 Actions in Response to Failure to Stop an Instance"](#).

3. Confirm the restore point

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

Specify the following values in the `pgx_rcvall` command:

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

Example

```
> pgx_rcvall -D /database/inst1 -B /backup/inst1 -l
Date           Status        Dir
2022-03-01 10:00:00  COMPLETE    /backup/inst1/2022-03-01_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_20220303_1".

```
> pgx_rcvall -D /database/inst1 -B /backup/inst1 -n batch_20220303_1
```

If recovery fails

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.

Information

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

```
FATAL: the database system is starting up
```

6. Start the instance

Start the instance.

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

7. Restart operation of the database

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

See

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

19.4 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 "19.4.1 When using the view (pg_stat_activity)")
- `ps` command (refer to "19.4.2 Using the ps Command")
- pgAdmin (refer to "19.4.3 Using pgAdmin")

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

19.4.1 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 (<x>) (*1)
Type "help" for help.
```

*1: <x> indicates the PostgreSQL version on which this product is based.

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

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

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

Example

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

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



See

- Refer to "System Administration Functions" under "The SQL Language" in the PostgreSQL Documentation for information on pg_terminate_backend.
- Refer to "Notes on Application Compatibility" in the Application Development Guide for information on how to maintain application compatibility.

19.4.2 Using the ps Command

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

1. Execute the ps command.

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

```
> ps axwfo user,pid,ppid,TTY,command | grep postgres
fsepuser 19174 18027 pts/1          \_ grep postgres
fsepuser 20517      1 ?          /opt/fsepv<x>server64/bin/postgres -D /disk01/data
fsepuser 20518 20517 ?          \_ postgres: logger
fsepuser 20520 20517 ?          \_ postgres: checkpointer
fsepuser 20521 20517 ?          \_ postgres: background writer
fsepuser 20522 20517 ?          \_ postgres: walwriter
fsepuser 20523 20517 ?          \_ postgres: autovacuum launcher
fsepuser 20524 20517 ?          \_ postgres: archiver
fsepuser 20525 20517 ?          \_ postgres: logical replication launcher
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)
```



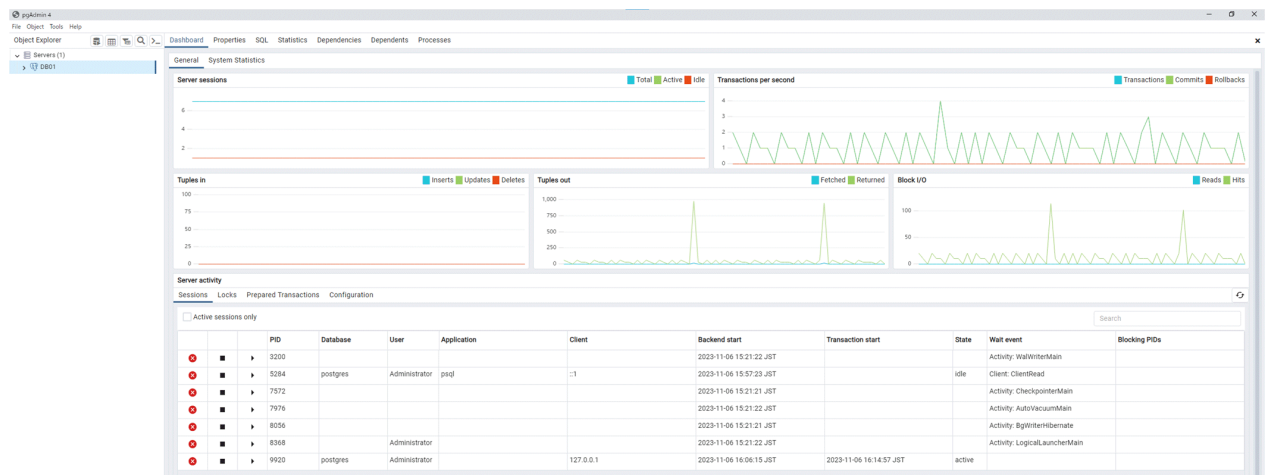
- Refer to "System Administration Functions" under "The SQL Language" in the PostgreSQL Documentation for information on `pg_terminate_backend`.
- Refer to "Notes on Application Compatibility" in the Application Development Guide for information on how to maintain application compatibility.

19.4.3 Using pgAdmin

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

1. In the [Browser] pane, click the relevant database server.
2. Close the client connections that have been in a wait state for an extended period.

Click the [Dashboard] tab. In the [Server activity] section, select the connections that have been in an "idle" or "idle in transaction" state for an extended period. For each of these connections, click  to close the session.



19.5 Actions in Response to an Access Error

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

- Data storage destination

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



Refer to "Preparing Directories to Deploy Resources" under "Setup" in the Installation and Setup Guide for Server for information on the privileges required for the directory.

19.6 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:

- [19.6.1 Using a Tablespace](#)
- [19.6.2 Replacing the Disk with a Larger Capacity Disk](#)

19.6.1 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.

19.6.2 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:

- [19.6.2.1 Using WebAdmin](#)
- [19.6.2.2 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.



.....
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.
.....

19.6.2.1 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 "[19.1.1 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.

19.6.2.2 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 "[19.11 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

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.

 Information

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

```
FATAL: the database system is starting up
```

 See

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

9. Start the instance

Start the instance.

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

10. Resume applications

Resume applications that are using the database.

11. Restore files

Restore the files backed up in step 1.

19.7 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:

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

19.7.1 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:

- [19.7.1.1 Using WebAdmin](#)
- [19.7.1.2 Using Server Commands](#)

19.7.1.1 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
```

19.7.1.2 Using Server Commands

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

There are two methods of taking action:

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

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

No	Step	Instance stopped	
		No	Yes
1	Stop transaction log mirroring	Y	N
2	Stop output of archive logs	Y	N
3	Stop applications	N	Y
4	Stop the instance	N	Y
5	Temporarily save backup data	Y	Y
6	Resume output of archive logs	Y	N
7	Resume transaction log mirroring	Y	N
8	Start an instance	N	Y
9	Run backup	Y	Y
10	Resume applications	N	Y
11	Delete temporarily saved backup data	Y	Y

Y: Required

N: Not required

The procedure is as follows:

Performing recovery while the instance is active

1. Stop transaction log mirroring

Stop transaction log mirroring.

```
postgres=# SELECT pgx_pause_wal_multiplexing();
LOG:  multiplexing of transaction log files has been stopped
pgx_pause_wal_multiplexing
-----
(1 row)
```

2. Stop output of archive logs

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

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

- Changing the archive_command parameter

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

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

- Reloading the configuration file

Run the pg_ctl reload command or the pg_reload_conf SQL function.

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

3. Temporarily save backup data

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

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

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

Example

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

4. Resume output of archive logs

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

5. Resume transaction log mirroring

Execute the `pgx_resume_wal_multiplexing` SQL function.

Example

```
SELECT pgx_resume_wal_multiplexing()
```

6. Run backup

Use the `pgx_dmpall` command to back up the database cluster.

Specify the following option in the `pgx_dmpall` command:

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

Example

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

7. Delete temporarily saved backup data

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

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

Example

```
> rm -rf /mnt/usb/backup
```

If an instance has been stopped

1. Stop applications

Stop applications that are using the database.

2. Stop the instance

Stop the instance. Refer to "[2.1.2 Using Server Commands](#)" for details.

If the instance fails to stop, refer to "[19.11 Actions in Response to Failure to Stop an Instance](#)".

3. Temporarily save backup data

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

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

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

Example

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

4. Start the instance

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

5. Run backup

Use the `pgx_dmpall` command to back up the database cluster.

Specify the following value in the `pgx_dmpall` command:

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

Example

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

6. Resume applications

Resume applications that are using the database.

7. Delete temporarily saved backup data

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

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

Example

```
> rm -rf /mnt/usb/backup
```



See

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

19.7.2 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:

- [19.7.2.1 Using WebAdmin](#)
- [19.7.2.2 Using Server Commands](#)

19.7.2.1 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. Stop applications

Stop applications that are using the database.

4. Replace with a larger capacity disk

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

5. 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 "19.1.1 Using WebAdmin".

6. Restore files

Restore the files backed up in step 1.

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
```

19.7.2.2 Using Server Commands

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

There are two methods of taking action:

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

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

No	Step	Instance stopped	
		No	Yes
1	Back up files	Y	Y
2	Temporarily save backup data	Y	Y
3	Confirm that transaction log mirroring has stopped	Y	N
4	Stop output of archive logs	Y	N
5	Stop applications	N	Y

No	Step	Instance stopped	
		No	Yes
6	Stop the instance	N	Y
7	Replace with a larger capacity disk	Y	Y
8	Create a backup storage directory	Y	Y
9	Resume output of archive logs	Y	N
10	Resume transaction log mirroring	Y	N
11	Start the instance	N	Y
12	Run backup	Y	Y
13	Resume applications	N	Y
14	Restore files	Y	Y
15	Delete temporarily saved backup data	Y	Y

Y: Required

N: Not required

The procedure is as follows:

If an instance has not been stopped

1. Back up files

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

2. Temporarily save backup data

Save the backup data to a different directory.

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

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

Example

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

3. Confirm that transaction log mirroring has stopped

Use the following SQL function to confirm that transaction log mirroring has stopped.

```
postgres=# SELECT pgx_is_wal_multiplexing_paused();
pgx_is_wal_multiplexing_paused
-----
t
(1 row)
```

If transaction log mirroring has not stopped, then stop it using the following SQL function.

```
postgres=# SELECT pgx_pause_wal_multiplexing();
LOG: multiplexing of transaction log files has been stopped
pgx_pause_wal_multiplexing
-----
```

```
(1 row)
```

4. Stop output of archive logs

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

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

- Changing the `archive_command` parameter

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

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

- Reloading the configuration file

Run the `pg_ctl reload` command or the `pg_reload_conf` SQL function.

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

5. Replace with a larger capacity disk

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

6. Create a backup data storage destination

Create a backup data storage destination.

Example

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

Refer to "[3.2.2 Using Server Commands](#)" for details.

7. Resume output of archive logs

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

8. Resume transaction log mirroring

Execute the `pgx_resume_wal_multiplexing` SQL function.

Example

```
SELECT pgx_resume_wal_multiplexing()
```

9. Run backup

Use the `pgx_dmpall` command to back up the database cluster.

Specify the following value in the `pgx_dmpall` command:

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

Example

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

10. Restore files

Restore the files backed up in step 1.

11. Delete temporarily saved backup data

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

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

Example

```
> rm -rf /mnt/usb/backup
```

If an instance has been stopped

1. Back up files

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

2. Temporarily save backup data

Save the backup data to a different directory.

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

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

Example

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

3. Stop applications

Stop applications that are using the database.

4. Stop the instance

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

If the instance fails to stop, refer to "[19.11 Actions in Response to Failure to Stop an Instance](#)".

5. Replace with a larger capacity disk

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

6. Create a backup data storage destination

Create a backup data storage destination.

Example

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

Refer to "[3.2.2 Using Server Commands](#)" for details.

7. Start the instance

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

8. Run backup

Use the `pgx_dmpall` command to back up the database cluster.

Specify the following value in the `pgx_dmpall` command:

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

Example

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

9. Resume applications

Resume applications that are using the database.

10. Restore files

Restore the files backed up in step 1.

11. Delete temporarily saved backup data

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

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

Example

```
> rm -rf /mnt/usb/backup
```



See

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

19.8 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.

19.8.1 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:

- [19.8.1.1 Using WebAdmin](#)
- [19.8.1.2 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.



Point

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.

19.8.1.1 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 "[19.1.1 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.

19.8.1.2 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 "[19.11 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

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.

 Information

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

```
FATAL: the database system is starting up
```

 See

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

9. Start the instance

Start the instance.

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

10. Resume applications

Resume applications that are using the database.

11. Restore files

Restore the files backed up in step 1.

19.9 Errors in More Than One Storage Disk

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

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

Recreate the instance, and rebuild the runtime environment.

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



See

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

19.10 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.

19.10.1 Errors in the Configuration File

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

- postgresql.conf
- pg_hba.conf



See

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

- "Configuring Parameters" in the Installation and Setup Guide for Server
 - "[Appendix A Parameters](#)"
 - "Server Configuration" and "Client Authentication" under "Server Administration" in the PostgreSQL Documentation
-

19.10.2 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 "[19.14.2 Errors Caused by Power Failure or Mounting Issues](#)", and take actions accordingly.

19.10.3 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:

- [19.10.3.1 Using WebAdmin](#)
- [19.10.3.2 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.

19.10.3.1 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 "[19.2.1 Using WebAdmin](#)" for details.

19.10.3.2 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 "[19.2.2 Using the pgx_rcvall Command](#)" for details.

19.11 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:

- [19.11.1 Using WebAdmin](#)
- [19.11.2 Using Server Commands](#)

19.11.1 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.

19.11.2 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.

19.11.2.1 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 "19.11.2.2 Stopping the Instance Using the Immediate Mode" or "19.11.2.3 Forcibly Stopping the Server Process".



Example

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

19.11.2.2 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 "19.11.2.3 Forcibly Stopping the Server Process".



Example

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

19.11.2.3 Forcibly Stopping the Server Process

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

The procedure is as follows:

1. Execute the ps command

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

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

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
```

19.12 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.

19.13 Actions in Response to Error in a Distributed Transaction

If a system failure (such as server failure) occurs in an application that uses distributed transactions (such as .NET TransactionScope), then transactions may be changed to the in-doubt state.

At that point, resources accessed by the transaction will be locked, and rendered unusable by other transactions.

The following describes how to check for in-doubt transactions, and how to resolve them.

How to check for in-doubt transactions

The following shows how to check for them:

If the server fails

1. An in-doubt transaction will have occurred if a message similar to the one below is output to the log when the server is restarted.

Example

```
LOG: Restoring prepared transaction 2103.
```

2. Refer to system view pg_prepared_xacts to obtain information about the prepared transaction.

If the transaction identifier of the prepared transaction in the list (in the transaction column of pg_prepared_xacts) is the same as the identifier of the in-doubt transaction obtained from the log output when the server was restarted, then that row is the information about the in-doubt transaction.

Example

```
postgres=# select * from pg_prepared_xacts;
 transaction | gid      | prepared | owner  | database
-----+-----+-----+-----+-----
```

```
2103 | 374cc221-f6dc-4b73-9d62-d4fec9b430cd | 2022-03-06 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.

Point

As you can see from the explanations in this section, there is no one way to definitively determine in-doubt transactions.

Consider collecting other supplementary information (for example, logging on the client) or performing other operations (for example, allocating database users per job).

How to resolve in-doubt transactions

From the system view `pg_prepared_xacts` mentioned above, obtain the global transaction identifier (in the `gid` column of `pg_prepared_xacts`) for the in-doubt transaction, and issue either a `ROLLBACK PREPARED` statement or `COMMIT PREPARED` statement to resolve the in-doubt transaction.

Example

- Rolling back in-doubt transactions

```
postgres=# rollback prepared '374cc221-f6dc-4b73-9d62-d4fec9b430cd';
ROLLBACK PREPARED
```

- Committing in-doubt transactions

```
postgres=# commit prepared '374cc221-f6dc-4b73-9d62-d4fec9b430cd';
COMMIT PREPARED
```

19.14 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.

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

19.14.1 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.

19.14.2 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.

19.15 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.

19.15.1 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

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

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



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

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

19.15.2 Mirroring Controller Anomalies

The following conditions will cause a Mirroring Controller anomaly:

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

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

The following occurs when a Mirroring Controller anomaly is detected:

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

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

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

Appendix A Parameters

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

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

Information

The maximum value that can be expressed as a 4-byte signed integer changes according to the operating system. Follow the definition of the operating system in use.

- core_directory (string)

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

- core_contents (string)

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

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

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

- keystore_location (string)

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

Cannot be specified with the tde_kms.kms_conninfo_file parameter.

- tde_kms.plugin_path (string)

When connecting to the key management system using a plug-in, specify the directory that stores the plug-in with an absolute path. Only database administrators should be able to store plugins in this directory. It can only be set by specifying a parameter when starting the instance.

- tde_kms.enable_shared_dek (boolean)

Enables or disables sharing of data encryption keys between backend processes for each encrypted tablespace in transparent data encryption. Default is off. It can only be set by specifying a parameter when starting the instance.

- tde_kms.max_shared_dek (numerical value)

Specify the maximum number of shared data encryption keys when sharing data encryption keys per tablespace in transparent data encryption. Default is 1000. It can only be set by specifying a parameter when starting the instance.

- tde_kms.kms_conninfo_file (string)

When using the key management service as a key store, specifies the file that contains the connection information for the key management system. Cannot be specified with the keystore_location parameter.

Create a connection information file for the key management system in one of the following format:

```
kmipl kms-name address port auth-method [auth-options]
custom kms-name plugin-name [plugin-options] [extra-args]
```

Specify one of the following methods to connect to the key management system.

- kmip

Access this key management system using the KMIP protocol. Take out and use the encryption key in this key management system.

- custom

Access the key management system using a plugin module. The encryption/decryption process is done inside the plugin or inside the key management system. Fujitsu Enterprise Postgres does not use cryptographic keys directly.

For type kmip

- kms-name

The key management system name assigned to the key management system and specified when declaring the master encryption key or opening the keystore. The name of the key management system must be unique within this file. The key management system name must be a string of no more than 63 characters beginning with a-z, consisting of a-z, a number (0-9), and an underscore. Upper and lower case letters are the same.

- address

Specifies the host name or IP address of the key management service.

- port

Specifies the port number on which the key management service listens for services.

- auth-method

Specifies the authentication method for the key management service.

- auth-options

The auth-method is followed by the authentication method options. You can specify multiple options in a name = value field.

Authentication method when using a key management service of type kmip

cert

A certificate is used to authenticate the KMIP server and the client, Fujitsu Enterprise Postgres, to each other. The auth-options can be.

- sslcert

Specifies the file name of the client certificate. The corresponding format is PEM format.

- sslkey

Specifies the file name of the private key used for the client certificate. The corresponding format is PEM format. If you choose to encrypt the file with a passphrase, use a passphrase that is no more than 1023 bytes long.

- sslkeypassphrase-obf

Specifies the file that contains the obfuscated passphrase for the private key file specified by sslkey. This option allows the keystore to be opened automatically when the server starts. The `pgx _ keystore` command creates obfuscated files. It can be omitted.

- sslrootcert

Specifies the file name of the SSL Certificate Authority certificate. The corresponding format is PEM format. Used to verify the server certificate of the connection destination.

Example

```
kmip mykmipsvr mykmipsvr.example.com 5696 cert sslcert=postgres.crt  
sslkey=postgres.key sslrootcert=root.crt
```

Information

cert authentication does not verify that the server you are connecting to is the same server you are trying to connect to. Any server using a server certificate that is signed with the certificate of the certificate authority specified in sslrootcert is considered the correct destination. To avoid problems with this behavior, consider using your own CA or self-signed certificate for the KMIP server.

For type custom

- kms-name

The key management system name given to the key management system, specified when declaring a master encryption key or opening a keystore. The name of the key management system must be unique within this file. The key management system name must be a string of up to 63 characters starting with a-z and consisting of a-z, digits (0-9) and underscores. Uppercase and lowercase letters are considered the same.

- plugin-name

Specify the name of the plugin. Use the file with the same name as the plugin name in the directory specified by tde_kms.plugin_path as the plugin module. The file must have execution privilege for the OS user that starts the Fujitsu Enterprise Postgres server.

- plugin-options

Specify other options for the plugin. Multiple options can be specified in name=value format fields. You can specify the following option names.

- kms-secret-obf

Specify the file containing the obfuscated KMS secret when enabling automatic opening of the keystore using transparent data encryption. The obfuscated file is created with the pgx_keystore command. The obfuscated file contents are decrypted by Fujitsu Enterprise Postgres and passed to the plugin. Can be omitted if you are not using transparent data encryption to enable automatic opening of the keystore.

- extra-args

Specify additional arguments to pass to the specified shell command in the form arg=value. If you specify multiple extra-args, the values are passed to the shell command in that order.

Example

If you pass the plugin an additional argument: --profile user1

```
custom mykms mykms arg=--profile arg=user1
```

- tablespace_encryption_algorithm (string)

This parameter specifies the encryption algorithm for tablespaces that will be created. Valid values are "AES128", "AES256", and "none". If you specify "none", encryption is not performed. The default value is "none". To perform encryption, it is recommended that you specify "AES256". Only superusers can change this setting.

- backup_destination (string)

This parameter specifies the absolute path of the directory where pgx_dmpall will store the backup data. Specify a different location from other database clusters. This parameter can only be set when specified on starting an instance. It cannot be changed dynamically, while an instance is active.

Place this directory on a different disk from the data directory to be backed up and the tablespace directory. Ensure that users do not store arbitrary files in this directory, because the contents of this directory are managed by the database system.

- search_path (string)

When using the SUBSTR function compatible with Oracle databases, set "oracle" and "pg_catalog" in the search_path parameter. You must specify "oracle" before "pg_catalog".

Example

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

Information

- The search_path feature specifies the priority of the schema search path. The SUBSTR function in Oracle database is defined in the oracle schema.
- Refer to "Statement Behavior" under "Server Administration" in the PostgreSQL Documentation for information on search_path.

- track_waits (string)

This parameter enables collection of statistics for pgx_stat_lwlock and pgx_stat_latch.

- on: Enables collection of statistics.
- off: Disables collection of statistics.

If this parameter is omitted, "on" is assumed.

Only superusers can change this setting.

- track_sql (string)

This parameter enables collection of statistics for pgx_stat_sql.

- on: Enables collection of statistics.
- off: Disables collection of statistics.

If this parameter is omitted, "on" is assumed.

Only superusers can change this setting.

Parameters for the in-memory feature

- reserve_buffer_ratio (numerical value)

This parameter specifies the proportion of shared memory to be used for a stable buffer table.

- Minimum value: 0
- Maximum value: 80

If this parameter is omitted, 0 will be used.

- vci.cost_threshold (numerical value)

This parameter specifies the lowest cost that selects an execution plan that uses a VCI. If the cost of the best execution plan that does not use a VCI is lower than this value, that execution plan will be selected.

- Minimum value: 0
- Maximum value: Maximum value that can be expressed as a 4-byte signed integer

If this parameter is omitted or a value outside this range is specified, 18000 will be used.

- vci.control_max_workers (numerical value)

This parameter specifies the number of background workers that manage VCI. The number of workers for the entire instance is limited by max_worker_processes, so add the value specified here to max_worker_processes.

- Minimum value: 1
- Maximum value: 8388607

If this parameter is omitted or a value outside this range is specified, 8 will be used.

- vci.enable (string)

This parameter enables or disables VCI.

- on: Enables VCI.
- off: Disables VCI.

If this parameter is omitted, "on" will be used.

- vci.log_query (string)

This parameter enables or disables log output when VCI is not used due to insufficient memory specified by vci.max_local_ros.

- on: Enables log output.
- off: Disables log output.

If this parameter is omitted, "off" will be used.

- vci.maintenance_work_mem (numerical value)

This parameter specifies the maximum memory size used for maintenance of VCI (when executing CREATE INDEX, for example).

- Minimum value: 1 MB
- Maximum value: Maximum value that can be expressed as a 4-byte signed integer

If this parameter is omitted or a value outside this range is specified, 256 MB will be used.

- vci.max_local_ros (numerical value)

This parameter specifies the maximum memory size used for VCI scan.

- Minimum value: 64 MB
- Maximum value: Maximum value that can be expressed as a 4-byte signed integer

If this parameter is omitted or a value outside this range is specified, 64 MB will be used.

- vci.max_parallel_degree (numerical value)

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

A value from -8388607 to 8388607 can be specified.

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

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

- vci.shared_work_mem (numerical value)

This parameter specifies the maximum memory size used for VCI parallel scan.

- Minimum value: 32 MB
- Maximum value: Maximum value that can be expressed as a 4-byte signed integer

If this parameter is omitted or a value outside this range is specified, 1 GB will be used.

Parameters for the Global Meta Cache feature

- pgx_global_metacache (numerical value)

Specifies the memory size of the GMC area.

Specify a value calculated by the formula below.

A value lower than the calculated value will still work, but the meta cache may not be able to fit into the GMC area.

In this case, the system will discard the meta cache it thinks it is no longer needed, but if it is needed again, the meta cache will need to be expanded and will not perform well.

If the value is less than 10 MB and is set to a nonzero value that disables the feature, the database startup fails because the Global Meta Cache feature cannot operate.

A setting of 0 disables the Global Meta Cache feature. The default is 0.

Changing this setting requires restarting the database.

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

*1) Safety Factor (1.5)

This value takes into account the case where both GMC before and after the change temporarily exist at the same time in shared memory when the table definition is changed or the row of the system catalog is changed.

- track_gmc (string)

This parameter enables collection of statistics for pgx_stat_gmc.

- on: Enables collection of statistics.
- off: Disables collection of statistics.

If this parameter is omitted, "on" is used.

Only superusers can change this setting.

Parameters for the Local Meta Cache Limit feature

- pgx_catalog_cache_max_size(numerical value)

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

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

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

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

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

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

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

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

When Global Meta Cache is enabled:

```
(Number of tables to access + Number of indexes to access + Number of columns to access)
```

× 0.1KB × 1.5 (*1)

When Global Meta Cache is disabled:

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

*1) Safety Factor (1.5)

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

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

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

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

*1) Safety Factor (1.5)

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

4. If the value of 1. the value of 2. + the value of 3., the backend process can keep all caches to the extent allowed, so there is no need to limit the caches. If you want to cap for safety, set the value of 2. to pgx_catalog_cache_max_size and the value of 3. to pgx_relation_cache_max_size.
5. If the value of 1. < the value of 2. + the value of 3. then you need to limit the cache. However, this parameter does not limit the size of the cache used by a transaction. Therefore, take the following steps.
6. Calculate the catalog cache used by a transaction using the formula in 2.
7. Calculate the relation cache used by a transaction using the formula in 3.
8. If the value of 1. < the value of 6. + the value of 7., then the value of 1. needs to be increased. In other words, in some cases, it may be necessary to increase the installed memory or reduce the number of connections.
9. If the value of 1. the value of 6. + the value of 7., the condition of 1. can be satisfied by limiting the cache with this parameter. Divide the value of 1. by the ratio of 2. and 3. and set it as a parameter. Set the value distributed to 2. to pgx_catalog_cache_max_size and the value distributed to 3. to pgx_relation_cache_max_size.
10. The value calculated in 9. is a provisional value. If you cannot meet your target performance, first try to shift the focus of allocation to the relation cache. This is because when executing SQL, the relation cache generated based on the catalog cache is mainly referenced, so it is advantageous to leave a large amount of relation cache. If the performance is still not satisfied, adjust the parameters by referring to "[17.1.4 Performance Impact and Parameter Tuning of the Local Meta Cache Limit Feature](#)".

For Partitioning

Be careful when partitioning the table.

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

When specifying the parent table:

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

When specifying the child table directly:

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

Example)

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

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

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

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

- `pgx_relation_cache_max_size`(numerical value)

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

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

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

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

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

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

- `pgx_cache_hit_log_interval`(numerical value)

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

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

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

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

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

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

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

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

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

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

- Minimum value: 0
- Maximum value: 2147483647ms



See

.....
Refer to "Server Configuration" under "Server Administration" in the PostgreSQL Documentation for information on other `postgresql.conf` parameters.
.....

Parameters for the Policy-based Login Security

- userprofile_database(string)

Specifies the name of the database to which the background worker process connects. This parameter must specify a connectable database name. If you omit this parameter, then it is assumed to be "postgres".

The pseudo database "replication" cannot be specified for streaming replication.

If you change this parameter, reload the configuration file.

Parameters for using OpenSSL legacy algorithms

- openssl_conf(string)

Specifies the OpenSSL configuration file. Specifying a valid configuration file makes legacy algorithms available. Use the example below to prepare the configuration file in any directory.

If this parameter is not specified, the empty string is assumed.

This parameter can only be set by specifying the parameter at instance startup.

You cannot make dynamic changes during instance startup.



Example

```
openssl_conf = '/path/to/openssl.conf'
```

[OpenSSL Configuration File (openssl.conf) Example]

```
=====  
openssl_conf = openssl_init
```

```
[openssl_init]  
providers = provider_sect
```

```
[provider_sect]  
default = default_sect  
legacy = legacy_sect
```

```
[default_sect]  
activate = 1
```

```
[legacy_sect]  
activate = 1
```

```
=====  
.....
```

- openssl_modules(string)

Specifies the directory that contains additional OpenSSL modules.

Legacy algorithms are available by specifying 'server installation directory/lib/openssl-modules'.

If this parameter is not specified, the empty string is assumed.

This parameter can only be set by specifying the parameter at instance startup.

You cannot make dynamic changes during instance startup.

This parameter sets the OPENSSL_MODULES environment variable to be applied to the server process. Do not set the OPENSSL_MODULES environment variable in any way other than setting this parameter, as this may cause abnormal behavior.



Example

```
openssl_modules = '/opt/fsepv<X>server64/lib/openssl-modules'
```

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

Information

Legacy algorithms include the following encryption algorithms:

- BF
- CAST5
- DES-ECB
- DES-CBC
- MD4
- Whirlpool

Parameters for scheduling aggressive freeze for tuples

- `pgx_stat_vacuum_freeze.track_freeze_info` (boolean)

Collects and maintains information about aggressive freeze for tuples. The default is on.

- on: Enables information collection.
- off: Disables information collection.

- `pgx_stat_vacuum_freeze.monitoring_interval` (numerical value)

Aggregate statistics for aggressive freeze for tuples in this interval, so that statistics for each interval can be obtained in the `pgx_stat_freeze_results` table. Note that an aggressive freeze for tuples across this interval will still account for the overall tuple freeze at completion time. The default is 1 hour.

Set the aggressive freeze for tuples script to run at weekly intervals of days, daily intervals of hours, and hourly intervals of minutes.

- Minimum value: 1min
- Maximum value: 1d

If a value out of range is specified, 1h is set.

If this parameter is omitted, 1h is assumed.

- `pgx_stat_vacuum_freeze.log_retention_period` (numerical value)

Specify how long statistics are kept. The default is 90 days.

Set the period according to the work cycle. For example, if a monthly batch is to be processed every 1 month, set it to 90 days, which is approximately 3 times that period. to capture changes in the aggressive freeze for tuples over the duration of the cycle.

- Minimum value: 1d
- Maximum value: 180d

If an out-of-range value is specified, 90d is set.

If this parameter is omitted, 90d is assumed.

Appendix B System Administration Functions

This appendix describes the system administration functions of Fujitsu Enterprise Postgres.



Refer to "System Administration Functions" under "The SQL Language" in the PostgreSQL Documentation for information on other system administration functions.

B.1 WAL Mirroring Control Functions

The following table lists the functions that can be used for backup and recovery based on WAL mirroring.

Table B.1 WAL mirroring control functions

Name	Return type	Description
<code>pgx_pause_wal_multiplexing()</code>	void	Stops WAL multiplexing
<code>pgx_resume_wal_multiplexing()</code>	void	Resumes WAL multiplexing
<code>pgx_is_wal_multiplexing_paused()</code>	boolean	Returns true if WAL multiplexing has stopped

If WAL multiplexing has not been configured, these functions return an error. Setting the `backup_destination` parameter in `postgresql.conf` configures WAL multiplexing.

Only superusers can execute these functions.

B.2 Transparent Data Encryption Control Functions

The following table lists the functions that can be used for transparent data encryption.

Table B.2 Transparent data encryption control functions

Name	Return type	Description
<code>pgx_open_keystore(<i>passphrase</i>)</code> <code>pgx_open_keystore(sslpassphrase => text)</code>	void	Opens the keystore
<code>pgx_set_master_key(<i>passphrase</i>)</code>	void	Sets the master encryption key
<code>pgx_declare_external_master_key(kms_name => text, key_id => text, sslpassphrase => text)</code>	void	To set an encryption key existing in a key management system as a master encryption key for transparent data encryption.
<code>pgx_set_keystore_passphrase(<i>oldPassphrase</i>, <i>newPassphrase</i>)</code>	void	Changes the keystore passphrase

B.2.1 `pgx_open_keystore`

`pgx_open_keystore` opens the keystore.

Only superusers can execute this function. Also, this function cannot be executed within a transaction block.

File-based keystores:

The `pgx_open_keystore` function uses the specified passphrase to open the keystore. When the keystore is opened, the master encryption key is loaded into the database server memory. In this way, you can access the encrypted data and create encrypted tablespaces. If the keystore is already open, this function returns an error.

Using the key management system as a keystore

`pgx_open_keystore` makes available (opens a keystore) a master encryption key on a key management system that has already been declared for use. The keystore cannot be opened unless it has been declared to use a master encryption key.

If the keystore is already open, use the credentials you entered to reconnect to the key management system.

Specify the authentication information for connecting to the key management system. Arguments must be specified in naming notation. The information you pass in the argument depends on the key management system you use.

If the key management system information file specifies an obfuscated credentials file, the file is recreated with the new credentials.

Using the key management service of type kmip

The following arguments are specified in naming notation.

- `sslpassphrase` text

Specifies the passphrase of the client certificate private key file when connecting to the KMIP server. This can be omitted if no passphrase is set in the private key file.

Using the key management service of type custom

The following arguments are specified in naming notation.

- `kms_secret` text

Confidential information passed to the plugin. It can be omitted if it is not necessary for using the key management system. Whether or not it can be omitted depends on the implementation of the plugin.

Example

To specify the passphrase `mykmippassphrase` for the client certificate private key file in naming notation:

```
SELECT pgx_open_keystore( sslpassphrase => 'mykmippassphrase' );
```

B.2.2 `pgx_set_master_key`

The `pgx_set_master_key` function generates a master encryption key and stores it in the file-based keystore.

If the keystore does not exist, this function creates a keystore. If the keystore already exists, this function modifies the master encryption key. If the keystore has not been opened, this function opens it.

The passphrase is a string of 8 to 200 bytes.

Only superusers can execute this function. Also, this function cannot be executed within a transaction block. Processing is not affected by whether the keystore is open.

B.2.3 `pgx_declare_external_master_key`

`pgx_declare_external_master_key` declares the use of an encryption key that exists in the key management system as the master encryption key for transparent data encryption. If the master encryption key already exists, change the master encryption key. If the master encryption key already exists, the keystore must be open.

The argument specifies information that identifies the master encryption key. Arguments must be specified in naming notation. The information you pass in the argument depends on the key management system you use.

This function can only be executed by superuser. Also, you cannot execute this function within a transaction block.

This function is available if you have installed the extension `'tde_kms'`.

The following arguments are specified in naming notation:

- `kms_name` text

Specify the key management system name specified in the key management system connection information file. Required.

- key_id text

Specify the key ID assigned to the encryption key. Cannot be omitted.

- sslpassphrase text

Specify the passphrase of the client certificate private key file when connecting to the KMIP server. This can be omitted if the private key file does not have a passphrase. Ignored if the key management system type specified by kms_name is not kmip.

- kms_secret text

Confidential information passed to the plugin. It can be omitted if it is not necessary for using the key management system. Whether or not it can be omitted depends on the implementation of the plugin. Ignored if the key management system type specified by kms_name is not custom.

Example

```
SELECT pgx_declare_external_master_key( kms_name => 'mykmipsvr', key_id =>
'a0eebc99-9c0b-0000-0000-000000000000', sslpassphrase => 'mykmippassphrase' );
```

B.2.4 pgx_set_keystore_passphrase

The pgx_set_keystore_passphrase function changes the file-based 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 Profile Management Functions and User Management Functions

The table below lists the profile management functions and user management functions used in policy-based login security.

B.3.1 Profile Management Functions

Name	Return type	Description
pgx_create_profile(profile_name name, password_parameter json)	void	<p>Create a new profile.</p> <p>For profile_name, specify the profile name. An error will occur if an existing profile name is specified.</p> <p>For password_parameter, specify parameters and values in key-value format as follows.</p> <pre>'{ "INACTIVE_USER_TIME": 15, "PASSWORD_LIFE_TIME": 30, "PASSWORD_GRACE_TIME": "UNLIMITED", "PASSWORD_REUSE_TIME": 10, "PASSWORD_REUSE_MAX": 5, "PASSWORD_LOCK_TIME": 0.5, "FAILED_LOGIN_ATTEMPTS": "DEFAULT", "PASSWORD_ALLOW_HASHED": true, "PASSWORD_ROLLOVER_TIME": 0.125 }');</pre> <p>The json value accepts integer, numeric, boolean depending on the parameter, but only accepts strings "DEFAULT" and "UNLIMITED" as special values.</p>

Name	Return type	Description
		For parameters that are omitted in json or that specify null for value, they follow the values in the default profile. Also, password_parameter can be omitted, and if omitted, a profile with all parameters conforming to the default profile will be created.
pgx_alter_profile(profile_name name, alter_parameter json)	void	<p>Update the contents of an existing profile.</p> <p>For profile_name, specify the name of the profile to update. If you specify a profile name that does not exist, an error will occur.</p> <p>In alter_parameter, specify the variable name and value you want to change in key-value format as follows. Does not change the profile contents for parameters that are omitted or for which value is null.</p> <pre>{ "name": "new_name", "PASSWORD_LIFE_TIME": 50, "PASSWORD_GRACE_TIME": 10 }</pre> <p>name: Specify the modified profile name. null cannot be specified.</p> <p>If the default profile is specified in profile_name, the following changes cannot be made.</p> <ul style="list-style-type: none"> - Change name - Change each password_parameter to "DEFAULT"
pgx_drop_profile(profile_name name, if_exists boolean, cascade boolean)	void	<p>Delete an existing profile.</p> <p>Specify the profile name to be deleted in profile_name.</p> <p>Specifying the default profile will result in an error.</p> <p>if_exists specifies whether an error occurs when a nonexistent profile name is specified in profile_name. If true, no error. if_exists is optional. Default value is false.</p> <p>cascade specifies whether an error occurs when a profile assigned to a user is specified in profile_name. If false, an error will occur. If true, unassign all of the profiles before deleting them, and any unassigned users will be assigned the default profile instead. cascade is optional. Default value is false.</p> <p>[Example when specifying true] Specify arguments explicitly. (Example of named notation)</p> <pre>SELECT pgx_drop_profile('test_profile', cascade => true);</pre>

B.3.2 User Management Functions

Name	Return type	Description
pgx_assign_profile_to_user(user_name name, profile_name name)	void	<p>Assigns an existing profile to a user.</p> <p>Specify the assignee user in user_name. If you specify a user name that does not exist, an error will occur.</p> <p>Specify the profile name to be assigned in profile_name. If you specify a profile name that does not exist, an error will occur.</p>
pgx_lock_user(user_name name)	void	Lock the user explicitly.

Name	Return type	Description
		Specify the user name in user_name. If you specify a user name that does not exist, an error will occur. If locked by this function, the lock period is indefinite. If you want to unlock it, you need to unlock it with the pgx_unlock_user function. If you lock a user that is already locked, the lock state is overwritten and becomes an indefinite lock even if it was previously a finite lock.
pgx_unlock_user(user_name name)	void	Explicitly unlock the user. The type of lock (profile reference lock or permanent lock) does not matter. Specify the user name in user_name. If you specify a user name that does not exist, an error will occur. Specifying an unlocked user does not result in an error.
pgx_make_password_expire(user_name name, expireat timestampz)	void	Expires the specified user's password immediately. Alternatively, specify the time to expire. Specify the user name in user_name. If you specify a user name that does not exist, an error will occur. expireat can be a time that expires.expireat is optional. If omitted, execution time is specified.
pgx_make_password_rollover_expire(user_name name)	void	Expires immediately the period during which the specified user's old and new passwords can be used together so that the user cannot log in with the old password. Specify the user name in user_name. If you specify a user name that does not exist, an error will occur. If this command is executed outside the period for using the old and new passwords, it is ignored.

B.4 Data Masking Control Functions

The table below lists the functions that can be used for data masking.

Table B.3 Data masking control functions

Name	Return type	Description
pgx_alter_confidential_policy	boolean	Changes masking policies
pgx_create_confidential_policy	boolean	Creates masking policies
pgx_drop_confidential_policy	boolean	Deletes masking policies
pgx_enable_confidential_policy	boolean	Enables or disables masking policies
pgx_update_confidential_values	boolean	Changes replacement characters when full masking is specified for masking type

B.4.1 [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_COLUMN'
	column_name	varchar(63)	Masking target name	Mandatory
	function_type	varchar(63)	Masking type - 'FULL': Full masking - 'PARTIAL': Partial masking - 'REGEXP': Regular expression masking	'FULL'
Partial masking	function_parameters	varchar(1024)	Masking format for partial masking	Mandatory
Regular expression masking	regexp_pattern	varchar(1024)	Search pattern for regular expression masking	Mandatory
	regexp_replacement	varchar(1024)	Replacement character/string for regular expression masking	Mandatory
	regexp_flags	varchar(20)	Regular expression flags	NULL

- Delete a masking target from a masking policy

Masking type for which an argument can be specified	Argument	Data type	Description	Default value
All	action	varchar(63)	'DROP_COLUMN'	Mandatory
	column_name	varchar(63)	Masking target name	Mandatory

- Change the masking condition

Masking type for which an argument can be specified	Argument	Data type	Description	Default value
All	action	varchar(63)	'MODIFY_EXPRESSION'	Mandatory
	expression	varchar(1024)	Masking condition to be changed	Mandatory

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

Masking type for which an argument can be specified	Argument	Data type	Description	Default value
All	action	varchar(63)	'MODIFY_COLUMN'	Mandatory
	column_name	varchar(63)	Masking target name	Mandatory
	function_type	varchar(63)	Masking type - 'FULL': Full masking - 'PARTIAL': Partial masking - 'REGEXP': Regular expression masking	'FULL'
Partial masking	function_parameters	varchar(1024)	Masking format for partial masking	Mandatory
Regular expression masking	regexp_pattern	varchar(1024)	Search pattern for regular expression masking	Mandatory
	regexp_replacement	varchar(1024)	Replacement character/string for regular expression masking	Mandatory
	regexp_flags	varchar(20)	Regular expression flags	NULL

- Change the masking policy description

Masking type for which an argument can be specified	Argument	Data type	Description	Default value
All	action	varchar(63)	'SET_POLICY_DESCRIPTION'	Mandatory
	policy_description	varchar(1024)	Masking policy description	Mandatory

- Change the masking target description

Masking type for which an argument can be specified	Argument	Data type	Description	Default value
All	action	varchar(63)	'SET_COLUMN_DESCRIPTION'	Mandatory
	column_name	varchar(63)	Masking target name	Mandatory
	column_description	varchar(1024)	Masking target description	Mandatory

Details about whether arguments can be omitted are as follows.

Argument	Mandatory or optional									
	ADD_COLUMN			DROP_COLUMN	MODIFY_EXPRESSION	MODIFY_COLUMN			SET_POLICY_DESCRIPTION	SET_COLUMN_DESCRIPTION
	Full masking	Partial masking	Regular expression masking			Full masking	Partial masking	Regular expression masking		
schema_name	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
table_name	N	N	N	N	N	N	N	N	N	N
policy_name	N	N	N	N	N	N	N	N	N	N
action	Y	Y	Y	N	N	N	N	N	N	N
column_name	N	N	N	N	-	N	N	N	-	N
function_type	Y	N	N	-	-	Y	N	N	-	-
expression	-	-	-	-	N	-	-	-	-	-
policy_description	-	-	-	-	-	-	-	-	N	-
column_description	-	-	-	-	-	-	-	-	-	N
function_parameters	-	N	-	-	-	-	N	-	-	-
regexp_pattern	-	-	N	-	-	-	-	N	-	-
regexp_replacement	-	-	N	-	-	-	-	N	-	-
regexp_flags	-	-	Y	-	-	-	-	Y	-	-

Y: Can be omitted; N: Cannot be omitted; -: Ignored when specified

Return value

Return value	Description
TRUE	Ended normally
FALSE	Ended abnormally

Execution example 1

Adding masking policy p1 to masking target c2

```
postgres=# select pgx_alter_confidential_policy(table_name := 't1', policy_name := 'p1', action :=
'ADD_COLUMN', column_name := 'c2', function_type := 'PARTIAL', function_parameters := 'VVVFVVVFVVVV,
VVV-VVVV-VVVV, *, 4, 11');
pgx_alter_confidential_policy
-----
```

```
t
(1 row)
```

Execution example 2

Deleting masking target c1 from masking policy p1

```
postgres=# select pgx_alter_confidential_policy(table_name := 't1', policy_name := 'p1', action :=
'DROP_COLUMN', column_name := 'c1');
pgx_alter_confidential_policy
-----
t
(1 row)
```

Execution example 3

Changing the masking condition for masking policy p1

```
postgres=# select pgx_alter_confidential_policy(table_name := 't1', policy_name := 'p1', action :=
'MODIFY_EXPRESSION', expression := 'false');
pgx_alter_confidential_policy
-----
t
(1 row)
```

Execution example 4

Changing the content of masking policy p1 set for masking target c2

```
postgres=# select pgx_alter_confidential_policy(table_name := 't1', policy_name := 'p1', action :=
'MODIFY_COLUMN', column_name := 'c2', function_type := 'FULL');
pgx_alter_confidential_policy
-----
t
(1 row)
```

Execution example 5

Changing the description of masking policy p1

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

Execution example 6

Changing the description of masking target c2

```
postgres=# select pgx_alter_confidential_policy(table_name := 't1', policy_name := 'p1', action :=
'SET_COLUMN_DESCRIPTION', column_name := 'c2', column_description := 'c2 column is FULL. ');
pgx_alter_confidential_policy
-----
t
(1 row)
```

Description

- The arguments for the `pgx_alter_confidential_policy` system management function can be specified in any order.

- The action parameters below can be specified. When action parameters are omitted, ADD_COLUMN is applied.

Parameter	Description
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.4.2 pgx_create_confidential_policy

Description

Creates masking policies

Format

The format varies depending on the masking type. The format is shown below.

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

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

```

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

```

Argument

Details are as follows.

Masking type for which an argument can be specified	Argument	Data type	Description	Default value
All	schema_name	vchar(63)	Schema name of table for which the masking policy is created	'public'
	table_name	vchar(63)	Name of table for which the masking policy is created	Mandatory
	policy_name	vchar(63)	Masking policy name	Mandatory
	expression	vchar(1024)	Masking condition	Mandatory
	enable	boolean	Masking policy status - 't': Enabled - 'f': Disabled	't'
	policy_description	vchar(1024)	Masking policy description	NULL
	column_name	vchar(63)	Masking target name	NULL
	function_type	vchar(63)	Masking type - 'FULL': Full masking - 'PARTIAL': Partial masking - 'REGEXP': Regular expression masking	'FULL'
	column_description	vchar(1024)	Masking target description	NULL
Partial masking	function_parameters	vchar(1024)	Masking format for partial masking	Mandatory
Regular expression masking	regexp_pattern	vchar(1024)	Search pattern for regular expression masking	Mandatory
	regexp_replacement	vchar(1024)	Replacement character/string for regular expression masking	Mandatory
	regexp_flags	vchar(20)	Regular expression flags	NULL

Details about whether arguments can be omitted are as follows.

Argument	Mandatory or optional		
	Full masking	Partial masking	Regular expression masking
schema_name	Y	Y	Y
table_name	N	N	N
policy_name	N	N	N
expression	N	N	N
enable	Y	Y	Y
policy_description	Y	Y	Y

Argument	Mandatory or optional		
	Full masking	Partial masking	Regular expression masking
column_name	Y	Y	Y
function_type	Y	Y	Y
column_description	Y	Y	Y
function_parameters	-	N	-
regexp_pattern	-	-	N
regexp_replacement	-	-	N
regexp_flags	-	-	Y

Y: Can be omitted; N: Cannot be omitted; -: Ignored when specified

Return value

Return value	Description
TRUE	Ended normally
FALSE	Ended abnormally

Execution example 1

Creating masking policy p1 that does not contain a masking target

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

Execution example 2

Creating masking policy p1 that contains masking target c1 of which the masking type is full masking

```
postgres=# select pgx_create_confidential_policy(schema_name := 'public', table_name := 't1',
policy_name := 'p1', expression := 'l=1', enable := 't', policy_description := 'this policy is an
example.', column_name := 'c1', function_type := 'FULL', column_description := 'c1 column is FULL.');
```

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

Execution example 3

Creating masking policy p1 that contains masking target c2 of which the masking type is partial masking

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

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

Execution example 4

Creating masking policy p1 that contains masking target c3 of which the masking type is regular expression masking

```
postgres=# select pgx_create_confidential_policy( table_name := 't1', policy_name := 'p1',
expression := 'l=1', column_name := 'c3', function_type := 'REGEXP', regexp_pattern := '(.*)(@.*)',
regexp_replacement := 'xxx\2', regexp_flags := 'g');
 pgx_create_confidential_policy
-----
 t
(1 row)
```

Description

- The arguments for the `pgx_create_confidential_policy` system management function can be specified in any order.
- If `column_name` is omitted, only masking policies that do not contain masking target will be created.
- One masking policy can be created for each table. Use the `pgx_alter_confidential_policy` system management function to add a masking target to a masking policy.
- The `function_parameters` argument is enabled when the `function_type` is `PARTIAL`. If the `function_type` is other than `PARTIAL`, it will be ignored.
- The arguments below are enabled when the `function_type` is `REGEXP`. If the `function_type` is other than `REGEXP`, these arguments will be ignored.
 - `regexp_pattern`
 - `regexp_replacement`
 - `regexp_flags`
- If a table for which a masking policy is to be applied is deleted, delete the masking policy as well.



See

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

B.4.3 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
<code>schema_name</code>	<code>varchar(63)</code>	Schema name of table for which a masking policy is deleted	'public'

Argument	Data type	Description	Default value
table_name	varchar(63)	Name of table for which a masking policy is deleted	Mandatory
policy_name	varchar(63)	Masking policy name	Mandatory

Details about whether arguments can be omitted are as follows.

Argument	Mandatory or optional
schema_name	Y
table_name	N
policy_name	N

Y: Can be omitted; N: Cannot be omitted

Return value

Return value	Description
TRUE	Ended normally
FALSE	Ended abnormally

Execution example

Deleting masking policy p1

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

Description

- The arguments for the `pgx_drop_confidential_policy` system management function can be specified in any order.
- If a table for which a masking policy is to be applied is deleted, delete the masking policy as well.



See

Refer to "String Constants" in the PostgreSQL Documentation for information on the strings to specify for arguments.

B.4.4 `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 - 't': Enabled - 'f': Disabled	Mandatory

Details about whether arguments can be omitted are as follows.

Argument	Mandatory or optional
schema_name	Y
table_name	N
policy_name	N
enable	N

Y: Can be omitted; N: Cannot be omitted

Return value

Return value	Description
TRUE	Ended normally
FALSE	Ended abnormally

Execution example

Enabling masking policy p1

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

Description

The arguments for the `pgx_enable_confidential_policy` system management function can be specified in any order.



Refer to "String Constants" in the PostgreSQL Documentation for information on the strings to specify for arguments.

B.4.5 `pgx_update_confidential_values`

Description

Changes replacement characters when full masking is specified for masking type

Format

```
pgx_update_confidential_values(  
  [number_value := 'numberValue']  
  [, char_value := 'charValue']  
  [, varchar_value := 'varcharValue']  
  [, date_value := 'dateValue']  
  [, ts_value := 'tsValue']  
)
```

Argument

Details are as follows.

Argument	Data type	Description
number_value	integer	Replacement character in numeric type
char_value	varchar(1)	Replacement character in char type
varchar_value	varchar(1)	Replacement character in varchar type
date_value	date	Replacement character in date type
ts_value	timestamp	Replacement character in timestamp type

Return value

Return value	Description
TRUE	Ended normally
FALSE	Ended abnormally

Execution example

Using '*' as a replacement character in char type and varchar type

```
postgres=# select pgx_update_confidential_values(char_value := '*', varchar_value := '*');  
pgx_update_confidential_values  
-----  
t  
(1 row)
```

Description

- The arguments for the `pgx_update_confidential_values` system management function can be specified in any order.
- Specify one or more arguments for the `pgx_update_confidential_values` system management function. A replacement character is not changed for an omitted argument.



See

Refer to "String Constants" in the PostgreSQL Documentation for information on the strings to specify for arguments.

B.5 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.6 High-Speed Data Load Control Functions

The table below lists the functions that can be used for high-speed data load.

Table B.5 High-speed data load control functions

Name	Return type	Description
pgx_loader	bigint	Creates dynamic shared memory, starts parallel workers and loads data
pgx_loader_recovery	smallint	Resolves in-doubt transactions

The pgx_loader command executes the above functions internally.

Appendix C System Catalogs

Describes the system catalog of Fujitsu Enterprise Postgres.



See

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

C.1 pgx_profile

A system catalog for managing profile information. -1 is stored if DEFAULT is specified for each parameter in the profile, and -2 is stored if UNLIMITED is specified.

Column	Type	Description
oid	oid	Profile identifier (Primary key constraint)
prfname	name	Profile name (Unique constraint, NOT NULL constraint)
prfpasswordlifetime	integer	Value of PASSWORD_LIFE_TIME (seconds) (NOT NULL constraint)
prfpasswordgracetime	integer	Value of PASSWORD_GRACE_TIME (seconds) (NOT NULL constraint)
prfpasswordreusetime	integer	Value of PASSWORD_REUSE_TIME (seconds) (NOT NULL constraint)
prfpasswordreusemax	integer	Value of PASSWORD_REUSE_MAX (seconds) (NOT NULL constraint)
prfpasswordlocktime	integer	Value of PASSWORD_LOCK_TIME (seconds) (NOT NULL constraint)
prffailedloginattempts	integer	Value of FAILED_LOGIN_ATTEMPTS (seconds) (NOT NULL constraint)
prfpasswordallowhashed	integer	Value of PASSWORD_ALLOW_HASHED (NOT NULL constraint)
prfinactiveusertime	integer	Value of INACTIVE_USER_TIME (seconds) (NOT NULL constraint)
prfpasswordrollovertime	integer	Value of PASSWORD_ROLLOVER_TIME (seconds) (NOT NULL constraint)

C.2 pgx_user_profile

A system catalog that manages profile-related information associated with users. When DROP ROLE, the corresponding row is automatically deleted.

Column	Type	References	Description
userprfroleid	oid	pg_authid.oid	User's identifier (Primary key constraint)
userprfprfid	oid	pgx_profile.oid	Identifier of the profile assigned to the user
userprfaccountlock	smallint		User lock state (NOT NULL constraint) 0: Not locked

Column	Type	References	Description
			1: Lock (Refer to PASSWORD_LOCK_TIME in profile for duration) 2: Lock (Unlimited)
userprfpasswordstatus	char		Password status Updated if there are any changes at login (NOT NULL constraint) o: Current valid password g: In grace period e: Expired
userprflockdate	timestamp with time zone		Time the user was locked
userprfpasswordsetat	timestamp with time zone		Time you updated your current password
userprfpasswordexpire	timestamp with time zone		Password expiration time
userprflastactivetime	timestamp with time zone		Time the user's session was last checked
userprfpasswordrolloverexpire	timestamp with time zone		Time at which the period in which old and new passwords can be used ends

C.3 pgx_auth_password

The system catalog for storing user-associated password information. Used for password rollover.

Since this catalog contains passwords, it must not be readable by third parties.

Column	Type	References	Description
authpwdroleid	oid	pg_authid.oid	User's identifier (Primary key constraint)
authpwdoldpassword	text		Hashed password Used for authentication when password rollover is enabled.
authpwdoldpasswordsetat	timestamp with time zone		Time at which authpwdoldpassword was set
authpwdnewpassword	text		Hashed password Used for authentication when password rollover is enabled.
authpwdnewpasswordsetat	timestamp with time zone		Time at which authpwdnewpassword was set

C.4 pgx_password_history

A system catalog for managing password history. Used for password reuse confirmation and password rollover. When DROP ROLE is executed, the corresponding row is automatically deleted.

Since this catalog contains passwords, it must not be readable by third parties.

This system catalog is updated when the password is updated, and if the password reuse limit is finite, unnecessary history deletion is also performed. Old passwords that have been removed from the catalog are treated as having never been used.

passhistpassword is saved in the format specified by the password_encryption parameter.

Column	Type	Description
passhistroleid	oid	Identifier of the user whose password was set Set unique constraint with (passhistroleid, passhistpassword).
passhistpassword	text	Encrypted password Save rolpassword value in pg_authid as history (NOT NULL constraint)
passhistpasswordsetat	timestamp with time zone	Time when the password was updated (NOT NULL constraint)

Appendix D 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.

D.1 pgx_tablespaces

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

Table D.1 `pgx_tablespaces` view

Column	Type	References	Description
<code>spctablespace</code>	<code>oid</code>	<code>pg_tablespace.oid</code>	Tablespace OID
<code>spcencalgo</code>	<code>text</code>		Tablespace encryption algorithm

The `spcencalgo` string displays one of the following values:

- none: Tablespace is not encrypted
- AES128: AES with key length of 128 bits
- AES256: AES with key length of 256 bits

D.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 D.2 `pgx_stat_lwlock` view

Column	Type	Description
<code>lwlock_name</code>	<code>name</code>	Name of the lightweight lock
<code>total_waits</code>	<code>bigint</code>	Number of waits caused by the lightweight lock
<code>total_wait_time</code>	<code>double precision</code>	Number of milliseconds spent in waits caused by the lightweight lock
<code>stats_reset</code>	<code>timestamp with timezone</code>	Last time at which this statistics was reset

D.3 pgx_stat_latch

The `pgx_stat_latch` view displays statistics related to latches, with each type of wait information within Fujitsu Enterprise Postgres displayed on a separate line.

Table D.3 `pgx_stat_latch` view

Column	Type	Description
<code>latch_name</code>	<code>name</code>	Name of the latch
<code>total_waits</code>	<code>bigint</code>	Number of waits caused a wait
<code>total_wait_time</code>	<code>double precision</code>	Number of milliseconds spent in waits caused by the latch
<code>stats_reset</code>	<code>timestamp with timezone</code>	Last time at which this statistic was reset

D.4 pgx_stat_walwriter

The `pgx_stat_walwriter` view displays statistics related to WAL writing, in a single line.

Table D.4 `pgx_stat_walwriter` view

Column	Type	Description
<code>dirty_writes</code>	<code>bigint</code>	Number of times old WAL buffers were written to the disk because the WAL buffer was full when WAL records were added
<code>writes</code>	<code>bigint</code>	Number of WAL writes
<code>write_blocks</code>	<code>bigint</code>	Number of WAL write blocks
<code>total_write_time</code>	<code>double precision</code>	Number of milliseconds spent on WAL writing
<code>stats_reset</code>	<code>timestamp with timezone</code>	Last time at which this statistic was reset

D.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 D.5 `pgx_stat_sql` view

Column	Type	Description
<code>selects</code>	<code>bigint</code>	Number of SELECT statements executed In database multiplexing mode, this number includes the SELECT statements executed in Mirroring Controller. Mirroring Controller executes the SELECT statement using the interval specified for the <code>heartbeat_interval</code> of the server definition file (milliseconds).
<code>inserts</code>	<code>bigint</code>	Number of INSERT statements executed
<code>deletes</code>	<code>bigint</code>	Number of DELETE statements executed
<code>updates</code>	<code>bigint</code>	Number of UPDATE statements executed
<code>selects_with_parallelism</code>	<code>bigint</code>	Number of times parallel scan was used in SELECT statements
<code>inserts_with_parallelism</code>	<code>bigint</code>	Not used
<code>deletes_with_parallelism</code>	<code>bigint</code>	Not used
<code>updates_with_parallelism</code>	<code>bigint</code>	Not used
<code>copies_with_parallelism</code>	<code>bigint</code>	Not used
<code>declares</code>	<code>bigint</code>	Number of DECLARE statements executed (number of cursor OPENS)
<code>fetches</code>	<code>bigint</code>	Number of FETCH statements executed
<code>checkpoints</code>	<code>bigint</code>	Number of CHECKPOINT statements executed
<code>clusters</code>	<code>bigint</code>	Number of CLUSTER statements executed
<code>copies</code>	<code>bigint</code>	Number of COPY statements executed
<code>reindexes</code>	<code>bigint</code>	Number of REINDEX statements executed
<code>truncates</code>	<code>bigint</code>	Number of TRUNCATE statements executed
<code>locks</code>	<code>bigint</code>	Number of times a lock occurred
<code>stats_reset</code>	<code>timestamp with timezone</code>	Last time at which this statistic was reset

D.6 pgx_stat_gmc

The `pgx_stat_gmc` view provides information about the GMC areas.

Table D.6 `pgx_stat_gmc` view

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

D.7 pgx_stat_progress_loader

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

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

Table D.7 `pgx_stat_progress_loader` view

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

Appendix E Tables Used by Transparent Data Encryption

This appendix explains tables used by the transparent data encryption feature.

E.1 pgx_tde_master_key

Provides information about the master encryption key being used when using the key management system as a keystore.

Column	Type	Description
local_key_id	integer	Sequence of encryption keys used internally by Fujitsu Enterprise Postgres
kms_name	text	Key management system name
key_name	text	Not currently used
key_id	text	Key ID in key management system
start_time	timestamp	Time the key was activated
status	enum	in-use: Current master encryption key in use used: Master encryption key used in the past scheduled: A new encryption key that was requested to update the master encryption key and is to be used.

Execution example

```
postgres=# select * from pgx_tde_master_key;
local_key_id | kms_name | key_name | key_id | start_time | status
-----+-----+-----+-----+-----+-----
1 | mykmipsvr | | a0eebc99-9c0b-0000-0000-000000000000 | 2022-12-16 05:43:49 | in-use
(1 row)
```

Information

Do not use queries with "*" in the selection list, as the order of the columns may change or columns may be added.

Appendix F Tables Used by Data Masking

This appendix explains tables used by the data masking feature.

Information

These tables are updated by the data masking control function, so do not use SQL statements to directly update these tables.

F.1 pgx_confidential_columns

This table provides information on masking target for which masking policies are set.

Column	Type	Description
schema_name	varchar(63)	Schema name of table for which a masking policy is applied
table_name	varchar(63)	Name of table for which a masking policy is applied
policy_name	varchar(63)	Masking policy name
column_name	varchar(63)	Masking target name
function_type	varchar(63)	Masking type <ul style="list-style-type: none"> - 'FULL': Full masking - 'PARTIAL': Partial masking - 'REGEXP': Regular expression masking
function_parameters	varchar(1024)	Masking format for partial masking
regexp_pattern	varchar(1024)	Search pattern for regular expression masking
regexp_replacement	varchar(1024)	Replacement character/string for regular expression masking
regexp_flags	varchar(20)	Regular expression flags
column_description	varchar(1024)	Masking target description

Execution example

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

F.2 pgx_confidential_policies

This table provides information on masking policies.

Column	Type	Description
schema_name	varchar(63)	Schema name of table for which a masking policy is applied

Column	Type	Description
table_name	varchar(63)	Name of table for which a masking policy is applied
policy_name	varchar(63)	Masking policy name
expression	varchar(1024)	Masking condition
enable	boolean	Masking policy status - 't': Enabled - 'f': Disabled
policy_description	varchar(1024)	Masking policy description

Execution example

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

F.3 pgx_confidential_values

This table provides information on replacement characters when full masking is specified for masking type.

Column	Data type	Description	Default value
number_value	integer	Numeric	0
char_value	varchar(1)	char type	Spaces
varchar_value	varchar(1)	varchar type	Spaces
date_value	date	date type	'1970-01-01'
timestamp_value	timestamp	timestamp type	'1970-01-01 00:00:00'

Execution example

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

Appendix G Tables Used by Aggressive Freeze for Tuples

This appendix explains tables used by the aggressive freeze for tuples feature.

G.1 `pgx_stat_freeze_results`

The table `pgx_stat_freeze_results` logs statistics about aggressive freeze for tuples. If an aggressive freeze for tuples is performed, that information will be added as an accumulation. The cumulative information is summarized for the time period specified by the `pgx_stat_vacuum_freeze.monitoring_interval` parameter, and the information for the time period specified by the `pgx_stat_vacuum_freeze.log_retention_period` parameter is kept. For the information about the `pgx_stat_vacuum_freeze.monitoring_interval` and `pgx_stat_vacuum_freeze.log_retention_period` parameters, refer to "[Parameters for scheduling aggressive freeze for tuples](#)".

You can gather information to determine how long the aggressive freeze for tuples process is taking, and help estimate the aggressive freeze for tuples schedule.

Column	Type	Description
<code>measurement_start_time</code>	timestamp with time zone	Start time of the measurement
<code>measurement_end_time</code>	timestamp with time zone	End time of the measurement
<code>vacuum_count</code>	bigint	Number of times the freeze operation was performed
<code>autovacuum</code>	boolean	true for autovacuum driven vacuum
<code>scan_pages</code>	bigint	Accumulation of pages scanned during freezing
<code>frozen_pages</code>	bigint	Accumulation of frozen pages
<code>frozen_tuples</code>	bigint	Accumulation of frozen tuples
<code>total_frozen_wal_records</code>	bigint	Accumulation of the number of WAL records produced by freezing
<code>frozen_full_page_image_wal_records</code>	bigint	Accumulation of the number of WAL full-page images produced by freezing
<code>wal_total_bytes</code>	numeric	Accumulation, in bytes, of WAL generated by freezing
<code>total_elapsed_time</code>	double precision	Accumulation of freezing time. Units are milliseconds
<code>sleep_times</code>	double precision	Accumulation of time spent sleeping during freezing. Units are milliseconds

Appendix H Tables Used by High-Speed Data Load

This appendix describes the tables used by high-speed data load.

H.1 pgx_loader_state

The pgx_loader_state table provides information about transactions prepared by high-speed data load.

Column	Type	Description
id	serial	Unique identifier. This value is assigned from the pgx_loader_state_id_seq sequence.
gid	text	Global transaction identifier assigned to a transaction.
state	text	State of the transaction. The value can be one of the following: <ul style="list-style-type: none">- commit: The prepared transaction has been committed.- rollback: The prepared transaction is in in-doubt state.
leader_pid	integer	Process ID of the backend process (leader process) that executed the pgx_loader control function.
role_oid	integer	Role identifier (OID). A prepared transaction can only be completed by the same user who executed the original transaction or by a superuser.
relation_oid	integer	Object identifier (OID).

Information

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 I Starting and Stopping the Web Server Feature of WebAdmin

To use WebAdmin for creating and managing a Fujitsu Enterprise Postgres instance on a server where Fujitsu Enterprise Postgres is installed, you must first start the Web server feature of WebAdmin.

- Using WebAdmin in a single-server configuration

You must start the Web server on the server on which Fujitsu Enterprise Postgres and WebAdmin are installed.

- Using WebAdmin in a multiserver configuration

You must start the Web server on all servers on which WebAdmin has been installed.

This appendix describes how to start and stop the Web server feature of WebAdmin.

Note that "<x>" in paths indicates the product version.



Refer to "Installing WebAdmin in a Multiserver Configuration" in the Installation and Setup Guide for Server for information on multiserver installation.

I.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
```

I.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 J 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.

J.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.

Credential name, User name and Password should not contain hazardous characters. Refer to “[Appendix K WebAdmin Disallow User Inputs Containing Hazardous Characters](#)”.

- [Credential name]: Name of the credential

The name must meet the conditions below:

- Maximum of 16 characters
- The first character must be an ASCII alphabetic character
- The other characters must be ASCII alphanumeric characters

- [User name]: The operating system user name or database instance user name that will be used later

- [Password]: Password for the user

- [Confirm password]: Reenter the password.

3. Click  to store the credential.

J.2 Using a Credential

Once a credential is created in the Wallet, it can be used during remote instance creation or standby instance creation.

If you select the credential saved in "J.1 Creating a Credential" in [Operating system credential], the user name and password will be automatically populated.

Appendix K 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 L Copy Command Samples that Use the Advanced Copy Feature of the ETERNUS Disk Array

Backup/recovery scripts that use OPC, an advanced copy feature of the Fujitsu Storage ETERNUS disk array (hereafter referred to as ETERNUS disk array), are supplied as copy command samples for use by the `pgx_dmpall` and `pgx_rcvall` commands.

Users can copy the samples to any file and make changes appropriate to their environment or operations.

The samples are stored in the directories below:

- Basic version

```
/installDir/share/copy_command.esf_acm1.sh.sample
```

- Advanced version

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



The samples use the replication management command of the Fujitsu Storage ETERNUS SF AdvancedCopy Manager (hereafter referred to as ACM) for operating the advanced copy features of the ETERNUS disk array. Refer to the relevant manual for details.

Prerequisite for using the samples

To use the samples, it is necessary to configure the settings of the advanced copy features of the ETERNUS disk array in advance. The samples assume that the replication source/volume has been configured, its contents have been physically copied to the replication volume, and that the tracking process has been started.

The samples cannot be used on SLES 15. The samples use the `sudo` command of the operating system so that operation of the commands and file systems of ACM is performed by the superuser of the operating system. Therefore, determine if these implementations satisfy the security standards on the database server, and if necessary, perform implementations using other means.

Sample content (advanced version)

The sample uses two replication volumes as the replication source on which the database cluster and tablespace are located, alternating between them for each backup. It then registers to the backup information file the replication volume (group) used for the latest backup.

In addition, the backup information file is also used to determine the replication volume to restore during recovery.

The processing for each operation mode is described below:

prepare mode

1. The backup information file is read, and the replication volume to be used for the current backup is determined.
2. The completion status of physical copies to all replication volumes is checked using the `swsrpstat` command of ACM.
3. The replication volume determined in step 1 is written to a temporary file for later use by the backup mode.

backup mode

1. The temporary file is read, and the replication volume information is retrieved.
2. The file system buffer is written using the `sync` command of the operating system.
3. The file system targeted for backup is frozen using the `fsfreeze` command of the operating system.
4. Snapshot retrieval (implementation of a logical copy) is performed using the `swsrpmake` command of ACM.
5. The file system frozen in step 3 is unfrozen using the `fsfreeze` command of the operating system.

Steps 3 to 5 are performed for all file systems targeted for backup.

finalize mode

1. The completion status of physical copies to all replication volumes is checked using the swsrpstat command of ACM.
2. Information about the replication volume used for the current backup is written to the backup information file.

restore mode

1. The backup information file is read, and the replication volume to be used for recovery is determined.
2. The file system targeted for recovery is unmounted using the umount command of the operating system.
3. Physical copy from the replication volume is performed using the swsrpmake command of ACM.
4. The file system targeted for recovery is mounted using the mount command of the operating system.
5. The completion status of physical copy from the replication volume is checked using the swsrpstat command of ACM.

Steps 2 to 5 are performed for all file systems targeted for recovery.

6. Files and directories not needed for archive recovery are deleted.



Information

The samples temporarily freeze the file system of the data storage destination (replication source volume) to protect the file system from copy processing by advanced copy features of the ETERNUS disk array. Therefore, consider the following:

- Consider freeze time during the timeout period of an SQL statement.

When using a feature that links with the database, such as a cluster feature that accesses the data storage destination, consider freeze time in relation to the timeout periods below.

- When performing database multiplexing

Misdetection may occur during abnormality monitoring, so it is necessary to consider the monitoring interval, timeout period and number of retries for abnormality monitoring, and consider temporarily stopping only the Mirroring Controller process during backup.

- When performing failover operations using PRIMECLUSTER

If a failure occurs while the file system is frozen, switching is triggered when PRIMECLUSTER detects an issue, but will be performed only after the system is unfrozen. In addition, if it takes time for the system to unfreeze, the active node may trigger operating system panic, resulting in a switch.

Appendix M Collecting Failure Investigation Data

For information on collecting failure investigation data, refer to "Collecting Failure Investigation Data" in the Installation and Setup Guide for Server.

Index

[A]	[E]
About using pgAdmin..... 3	Enabling and Disabling a Masking Policy..... 61
About using WebAdmin..... 2	Enabling Automatic Opening of the Keystore..... 30
Actions in Response to Instance Startup Failure..... 157	Encrypting a Tablespace..... 27,41
All user data within the specified tablespace..... 25	Encrypting Existing Data..... 34,45
Approximate backup time..... 16	Encryption mechanisms..... 25,38
Approximate recovery time..... 130	Errors in More Than One Storage Disk..... 157
Automatically opening the keystore..... 36	
[B]	[F]
Backing Up and Recovering the Keystore..... 30	Faster encryption and decryption based on hardware..... 25
Backing Up and Restoring/Recovering the Database..... 32	File system level backup and restore..... 33,44
Backup/Recovery Using the Copy Command..... 121	
Backup and recovery using the pgx_dmpall and pgx_rcvall commands..... 32,43	[H]
backup cycle..... 18	High-Speed Data Load..... 110
Backup data..... 25	
Backup operation..... 18	[I]
Backup operation (file backup)..... 19	If failure occurred in the data storage disk or the transaction log storage disk..... 131
Backup status..... 18,20	If failure occurred on the backup data storage disk..... 132,134
Backup using the backup information file..... 122	If failure occurred on the data storage disk or the transaction log storage directory..... 132
Backup Using the Copy Command..... 124	Importing and Exporting the Database..... 34,44
backup_destination (string)..... 166	Installing and Operating the In-memory Feature..... 100
Building and starting a standby server..... 36	
[C]	[K]
Changing a Masking Policy..... 60	keystore_location (string)..... 164
Changing the Keystore Passphrase..... 29	
Changing the Master Encryption Key..... 29	[L]
Changing the master encryption key and the passphrase..... 36	Logging in to WebAdmin..... 2
Checking an Encrypted Tablespace..... 28,42	log in..... 3
Checking backup status..... 125	
Checking the operating status of an instance..... 13,14	[M]
Collecting Failure Investigation Data 209	Managing the Keystore..... 29
Configuration of the Copy Command..... 121	Masking Condition..... 56
Configuration of the copy command for backup..... 123	Masking Format..... 56
Configuration of the copy command for recovery..... 123	Masking Policy..... 55
Confirming a Masking Policy..... 61	Masking Target..... 55
Continuous archiving and point-in-time recovery..... 33,44	Masking Type..... 56
Copy Command for Backup..... 126	Monitoring Database Activity..... 68
Copy Command for Recovery..... 128	
Copy Command Interface..... 126	[O]
Copy Command Samples that Use the Advanced Copy Feature of the ETERNUS Disk Array..... 207	Opening the Keystore..... 27
core_contents (string)..... 164	openssl_conf(string)..... 172
core_directory (string)..... 164	openssl_modules(string)..... 172
Creating a Masking Policy..... 60	Operating Fujitsu Enterprise Postgres..... 1
Cyclic usage of the backup area..... 121	
[D]	[P]
Data Masking..... 55	Parallel Query..... 108
Data Types for Masking..... 63	Performing backup..... 124
Deleting a Masking Policy..... 62	Perform recovery..... 125
Determining the backup area of the latest backup..... 125	Periodic Backup..... 18
	pgx_global_metacache (numerical value)..... 168
	pgx_stat_gmc view..... 197
	pgx_stat_latch view..... 195
	pgx_stat_lwlock view..... 195
	pgx_stat_progress_loader view..... 197
	pgx_stat_sql view..... 196

pgx_stat_vacuum_freeze.log_retention_period (numerical value).....	173
pgx_stat_vacuum_freeze.monitoring_interval (numerical value).....	173
pgx_stat_vacuum_freeze.track_freeze_info (boolean).....	173
pgx_stat_walwriter view.....	196
pgx_tablespaces.....	195
pgx_tablespaces view.....	195
Placement and automatic opening of the keystore file.....	35
Placing the keystore file.....	35
Preparing for backup.....	124

[R]

Recovery Using the Copy Command.....	125
reserve_buffer_ratio (numerical value).....	167

[S]

Scope of encryption.....	25
search_path (string).....	166
Security-Related Notes.....	36,46
Security Notes.....	63
Setting a restore point.....	20
Setting the Master Encryption Key.....	26,40
Starting and Stopping the Web Server Feature of WebAdmin.....	203
Starting an instance.....	12,13
Starting pgAdmin.....	3
Startup URL for WebAdmin.....	3
Stopping an instance.....	12,14
Streaming replication support.....	26
Streaming Replication Using WebAdmin.....	96
Strong encryption algorithms.....	25
System Administration Functions.....	174
System Views.....	195

[T]

tablespace_encryption_algorithm (string).....	166
Tables Used by Data Masking	199
Tables Used by Transparent Data Encryption.....	198
tde_kms.enable_shared_dek (boolean).....	164
tde_kms.kms_conninfo_file (string).....	164
tde_kms.max_shared_dek (numerical value).....	164
tde_kms.plugin_path (string).....	164
Tips for Installing Built Applications.....	37,46
track_gmc (string).....	169
track_sql (string).....	167
track_waits (string).....	167
Transparent Data Encryption Control Functions.....	174
Two-layer encryption key and the keystore.....	25,38

[U]

userprofile_database(string).....	172
Using Server Commands.....	13

[V]

vci.control_max_workers (numerical value).....	167
vci.cost_threshold (numeric).....	167
vci.enable (string).....	168

vci.log_query (string).....	168
vci.maintenance_work_mem (numerical value).....	168
vci.max_local_ros (numerical value).....	168
vci.max_parallel_degree (numerical value).....	168
vci.shared_work_mem (numerical value).....	168

[W]

WAL and temporary files.....	26
WAL Mirroring Control Functions.....	174
WebAdmin Wallet.....	205

Fujitsu Enterprise Postgres
18SP1 for x86

Operation Guide

Windows



Preface

Purpose of this document

The Fujitsu Enterprise Postgres database system extends the PostgreSQL features and runs on the Windows platform.

This document is the Fujitsu Enterprise Postgres Operation Guide.

Intended readers

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

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

- PostgreSQL
- SQL
- Windows

Structure of this document

This document is structured as follows:

[Chapter 1 Operating Fujitsu Enterprise Postgres](#)

Describes how to operate Fujitsu Enterprise Postgres.

[Chapter 2 Starting an Instance and Creating a Database](#)

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

[Chapter 3 Backing Up the Database](#)

Describes how to back up the database.

[Chapter 4 Configuring Secure Communication Using Secure Sockets Layer](#)

Describes communication data encryption between the client and the server.

[Chapter 5 Protecting Storage Data Using Transparent Data Encryption](#)

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

[Chapter 6 Policy-based Login Security](#)

Describes how to apply login security policies to users.

[Chapter 7 Data Masking](#)

Describes the data masking feature.

[Chapter 8 Periodic Operations](#)

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

[Chapter 9 Streaming Replication Using WebAdmin](#)

Describes how to create a streaming replication cluster using WebAdmin.

[Chapter 10 Installing and Operating the In-memory Feature](#)

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

[Chapter 11 Parallel Query](#)

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

[Chapter 12 High-Speed Data Load](#)

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

[Chapter 13 Global Meta Cache](#)

Describes how to use Global Meta Cache feature.

[Chapter 14 Local Meta Cache Limit](#)

Describes how to use Local Meta Cache Limit feature.

[Chapter 15 Backup/Recovery Using the Copy Command](#)

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

[Chapter 16 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 Catalogs](#)

Describes the system catalog used by Fujitsu Enterprise Postgres.

[Appendix D System Views](#)

Describes the system view used by Fujitsu Enterprise Postgres.

[Appendix E Tables Used by Data Masking](#)

Describes the tables used by the data masking feature.

[Appendix F Tables Used by Aggressive Freeze for Tuples](#)

Describes the tables used by the aggressive freeze for tuples.

[Appendix G Tables Used by High-Speed Data Load](#)

Describes the tables used by high-speed data load.

[Appendix H Starting and Stopping the Web Server Feature of WebAdmin](#)

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

[Appendix I WebAdmin Wallet](#)

Describes how to use the Wallet feature of WebAdmin.

[Appendix J WebAdmin Disallow User Inputs Containing Hazardous Characters](#)

Describes characters not allowed in WebAdmin.

[Appendix K Collecting Failure Investigation Data](#)

Describes how to collect information for initial investigation.

Export restrictions

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

Issue date and version

Edition 2.0: March 2026
Edition 1.0: December 2025

Copyright

Copyright 2015-2026 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.....	3
1.3 Starting pgAdmin.....	3
1.3.1 Starting pgAdmin.....	3
1.3.2 Adding an Instance.....	3
1.3.3 Connecting/Disconnecting an Instance.....	4
1.3.4 Changing the pgAdmin Display Language.....	5
1.4 Operations Using Commands.....	5
1.5 Operating Environment of Fujitsu Enterprise Postgres.....	6
1.5.1 Operating Environment.....	6
1.5.2 File Composition.....	8
1.5.3 Precautions when Using Antivirus Software.....	8
1.6 Notes on Compatibility of Applications Used for Operations.....	9
1.7 Notes on Upgrading Database Instances.....	9
1.7.1 Additional Steps for upgrading to Fujitsu Enterprise Postgres with Vertical Clustered Index (VCI) Enabled.....	10
1.7.2 Additional Steps to upgrading to Fujitsu Enterprise Postgres with Transparent Data Encryption (TDE) Enabled.....	11
Chapter 2 Starting an Instance and Creating a Database.....	12
2.1 Starting and Stopping an Instance.....	12
2.1.1 Using WebAdmin.....	12
2.1.2 Using Commands.....	13
2.2 Creating a Database.....	15
2.2.1 Using pgAdmin.....	15
2.2.2 Using Client Commands.....	15
Chapter 3 Backing Up the Database.....	16
3.1 Periodic Backup.....	18
3.2 Backup Methods.....	18
3.2.1 Using WebAdmin.....	18
3.2.2 Using Server Commands.....	18
Chapter 4 Configuring Secure Communication Using Secure Sockets Layer.....	22
4.1 Configuring Communication Data Encryption.....	22
4.1.1 Issuing a Certificate.....	23
4.1.2 Deploying a Server Certificate File and a Server Private Key File.....	23
4.1.3 Distributing a CA Certificate File to the Client.....	23
4.1.4 Configuring the Operating Environment for the Database Server.....	23
4.1.5 Configuring the Operating Environment for the Client.....	23
4.1.6 Performing Database Multiplexing.....	24
Chapter 5 Protecting Storage Data Using Transparent Data Encryption.....	25
5.1 Protecting Data Using Encryption.....	25
5.2 Setting the Master Encryption Key.....	26
5.3 Opening the Keystore.....	27
5.4 Encrypting a Tablespace.....	27
5.5 Checking an Encrypted Tablespace.....	28
5.6 Managing the Keystore.....	29
5.6.1 Changing the Master Encryption Key.....	29
5.6.2 Changing the Keystore Passphrase.....	29
5.6.3 Enabling Automatic Opening of the Keystore.....	30
5.6.4 Backing Up and Recovering the Keystore.....	30
5.7 Backing Up and Restoring/Recovering the Database.....	32
5.8 Importing and Exporting the Database.....	34
5.9 Encrypting Existing Data.....	34

5.10 Operations in Cluster Systems.....	34
5.10.1 HA Clusters that do not Use Database Multiplexing.....	35
5.10.2 Database Multiplexing Mode.....	35
5.11 Security-Related Notes.....	36
5.12 Tips for Installing Built Applications.....	36
Chapter 6 Policy-based Login Security.....	38
6.1 Advance Preparation.....	38
6.2 Changing the Contents of the default Profile.....	38
6.3 Creating and Assigning Profiles.....	39
6.4 Actions to be Taken in the Event of Deviation from Policy.....	39
6.5 Settings in Streaming Replication Configuration.....	40
6.6 Backup and Recovery.....	40
6.7 Profile parameters.....	41
6.8 Worker Processes.....	44
Chapter 7 Data Masking.....	45
7.1 Masking Policy.....	45
7.1.1 Masking Target.....	45
7.1.2 Masking Type.....	46
7.1.3 Masking Condition.....	46
7.1.4 Masking Format.....	46
7.2 Usage Method.....	49
7.2.1 Creating a Masking Policy.....	50
7.2.2 Changing a Masking Policy.....	50
7.2.3 Confirming a Masking Policy.....	51
7.2.4 Enabling and Disabling a Masking Policy.....	51
7.2.5 Deleting a Masking Policy.....	52
7.3 Data Types for Masking.....	53
7.4 Security Notes.....	53
Chapter 8 Periodic Operations.....	54
8.1 Configuring and Monitoring the Log.....	54
8.2 Monitoring Disk Usage and Securing Free Space.....	54
8.2.1 Monitoring Disk Usage.....	54
8.2.2 Securing Free Disk Space.....	54
8.3 Automatically Closing Connections.....	55
8.4 Monitoring the Connection State of an Application.....	55
8.4.1 Using the View (pg_stat_activity).....	56
8.4.2 Using pgAdmin.....	56
8.5 Reorganizing Indexes.....	57
8.6 Monitoring Database Activity.....	58
8.6.1 Information that can be Collected.....	59
8.6.2 Collection Configuration.....	59
8.6.3 Information Reset.....	60
8.7 Scheduling of an Aggressive Freeze for Tuples (VACUUM FREEZE).....	61
8.7.1 Monitoring Trends in Transaction ID Usage.....	61
8.7.2 How to Schedule Aggressive Freeze for Tuples.....	62
8.7.3 Tuning the Allocation Time for Aggressive Freeze for Tuples.....	62
8.8 Database Monitoring with Amazon CloudWatch.....	63
8.8.1 Metrics Collection.....	63
8.8.2 Log Collection.....	65
8.8.3 Metrics Collected by the Sample Script.....	65
8.9 Monitoring Deferred SQL and Periodically Backing up statistics.....	65
8.10 Performance Tuning.....	69
8.10.1 Enhanced Query Plan Stability.....	69
8.10.1.1 Fixing the Height of a Btree Index.....	70

Chapter 9 Streaming Replication Using WebAdmin.....	72
9.1 Creating a Standby Instance.....	72
9.2 Promoting a Standby Instance.....	73
9.3 Converting an Asynchronous Replication to Synchronous.....	73
9.4 Converting a Synchronous Replication to Asynchronous.....	74
9.5 Joining a Replication Cluster.....	74
Chapter 10 Installing and Operating the In-memory Feature.....	76
10.1 Installing Vertical Clustered Index (VCI).....	76
10.1.1 Evaluating whether to Install VCI.....	76
10.1.2 Estimating Resources.....	77
10.1.3 Setting up.....	77
10.1.3.1 Setting Parameters.....	77
10.1.3.2 Installing the Extensions.....	78
10.1.3.3 Creating a VCI.....	78
10.1.3.4 Confirming that the VCI has been Created.....	79
10.1.4 Data that can Use VCI.....	79
10.1.4.1 Relation Types.....	79
10.1.4.2 Data Types.....	80
10.2 Operating VCI.....	81
10.2.1 Commands that cannot be Used for VCI.....	81
10.2.2 Data Preload Feature.....	83
Chapter 11 Parallel Query.....	84
11.1 CPU Load Calculation.....	84
11.2 Increase of Workers during Runtime.....	84
11.3 Statistics View Displays the Action State.....	84
Chapter 12 High-Speed Data Load.....	86
12.1 Installing High-Speed Data Load.....	86
12.1.1 Deciding whether to Install.....	86
12.1.2 Estimating Resources.....	86
12.1.3 Setup.....	87
12.1.3.1 Setting Parameters.....	87
12.1.3.2 Installing the Extension.....	88
12.2 Using High-Speed Data Load.....	88
12.2.1 Loading Data.....	88
12.2.2 Checking Progress.....	89
12.2.3 Recovering from a Data Load that Ended Abnormally.....	90
12.3 Removing High-Speed Data Load.....	91
12.3.1 Removing the Extension.....	91
Chapter 13 Global Meta Cache.....	92
13.1 Usage.....	92
13.1.1 Deciding Whether to Enable the Global Meta Cache Feature.....	92
13.1.2 Estimating Memory for Global Meta Cache.....	92
13.1.3 How the GMC Memory Area Is Used.....	92
13.1.4 Enabling the Global Meta Cache Feature.....	92
13.1.5 Estimating Resources.....	93
13.2 Statistics.....	93
13.2.1 System View.....	93
Chapter 14 Local Meta Cache Limit.....	94
14.1 Usage.....	94
14.1.1 Deciding Whether to Enable the Local Meta Cache Limit Feature.....	94
14.1.2 How to Set Parameters for the Local Meta Cache Limit Feature.....	94
14.1.3 Cache Removal when Local Meta Cache Limit is Enabled.....	94
14.1.4 Performance Impact and Parameter Tuning of the Local Meta Cache Limit Feature.....	95

Chapter 15 Backup/Recovery Using the Copy Command.....	97
15.1 Configuration of the Copy Command.....	97
15.2 Backup Using the Copy Command.....	100
15.3 Recovery Using the Copy Command.....	101
15.4 Copy Command Interface.....	102
15.4.1 Copy Command for Backup.....	102
15.4.2 Copy Command for Recovery.....	104
Chapter 16 Actions when an Error Occurs.....	106
16.1 Recovering from Disk Failure (Hardware).....	107
16.1.1 Using WebAdmin.....	107
16.1.2 Using Server Command.....	108
16.2 Recovering from Data Corruption.....	112
16.2.1 Using WebAdmin.....	112
16.2.2 Using the pgx_rcvall Command.....	112
16.3 Recovering from an Incorrect User Operation.....	114
16.3.1 Using WebAdmin.....	114
16.3.2 Using the pgx_rcvall Command.....	115
16.4 Actions in Response to an Application Error.....	116
16.4.1 When using the view (pg_stat_activity).....	116
16.4.2 Using pgAdmin.....	117
16.5 Actions in Response to an Access Error.....	117
16.6 Actions in Response to Insufficient Space on the Data Storage Destination.....	117
16.6.1 Using a Tablespace.....	118
16.6.2 Replacing the Disk with a Larger Capacity Disk.....	118
16.6.2.1 Using WebAdmin.....	118
16.6.2.2 Using Server Commands.....	119
16.7 Actions in Response to Insufficient Space on the Backup Data Storage Destination.....	120
16.7.1 Temporarily Saving Backup Data.....	120
16.7.1.1 Using WebAdmin.....	121
16.7.1.2 Using Server Commands.....	121
16.7.2 Replacing the Disk with a Larger Capacity Disk.....	124
16.7.2.1 Using WebAdmin.....	124
16.7.2.2 Using Server Commands.....	125
16.8 Actions in Response to Insufficient Space on the Transaction Log Storage Destination.....	129
16.8.1 Replacing the Disk with a Larger Capacity Disk.....	129
16.8.1.1 Using WebAdmin.....	129
16.8.1.2 Using Server Commands.....	130
16.9 Errors in More Than One Storage Disk.....	131
16.10 Actions in Response to Instance Startup Failure.....	132
16.10.1 Errors in the Configuration File.....	132
16.10.2 Errors Caused by Power Failure or Mounting Issues.....	132
16.10.3 Other Errors.....	132
16.10.3.1 Using WebAdmin.....	132
16.10.3.2 Using Server Commands.....	133
16.11 Actions in Response to Failure to Stop an Instance.....	133
16.11.1 Using WebAdmin.....	133
16.11.2 Using Server Commands.....	133
16.11.2.1 Stopping the Instance Using the Fast Mode.....	133
16.11.2.2 Stopping the Instance Using the Immediate Mode.....	134
16.11.2.3 Forcibly Stopping the Server Process.....	134
16.12 Actions in Response to Failure to Create a Streaming Replication Standby Instance.....	134
16.13 Actions in Response to Error in a Distributed Transaction.....	135
16.14 I/O Errors Other than Disk Failure.....	136
16.14.1 Network Error with an External Disk.....	136
16.14.2 Errors Caused by Power Failure or Mounting Issues.....	136
16.15 Anomaly Detection and Resolution.....	136

16.15.1 Port Number and Backup Storage Path Anomalies.....	137
16.15.2 Mirroring Controller Anomalies.....	137
Appendix A Parameters.....	139
Appendix B System Administration Functions.....	147
B.1 WAL Mirroring Control Functions.....	147
B.2 Transparent Data Encryption Control Functions.....	147
B.2.1 pgx_open_keystore.....	147
B.2.2 pgx_set_master_key.....	147
B.2.3 pgx_set_keystore_passphrase.....	148
B.3 Profile Management Functions and User Management Functions.....	148
B.3.1 Profile Management Functions.....	148
B.3.2 User Management Functions.....	149
B.4 Data Masking Control Functions.....	150
B.4.1 pgx_alter_confidential_policy.....	150
B.4.2 pgx_create_confidential_policy.....	156
B.4.3 pgx_drop_confidential_policy.....	159
B.4.4 pgx_enable_confidential_policy.....	160
B.4.5 pgx_update_confidential_values.....	161
B.5 VCI Data Load Control Function.....	162
B.6 High-Speed Data Load Control Functions.....	162
Appendix C System Catalogs.....	164
C.1 pgx_profile.....	164
C.2 pgx_user_profile.....	164
C.3 pgx_auth_password.....	165
C.4 pgx_password_history.....	165
Appendix D System Views.....	167
D.1 pgx_tablespaces.....	167
D.2 pgx_stat_lwlock.....	167
D.3 pgx_stat_latch.....	167
D.4 pgx_stat_walwriter.....	168
D.5 pgx_stat_sql.....	168
D.6 pgx_stat_gmc.....	169
D.7 pgx_stat_progress_loader.....	169
Appendix E Tables Used by Data Masking.....	170
E.1 pgx_confidential_columns.....	170
E.2 pgx_confidential_policies.....	170
E.3 pgx_confidential_values.....	171
Appendix F Tables Used by Aggressive Freeze for Tuples.....	172
F.1 pgx_stat_freeze_results.....	172
Appendix G Tables Used by High-Speed Data Load.....	173
G.1 pgx_loader_state.....	173
Appendix H Starting and Stopping the Web Server Feature of WebAdmin.....	174
H.1 Starting the Web Server Feature of WebAdmin.....	174
H.2 Stopping the Web Server Feature of WebAdmin.....	174
Appendix I WebAdmin Wallet.....	175
I.1 Creating a Credential.....	175
I.2 Using a Credential.....	175
Appendix J WebAdmin Disallow User Inputs Containing Hazardous Characters.....	176
Appendix K Collecting Failure Investigation Data.....	177

Index.....178

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

Operation management using GUI tools

This involves managing operations using the WebAdmin and pgAdmin GUI tools.

- Management using WebAdmin

This removes the requirement for complex environment settings and operational design for backup and recovery that is usually required for running a database. It enables you to easily and reliably monitor the state of the database, create a streaming replication cluster, back up the database, and restore it even if you do not have expert knowledge of databases.

- Management using pgAdmin

When developing applications and maintaining the database, you can use pgAdmin to perform simple operations on database objects, such as:

- Rebuild indexes and update statistics
- Create, delete, and update database objects

In addition, from pgAdmin of Fujitsu Enterprise Postgres, you can use the expanded features provided by Fujitsu Enterprise Postgres on the PostgreSQL SQL commands.



See

Refer to pgAdmin Help for information on the expanded features of pgAdmin provided by Fujitsu Enterprise Postgres.

Operation management using commands

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

Points to consider when choosing an operation method

- You cannot combine WebAdmin and server commands to perform the following operations:
 - Use commands to operate an instance created using WebAdmin.
 - Use WebAdmin to recover a database backed up using commands.

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

- To operate an instance created using the `initdb` command in WebAdmin, the instance needs to be imported using WebAdmin.
- You can perform backup and restoration in pgAdmin, but the backup data obtained with WebAdmin and `pgx_dmpall` is not compatible with the backup data obtained with pgAdmin.

Features used in each phase

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

Operation		Operation with the GUI	Operation with commands
Setup	Creating an instance	WebAdmin is used. The server machine capacity, and the optimum parameter for operations using WebAdmin, are set automatically.	The configuration file is edited directly using the initdb command.
	Creating a standby instance	WebAdmin is used. WebAdmin performs a base backup of the source instance and creates a standby instance.	A standby instance is created using the pg_basebackup command.
	Changing the configuration files	WebAdmin is used.	The configuration file is edited directly.
Starting and stopping an instance		WebAdmin is used.	The net command or sc command of the operating system is used.
Creating a database		This is defined using pgAdmin of the GUI tool, or using the psql command or the application after specifying the DDL statement.	
Backing up the database		WebAdmin, or the pgx_dmpall command, is used.	It is recommended that the pgx_dmpall command be used. Recovery to the latest database can be performed.
Database recovery		WebAdmin is used.	To use the backup that was performed using the pgx_dmpall command, the pgx_rcvall command is used.
Monitoring	Database errors	The status in the WebAdmin window can be checked. (*1)	The messages that are output to the database server log are monitored (*1)
	Disk space	The status in the WebAdmin window can be checked. A warning will be displayed if the free space falls below 20%. (*1)	This is monitored using the fsutil command (check free space), and the dir command (check used space), of the operating system, for example. (*1)
	Connection status	This can be checked using pgAdmin of the GUI tool, or referencing pg_stat_activity of the standard statistics view from psql or the application.	

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

1.2 Starting WebAdmin

This section describes how to start and log in to WebAdmin.

About using WebAdmin

- It is recommended to use the following browsers with WebAdmin:

- Microsoft Edge

WebAdmin will work with other browsers, such as Firefox and Chrome, however, the look and feel may be slightly different.

- You must start the Web server feature of WebAdmin before using WebAdmin. Refer to "[Appendix H Starting and Stopping the Web Server Feature of WebAdmin](#)" for information on how to start the Web server feature of WebAdmin.

1.2.1 Logging in to WebAdmin

This section describes how to log in to WebAdmin.

Startup URL for WebAdmin

In the browser address bar, type the startup URL of the WebAdmin window in the following format:

```
http://hostNameOrIpAddress:portNumber/
```

- *hostNameOrIpAddress*: The host name or IP address of the server where WebAdmin is installed.
- *portNumber*: The port number of WebAdmin. The default port number is 27515.



Example

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

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

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

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

1.3 Starting pgAdmin

This section describes how to start pgAdmin, how to add an instance required for managing a database, and how to connect to and disconnect from the instance.

You can use pgAdmin on the Windows client.

About using pgAdmin

- You must start the instance to be connected to before using pgAdmin. Refer to "[2.1 Starting and Stopping an Instance](#)" for information on how to start an instance.
- When using pgAdmin4 with Microsoft Edge, enable network access by loopback in Microsoft Edge. Also, add Microsoft Edge to the loopback exclusion list.

1.3.1 Starting pgAdmin

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

Note that "<x>" indicates the product version. Also, the shortcuts displayed in the [Start] menu may be arranged differently depending on the operating system.

1.3.2 Adding an Instance

This section describes how to add an instance to be connected to.

If you use a link-local address with version 6 of the TCP/IP protocol, you may encounter the following error when registering the server with pgAdmin. Therefore, do not use link-local addresses.

```
unsupported format character ''' (0x22) at index 96
```

1. In the [Browser] pane, right-click [Servers], and then click [Register] >> [Server...].
2. In the [Register - Server] window, specify a value for each item.

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

- Name:** An empty text input field with a red error icon and a red border. A red error message at the bottom states: 'Name' cannot be empty.
- Server group:** A dropdown menu showing 'Servers'.
- Background:** A checkbox with an 'X' icon.
- Foreground:** A checkbox with an 'X' icon.
- Connect now?:** A checked toggle switch.
- Comments:** A large empty text area.

At the bottom of the dialog, there are buttons for 'Close', 'Reset', and 'Save', along with information and help icons.

[General] tab

- [Name]: Name of the instance to be managed

[Connection] tab

- [Host name/address]: Host name or IP address of the server where Fujitsu Enterprise Postgres is installed
- [Port]: Port number of the instance
- [Maintenance database]: Database name to which the client will connect
- [Username]: User name of the instance administrator
- [Password]: Password for the user name specified in [Username]

Point

.....

If you select [Save password], the Fujitsu Enterprise Postgres connection password is stored in the following location. Set the appropriate access permissions for the password file to protect it from unauthorized access.

- %APPDATA%\Roaming\pgAdmin\pgadmin4.db
-

1.3.3 Connecting/Disconnecting an Instance

This section describes how to connect pgAdmin to an instance, and how to disconnect it.

To connect to an instance created using WebAdmin, you must first configure the settings in the [Client authentication] window of WebAdmin to permit connection from pgAdmin.



See

Refer to "Changing the settings" in the Installation and Setup Guide for Server for information on the [Client authentication] window of WebAdmin.

Connecting to an instance

Starting pgAdmin does not connect it to any instance.

To connect to an instance, in the [Browser] pane, right-click the server, and then click [Connect Server].

If a password was not saved when the instance was added, enter a password in the password entry window that is displayed.

Disconnecting from an instance

To disconnect from an instance, in the [Browser] pane, right-click the server, and then click [Disconnect from Server].

1.3.4 Changing the pgAdmin Display Language

This section describes the procedure for changing the display language of pgAdmin.

1. Start pgAdmin.
2. Click [File] >> [Preferences].
3. Click [Miscellaneous] >> [User language] in the left pane of the [Preferences] dialog.
4. Select the language you want to display from the [User language] drop-down list displayed in the right pane, and click the [Save] button.
5. Refresh page by clicking the [Refresh] button.

1.4 Operations Using Commands

You can operate and manage the database using the following commands:

- Server commands

This group of commands includes commands for creating a database cluster and controlling the database. You can run these commands on the server where the database is operating.

To use these commands, you must configure the environment variables.



See

- Refer to "PostgreSQL Server Applications" under "Reference" in the PostgreSQL Documentation, or "Reference" for information on server commands.
- Refer to "Configure the environment variables" in the procedure to create instances in "Using the initdb Command" in the Installation and Setup Guide for Server for information on configuring the environment variables.

- Client commands

This group of commands includes the psql command and commands for extracting the database cluster to a script file. These commands can be executed on the client that can connect to the database, or on the server on which the database is running.

To use these commands, you must configure the environment variables.



See

- Refer to "PostgreSQL Client Applications" under "Reference" in the PostgreSQL Documentation, or "Reference" for information on client commands.


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

1.5 Operating Environment of Fujitsu Enterprise Postgres

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

1.5.1 Operating Environment

The following figure shows the configuration of the Fujitsu Enterprise Postgres operating environment. The tables given below list the roles of the OS resources and Fujitsu Enterprise Postgres resources.



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

Table 1.1 OS resources

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

Table 1.2 Fujitsu Enterprise Postgres client resources

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

Table 1.3 Server resources of Fujitsu Enterprise Postgres

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

1.5.2 File Composition

Fujitsu Enterprise Postgres consists of the following files for controlling and storing the database. The table below shows the relationship between the number of such files and their location within a single instance.

Table 1.4 Number of files within a single instance and how to specify their location

File type	Required	Quantity	How to specify the location
Program files	Y	Multiple	Note that "<x>" indicates the product version. By default, the following location is used: %Program Files%\Fujitsu\fsepv<x>server64
Database cluster	Y	1	Specify using WebAdmin or server commands.
Tablespace	Y	Multiple	Specify a space other than that under the database cluster, using pgAdmin or the DDL statement.
Backup	Y	Multiple	Specify using WebAdmin or server commands.
Core file	Y	Multiple	Specify using WebAdmin, server commands, or postgresql.conf.
Server certificate file (*1)	N	1	Specify using postgresql.conf.
Server private key file (*1)	N	1	Specify using postgresql.conf.
Master encryption key file (*1)	N	1	Specify the directory created as the key store using postgresql.conf.
Connection service file (*1)	N	1	Specify using environment variables.
Password file (*1)	N	1	Specify using environment variables.
CA certificate file (*1)	N	1	Specify using environment variables.

Y: Mandatory

N: Optional

*1: Set manually when using the applicable feature.

Information

If you are using antivirus software, refer to "1.5.3 Precautions when Using Antivirus Software".

1.5.3 Precautions when Using Antivirus Software

If anti-virus software is used, the instance might not start or not respond during operation, and proper operation of the database might not be guaranteed. To avoid such troubles, it is necessary to change the settings of the antivirus software so that virus scans are not executed on the active database.

Do scan exclusion for directories so that all files that comprise Fujitsu Enterprise Postgres are excluded from virus scanning. 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.

However, setting exclusions for scans poses a risk of missing malware that should be detected by antivirus software. To meet the security requirements necessary for system operation, implement risk countermeasures against malware threats across the entire system. Consider the following measures, for example:

- Introduce WAF, IPS, and IDS to prevent unauthorized access to the database server across the entire system.

- Disable unnecessary features, services, and protocols in the OS and network you use, and always obtain and apply the latest patches to prevent malware intrusion.
- Enable access log acquisition for the OS and database to detect data tampering or leakage caused by malware.
- Perform regular backups to ensure recovery from backups in case data is tampered with.
For backup methods of Fujitsu Enterprise Postgres, refer to "[Chapter 3 Backing Up the Database](#)".

1.6 Notes on Compatibility of Applications Used for Operations

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

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



See

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

1.7 Notes on Upgrading Database Instances

If you are upgrading a database instance that uses the extensions, follow the procedure below to upgrade according to the source product version.

When migrating from a database instance of FUJITSU Enterprise Postgres 10 or later

If you are using `pg_upgrade` to upgrade a database instance from FUJITSU Enterprise Postgres 10 or later that uses the extensions, follow the procedure below to upgrade.

It is strongly recommended to back up the database using `pg_dump` before performing `pg_upgrade`.

1. Create a database cluster at the destination
2. Stop the source and destination instances
3. Run the `pg_upgrade` command
4. Start the destination instance
5. Update the extensions used with `ALTER EXTENSION` to the latest version at the destination

```
ALTER EXTENSION extensionName UPDATE;
```

If you are using multiple extensions, update each extension separately.



See

For information about `pg_upgrade` and `ALTER EXTENSION`, refer to "Reference" in the PostgreSQL Documentation.

When migrating from a database instance of FUJITSU Enterprise Postgres 9.x

The following describes the tasks required to upgrade a database instance of FUJITSU Enterprise Postgres 9.x using the extension with `pg_upgrade`.

It is strongly recommended to back up the database using `pg_dump` before performing `pg_upgrade` or using `DROP EXTENSION`.

Before running `pg_upgrade`, remove the following extensions from all databases of the instance, except for "template0". If a column was created in the user table using the following extended data types, "DROP EXTENSION" will also drop that column. In such cases you might need to consider alternative upgrade methods rather than `pg_upgrade`. In that case, consider using `pg_dump` or `pg_restore`.

- `pg_stat_statements`

- pgx_io
- pgx_paging
- pgx_network
- pgx_network_err
- pgx_cpu
- pgx_memory
- pgx_swap
- pgx_disk
- pgx_process
- pgx_log
- oracle_compatible
- pg_dbms_stats
- pg_hint_plan

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

```
DROP EXTENSION extensionName;
```

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

```
CREATE EXTENSION extensionName;
```



See

For information about pg_upgrade, DROP EXTENSION and CREATE EXTENSION, refer to "Reference" in the PostgreSQL Documentation.

1.7.1 Additional Steps for upgrading to Fujitsu Enterprise Postgres with Vertical Clustered Index (VCI) Enabled

When upgrading FUJITSU Enterprise Postgres 11.0 or earlier instances that are using the VCI extension to FUJITSU Enterprise Postgres 12 or later using pg_upgrade, additional steps must be performed because of the incompatibility of the VCI extension between FUJITSU Enterprise Postgres 12 or later and FUJITSU Enterprise Postgres 11 or earlier.

Follow the procedure below in all databases in the FUJITSU Enterprise Postgres 11 or earlier instance, except "template0".

Before upgrading

1. Obtain the CREATE INDEX Definitions

Run the query below to list all the VCI indexes created in the database. Ensure that these indexes are re-created in the FUJITSU Enterprise Postgres 12 or later instance after pg_upgrade has finished.

```
SELECT nspname || '.' || relname AS index_relname,* FROM pg_class, pg_namespace
WHERE relnamespace = pg_namespace.oid AND relam IN (SELECT oid FROM pg_am WHERE amname='vci');
```

For each index_relname listed above, execute the commands below to obtain the CREATE INDEX definition (to use the same SQL syntax while re-creating the indexes on the FUJITSU Enterprise Postgres 12 or later instance).

```
SELECT pg_get_indexdef('indexName'::regclass);
```

2. Drop the VCI indexes and VCI extension along with all its dependencies.

To remove all the VCI indexes and VCI internal objects that are created in Fujitsu Enterprise Postgres, execute the commands below. VCI internal objects will be created in FUJITSU Enterprise Postgres 12 or later automatically when CREATE EXTENSION for VCI is executed.

```
DROP EXTENSION VCI CASCADE;
```

Information

To restore the VCI extension in the FUJITSU Enterprise Postgres 11 or earlier instance, execute CREATE EXTENSION.

After upgrading

Once the pg_upgrade operation is complete, for all databases except "template0", execute CREATE EXTENSION to create the VCI extension, and then execute CREATE INDEX for all the VCI indexes as required.

1.7.2 Additional Steps to upgrading to Fujitsu Enterprise Postgres with Transparent Data Encryption (TDE) Enabled

If you are using pg_upgrade to upgrade an instance of Fujitsu Enterprise Postgres that uses Transparent Data Encryption (TDE), there are steps you can take before upgrading.

If your old cluster operation was an HA cluster operation without database multiplexing operations, and you shared a single keystore file, see "[Before upgrading if you shared a keystore file](#)". For other operations, see "[Before upgrading](#)".

Before upgrading

Before upgrading, perform the following steps.

This is not necessary if you are using the old cluster keystore file location as the new cluster keystore file location. In that case, the old cluster cannot continue to be used. Enable automatic keystore opening for old and new clusters.

1. Copy Master Encryption Key

Copy the keystore file from the old cluster to the new cluster.

You do not need to use the pgx_set_master_key function to generate a new master encryption key on the new cluster; you must copy the keystore file from the old cluster.

As a database superuser, do the following:

```
$ mkdir <NEW-KEY-STORE>/
$ cp -p <OLD-KEY-STORE>/keystore.ks <NEW-KEY-STORE>/
```

NEW-KEY-STORE: The directory specified by the keystore_location parameter in postgresql.conf for the new cluster

OLD-KEY-STORE: The directory specified by the keystore_location parameter in postgresql.conf on the old cluster

Before upgrading if you shared a keystore file

If the primary and standby servers shared the same keystore file, copy the keystore file from the old environment and share it as the keystore file from the new environment.

For more secure management of keystore files, place them on a secure, isolated key management server or key management storage.

Enable automatic keystore opening for old and new clusters.

Chapter 2 Starting an Instance and Creating a Database

This chapter describes basic operations, from starting an instance to creating a database.

2.1 Starting and Stopping an Instance

This section describes how to start and stop an instance.

- [2.1.1 Using WebAdmin](#)
- [2.1.2 Using Commands](#)

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.
- The collected statistics are initialized if an instance is stopped in the "Immediate" mode or if it is abnormally terminated. To prepare for such initialization of statistics, consider regular collection of the statistics by using the SELECT statement. Refer to "The Statistics Collector" in "Server Administration" in the PostgreSQL Documentation for information on the statistics.


2.1.1 Using WebAdmin

WebAdmin enables you to start or stop an instance and check its operating status.

Starting an instance


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


 is displayed when an instance is stopped.

To start a stopped instance, click .

Stopping an instance

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

 is displayed when an instance is active.



To stop an active instance, click .

Stop mode

Select the mode in which to stop the instance. The following describes the operations of the modes:






Stop mode	Connected clients	Backup being executed using the command
Smart mode (*1)	Waits for all connected clients to be disconnected.	Waits for backups being executed using the command to finish.
Fast mode	Rolls back all transactions being executed and forcibly disconnects clients.	Terminates backups being executed using the command.
Immediate mode	All server processes are terminated immediately. Crash recovery is executed the next time the instance is started.	
Kill process mode	Send SIGKILL to the process and abort all active transactions. This will lead to a crash-recovery run at the next restart.	

*1: When the processing to stop the instance in the Smart mode has started and you want to stop immediately, use the following procedure:


1. Restart the Web server feature of WebAdmin.
2. In the [Instances] tab, click .
3. In the [Instances] tab, click , and select the Immediate mode to stop the instance.

Checking the operating status of an instance

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

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

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

When operating WebAdmin, click  to update the status. WebAdmin will reflect the latest status of the operation or the instance resources from the server.

If an error occurs while communicating with the server, there may be no response from WebAdmin. When this happens, close the browser and then log in again. If this does not resolve the issue, check the event log of the server and confirm whether a communication error has occurred.

The following message is output during startup of an instance when the startup process is operating normally, therefore, the user does not need to be aware of this message:

```
FATAL: the database system is starting up
```

2.1.2 Using Commands

The Windows service-related commands enable you to start or stop an instance and to check its operating state.

If you are to use Windows services, you should register instances in Windows services. It is possible to start and stop an instance by executing the pg_ctl command without registering the instance in Windows services, but because the instance is started as a user process, the postgres process will be forcibly stopped if you close the command prompt window in which the command was executed. For this reason, we recommend starting and stopping an instance using Windows services.

 See

.....
 Refer to "When an instance was created with WebAdmin" in "Configuring Automatic Start and Stop of an Instance" in the Installation Guide for Server for information on registering instances in Windows services.

Starting an instance

You can start an instance by specifying the service name in the net start command or sc start command.

Also, you can use the following procedure to start an instance in the Windows services window:

1. Display the [Services] window
 In the [Administrative Tools], click [Services].

2. Start a service

Select the instance name that you wish to start from the services list, and click [Start Service].

Stopping an instance

You can stop an instance by specifying the service name in the net stop command or sc stop command.

Also, you can use the following procedure to stop an instance in the Windows services window:

1. Display the [Services] window

In the [Administrative Tools], click [Services].

2. Stop the service

Select the instance name that you wish to stop from the services list, and click [Stop Service]. If you stop a service while applications and commands are running, Fujitsu Enterprise Postgres will force those applications and commands to close and will stop normally.

Checking the operating state of an instance

Use the following procedure to check if an instance is operating correctly immediately after performing the operation to start an instance:

1. Display the [Services] window

In the [Start] menu, select [Administrative Tools], and then click [Services].

2. Check the state of the service

In the services list, check the state of the services for the applicable Fujitsu Enterprise Postgres.

To check the operating state of an instance during operation, use the pg_ctl command.

Specify the following in the pg_ctl command:

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



Example

When the instance is active:

```
> pg_ctl status -D D:\database\inst1
pg_ctl: server is running (PID: 1234)
```

When the instance is inactive:

```
> pg_ctl status -D D:\database\inst1
pg_ctl: no server running
```



Information

You can also use the net start command or sc query command to check the operating state of an instance.



See

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

2.2 Creating a Database

This section explains how to create a database.

- [2.2.1 Using pgAdmin](#)
- [2.2.2 Using Client Commands](#)

2.2.1 Using pgAdmin

Follow the procedure below to define a database using pgAdmin.

1. In the pgAdmin window, right-click [Databases] in the [Browser] pane, and then click [Create] >> [Database...] to display a [Create - Database] window.
 - [General] tab
 - [Database]: Name of the database to be managed
3. Click [Save] to create the database.

2.2.2 Using Client Commands

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

An example of operations on the server is shown below.

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

```
> psql postgres
psql (<x>) (*1)
Type "help" for help.
```

*1: <x> indicates the PostgreSQL version on which this product is based.

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

```
postgres=# CREATE DATABASE db01;
CREATE DATABASE
```

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

```
postgres=# \l+
```

4. Disconnect from the postgres database.
Execute \q to terminate the psql command.

```
postgres=# \q
```

You can create a database using the createdb command.



See

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

Chapter 3 Backing Up the Database

This chapter describes how to back up the database.

Backup methods

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

- Backup using WebAdmin

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

WebAdmin is used for recovery.

If performing backups with WebAdmin, the following password file is temporarily created during backup for WebAdmin to connect to the database:

- `userProfileFolder\localSettingsFolder\Fujitsu\fsep_version\instanceName\pgpass.conf`

Therefore, when you are backing up corefiles created in the `core_directory` parameter of `postgresql.conf`, or log files created in the `log_directory` parameter of `postgresql.conf`, ensure not to back up the password files located in the same directories at the same time.

- 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.

Additionally, you can also take backups periodically by using any external scheduler.

The `pgx_rcvall` command is used for recovery.

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.

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 15 Backup/Recovery Using the Copy Command](#)" for details.

Approximate backup time

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

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

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

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

When defining a tablespace

If you have defined a tablespace, back it up. If you do not back it up, directories for the tablespace are not created during recovery, which may cause the recovery to fail. If the recovery fails, refer to the event log, create the tablespace, and then perform the recovery process again.

When encrypting data stored in the database

There are several considerations for the backup of the keystore and backup of the database in case the data stored in the database is encrypted. Refer to the following for details:

- [5.6.4 Backing Up and Recovering the Keystore](#)
- [5.7 Backing Up and Restoring/Recovering the Database](#)

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.

The following backup methods are available for the features provided by Enterprise Postgres:

Category	Backup method	Backup target	Enterprise Postgres					
			Transparent Data Encryption	Policy-based Login Security	Data Masking	Confidentiality Management	Audit Log Feature	Database Multiplexing
Physical Backup	WebAdmin	Database cluster	Y	Y	Y	Y	Y	Y
	pgx_dmpall command		Y	Y	Y	Y	Y	Y
	Backup by continuous archiving		Y	Y	Y	Y	Y	Y
	File system level backup		Y	Y	Y	Y	Y	Y
Logical backup	pg_dumpall command		N	Y	N	Y	N	Y
	pg_dump command	Database	N	N	N	Y	N	Y
	COPY command	Table	N	N	N	N	N	N

Y: Can be used

N: Cannot be used

In a logical backup, the restore process recreates the indexes. For example, if the index size is large, such as for vector data handled by pgvector, the restoration time may be long. If you want to restore quickly, use a physical backup.

To export large amounts of data from multiple tables at high speed, use multiple executions to export data in parallel.

Some features are important to keep in mind when backing up. For information about the features, see the following:

Transparent Data Encryption : "[5.7 Backing Up and Restoring/Recovering the Database](#)"

Policy-based Login Security: "[6.6 Backup and Recovery](#)"

Confidentiality Management: "Backup/Restore" in the Security Operation Guide

Database Multiplexing: "Backup Operation" in the Cluster Operation Guide(Database Multiplexing)

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.

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.

When encrypting data stored in the database

If the data to be stored in the database is to be encrypted, it is necessary to enable the automatic opening of the keystore before doing so. Refer to "[5.6.3 Enabling Automatic Opening of the Keystore](#)" for details.

3.2.2 Using Server Commands

Use the `pgx_dmpall` command and `pgx_rcvall` command to perform backup and check the backup result.

Preparing for backup

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

Follow the procedure below.



See

Refer to "Preparing Directories to Deploy Resources" in the Installation and Setup Guide for Server for information on the location of directories required for backup and for points to take into account.

1. Prepare the backup data storage disk

For backup, prepare a separate disk unit from the database storage disk and mount it using the operating system commands.

2. Create a directory where the backup data will be stored

Create an empty directory.

In [Properties] in Windows(R) Explorer, set appropriate permissions so that only the instance administrator can access the directory.



See

Refer to [Help and Support] in Windows(R) for information on [Properties].

3. Specify the settings required for backup

Stop the instance, and set the following parameters in the postgresql.conf file.

Start the instance after editing the postgresql.conf file.

Parameter name	Setting	Description
backup_destination	Name of the directory where the backup data will be stored	Specify the name of the directory where the backup data will be stored. Appropriate privileges that allow only the instance administrator to access the directory must already be set. Place the backup data storage destination directory outside the data storage destination directory, the tablespace directory, and the transaction log storage destination directory.
archive_mode	on	Specify the archive log mode. Specify [on] (execute).
archive_command	'cmd /c "" <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. Note the following when specifying the path: - Specify \\ as the path delimiter. - Enclose the path in double quotes ("") if it contains spaces.
archive_library	""(default)	Specify the archive library.

Parameter name	Setting	Description
		Specify [""](default).

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 D:\database\inst1
```

Backup stores the data obtained during the backup and the backup data of the data obtained during previous backup.

If the data to be stored in the database is encrypted, refer to the following and back up the keystore:

- [5.6.4 Backing Up and Recovering the Keystore](#)

Backup status

Use the `pgx_rcvall` command to check the backup status.

Specify the following values in the `pgx_rcvall` command:

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

```
> pgx_rcvall -l -D D:\database\inst1
Date                Status           Dir
2022-03-01 13:30:40 COMPLETE        E:/backup/inst1/2022-03-01_13-30-40
```

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

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



See

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

Setting a restore point

In case you want to recover your database to a certain point in time, you can name this particular point in time, which is referred to as the restore point, by using the `psql` command.

By setting a restore point before executing an application, it becomes easy to identify up to which point in time the data will be reverted.

A restore point can be set to any point in time after a backup is executed. However, if a restore point is set before a backup is executed, the database cannot be recovered to that point in time. This is because restore points are recorded in the archive logs, and the archive logs are discarded when backups are executed.

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.



Example

The following example uses the psql command to connect to the database and execute the SQL statement to set a restore point.

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

```
postgres=# SELECT pg_create_restore_point('batch_20220303_1');
LOG:  restore point "batch_20220303_1" created at 0/20000E8
STATEMENT:  select pg_create_restore_point('batch_20220303_1');
pg_create_restore_point
-----
0/20000E8
(1 row)
```

Refer to "16.3.2 Using the pgx_rcvall Command" for information on using a restore point to recover the database.



See

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

Chapter 4 Configuring Secure Communication Using Secure Sockets Layer

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

4.1 Configuring Communication Data Encryption


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

Table 4.1 Configuration procedure

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

The following figure illustrates the environment for communication data encryption.

Figure 4.1 Environment for communication data encryption



4.1.1 Issuing a Certificate

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

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

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

- a. Create a server private key file
- b. Disable the passphrase for the server private key file
- c. Create a CSR (signing request for obtaining a server certificate) from the server private key file
- d. Apply to the certificate authority (CA) for a server certificate
- e. Obtain a server certificate file and a CA certificate file from the certificate authority (CA)
- f. Store the server certificate file and the CA certificate file

Note: If you lose or destroy the certificates, you will need to have them re-issued.

The above procedure enables you to prepare the following files:

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

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

Create a directory on the local disk of the database server and store the server certificate file and the server private key file in it.

Use the operating system features to set access privileges for the server certificate file and the server private key file so that only the database administrator has load privileges.

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

4.1.3 Distributing a CA Certificate File to the Client

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

4.1.4 Configuring the Operating Environment for the Database Server



See

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

4.1.5 Configuring the Operating Environment for the Client



See

.....
Refer to the following sections in the Application Development Guide for details, depending on your application development environment:

- "Settings for Encrypting Communication Data" under "Setup" in "JDBC Driver"
 - "Settings for Encrypting Communication Data" under "Setup" in "C Library (libpq)"
 - "Settings for Encrypting Communication Data" under "Setup" in "Embedded SQL in C"
-

4.1.6 Performing Database Multiplexing

When you perform communication that uses database multiplexing and a Secure Socket Layer server certificate, certificates with the same "Common Name" must be used. To ensure this, take one of the following actions:

- Create one server certificate, replicate it, and place a copy on each server used for database multiplexing.
- Create a server certificate with the same "Common Name" for each server used for database multiplexing.



See

.....
Refer to "Using the Application Connection Switch Feature" in the Application Development Guide for information on how to specify applications on the client.
.....

Chapter 5 Protecting Storage Data Using Transparent Data Encryption

This chapter describes how to encrypt data to be stored in the database.

5.1 Protecting Data Using Encryption

With PostgreSQL, data in a database is protected from access by unauthorized database users through the use of authentication and access controls. However, the OS file is not protected from attackers who bypass the database server's authentication and access controls.

With Fujitsu Enterprise Postgres, data inside the OS file is encrypted, so valuable information is protected even if the file or disk is stolen.

Data to be stored in a database is encrypted when it is written to the data file, and decrypted when it is read.

This is performed automatically by the instance, so the user and the application need not be aware of key management and encryption or decryption. This process is called TDE (Transparent Data Encryption).

The characteristics of TDE are described below.

Encryption mechanisms

Two-layer encryption key and the keystore

In each tablespace, there is a tablespace encryption key that encrypts and decrypts all the data within. The tablespace encryption key is encrypted by the master encryption key and saved.

Only one master encryption key exists in a database cluster. It is encrypted based on a passphrase specified by the user and stored in a keystore. Fujitsu Enterprise Postgres provides a file-based keystore. Attackers who do not know the passphrase cannot read the master encryption key from the keystore.

Strong encryption algorithms

TDE uses the Advanced Encryption Standard (AES) as its encryption algorithm. AES was adopted as a standard in 2002 by the United States Federal Government, and is used throughout the world.

Faster encryption and decryption based on hardware

TDE minimizes the overhead of encryption and decryption by using the AES-NI (Advanced Encryption Standard New Instructions) built into Intel(R) Xeon(R) processors since the 5600 series. This means that even in situations where previously the minimum encryption target was selected as a tradeoff between performance and security, it is now possible to encrypt all the data of an application.

Refer to the Intel Corporation's website for information on the list of processors equipped with AES-NI.

Zero overhead storage areas

Encryption does not change the size of data stored in tables, indexes, or WAL. There is, therefore, no need for additional estimates or disks.

Scope of encryption

All user data within the specified tablespace

The tablespace is the unit for specifying encryption. All tables, indexes, temporary tables, and temporary indexes created in the encrypted tablespace are encrypted. There is no need for the user to consider which tables and strings to encrypt.

Refer to "[5.4 Encrypting a Tablespace](#)" for details.

Backup data

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

Refer to "[5.7 Backing Up and Restoring/Recovering the Database](#)" after enabling or reconfiguring encryption to back up the database.

WAL and temporary files

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

Streaming replication support

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

The following are not encrypted:

- `pg_dump` and `pg_dumpall` output files
- Files output by the `COPY` command
- Notification event payloads that communicate using the `LISTEN` or `NOTIFY` command
- Checksum validation is not performed on encrypted tablespaces during backup and when using the `pg_checksum` utility.

5.2 Setting the Master Encryption Key

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

1. In the `keystore_location` parameter of `postgresql.conf`, specify the directory to store the keystore.

Specify a different location for each database cluster. 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

```
keystore_location = 'C:\\key\\store\\location'
```

Refer to "[Appendix A Parameters](#)" for information on `postgresql.conf`.

After editing the `postgresql.conf` file, either start or restart the instance.

- Using WebAdmin
Refer to "[2.1.1 Using WebAdmin](#)", and restart the instance.
- Using commands
Refer to "[2.1.2 Using Commands](#)", and restart the instance.

2. Execute the `pgx_set_master_key` function to set the master encryption key.

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.

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.

Refer to "[B.2 Transparent Data Encryption Control Functions](#)" for information on the `pgx_set_master_key` function.



Note

Note that if you forget the passphrase, you will not be able to access the encrypted data. There is no method to retrieve a forgotten passphrase and decrypt data. Do not, under any circumstances, forget the passphrase.

5.3 Opening the Keystore

To create encrypted tablespaces and access the encrypted data, you must first open the keystore. When you open the keystore, the master encryption key is loaded into the database server memory and becomes usable for encryption and decryption.

You need to open the keystore each time you start the instance. To open the keystore, the database superuser must execute the following SQL function.

```
SELECT pgx_open_keystore('passphrase');
```

The value "passphrase" is the passphrase specified during creation of the keystore.

Refer to "[B.2 Transparent Data Encryption Control Functions](#)" for information on the `pgx_open_keystore` function.

Note that, in the following cases, the passphrase must be entered when starting the instance, because the encrypted WAL must be decrypted for recovery. In this case, the above-mentioned `pgx_open_keystore` function cannot be executed.

- If performing crash recovery at the time of starting the instance
- If performing recovery using continuous archiving

For the above cases, select one of the following methods:

- Use an automatically opening keystore

Select this method if ease of operation has priority over enhanced security. When using an automatically opening keystore, the content of the keystore file is decrypted and a copy of the keystore file is generated. Although the content of this file is obfuscated, the level of security becomes slightly weaker.

Select this method if performing operations using WebAdmin.

- Enter a passphrase when starting an instance

Select this method if enhanced security has priority over ease of operation.

Specify the `--keystore-passphrase` in the `pg_ctl` command and start the instance. This displays the prompt that asks for the passphrase to be entered.

```
> pg_ctl --keystore-passphrase start
Enter the passphrase:
The server is currently initiating
>
```

After performing the above operation, use the `pg_ctl` command to stop the instance.

Then start the instance in Windows services. Refer to "[2.1.2 Using Commands](#)" for information on how to start an instance in Windows services.

Point

When using an automatically opening keystore, you do not need to enter the passphrase and you can automatically open the keystore when the database server starts. Refer to "[5.6.3 Enabling Automatic Opening of the Keystore](#)" for details.

5.4 Encrypting a Tablespace

The keystore must be open before you can create an encrypted tablespace.

When creating a tablespace that will be encrypted, configure the encryption algorithm in the runtime parameters. For example, to create a tablespace with the name `secure_tablespace` using AES with a key length of 256 bits as the encryption algorithm, configure as shown below.

```
-- Specify the encryption algorithm for the tablespace to be created below
SET tablespace_encryption_algorithm = 'AES256';
CREATE TABLESPACE secure_tablespace LOCATION 'C:\My\Data\Dir';
-- Specify that the tablespace to be created below is not to be encrypted
SET tablespace_encryption_algorithm = 'none';
```

Or

```
CREATE TABLESPACE secure_tablespace LOCATION '\My\Data\Dir' WITH (tablespace_encryption_algorithm = 'AES256' );
```

When the tablespace is empty, the encryption algorithm can be modified with the command below.

```
ALTER TABLESPACE secure_tablespace SET (tablespace_encryption_algorithm=AES256);
```

Trying to set the encryption algorithm for a non-empty tablespace causes an error.

You can use AES with a key length of 128 bits or 256 bits as the encryption algorithm. It is recommended that you use 256-bit AES. Refer to "[Appendix A Parameters](#)" for information on how to specify the runtime parameters.

If user provides both GUC and command line options while creating the tablespace, the preference is given to the command line option.

The `pg_default` and `pg_global` tablespaces cannot be encrypted.

Create tables and indexes in the encrypted tablespace that you created. Relations created in the encrypted tablespace are automatically encrypted.



Example

Example 1: Specifying an encrypted tablespace when creating it

```
CREATE TABLE my_table (...)  
    TABLESPACE secure_tablespace;
```

Example 2: Not explicitly specifying a tablespace when creating it and instead using the default tablespace

```
SET default_tablespace = 'secure_tablespace';  
CREATE TABLE my_table (...);
```

The process is the same for encrypting temporary tables and temporary indexes. In other words, either explicitly specify the `TABLESPACE` clause or list encrypted tablespaces in the `temp_tablespaces` parameter, and then execute `CREATE TEMPORARY TABLE` or `CREATE INDEX`.



Point

If an encrypted tablespace is specified in the `TABLESPACE` clause of the `CREATE DATABASE` statement, relations created in the database without explicitly specifying a tablespace will be encrypted. Furthermore, the system catalog will also be encrypted, so the source code of user-defined functions is also protected.

Example: Specifying a tablespace in a database definition statement

```
CREATE DATABASE DB01 TABLESPACE=SP01 ... ;
```

Part of the data is also stored in the system catalog - to encrypt this data as well, specify an encrypted tablespace as above and create a database.

An encrypted tablespace cannot be created from the window used for creating the pgAdmin tablespace, or from the query tool using a multi-command string. To create an encrypted tablespace, use `psql` console or a single command in the query tool such as that shown in the example below.

```
CREATE TABLESPACE tspace LOCATION '/some/dir/in/the/path' WITH  
(tablespace_encryption_algorithm='AES256')
```

5.5 Checking an Encrypted Tablespace

The `pgx_tablespaces` system view displays information about whether each tablespace has been encrypted, and about the encryption algorithm. Refer to "[D.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;
```



Example

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

spcname	spcencalgo
pg_default	none
pg_global	none
secure_tablespace	AES256

(3 rows)



See

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

5.6 Managing the Keystore

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

5.6.1 Changing the Master Encryption Key

Using the same encryption key for an extended period gives attackers an opportunity to decipher the encrypted data. It is recommended that you change the key at regular intervals, or whenever the key is exposed to risk.

Adhere to the industry's best practices for encryption algorithms and key management when considering how often the key should be changed. For example, the NIST in the United States has published "NIST Special Publication 800-57". The PCI DSS also refers to this publication. This publication recommends changing the master encryption key once a year.

To change the master encryption key, execute the `pgx_set_master_key` function, which is the same function used for configuring the key. Refer to "[5.2 Setting the Master Encryption Key](#)" for details.

After changing the master encryption key, you must immediately back up the keystore.

5.6.2 Changing the Keystore Passphrase

In security policies for organizations, it is usually a requirement that the passphrase be changed whenever a security administrator who knows the passphrase is removed from duties due to transfer or retirement. It is also recommended that the passphrase be changed if it is ever exposed to risks due to deception such as social engineering.

To change the keystore passphrase, execute the following SQL function as a superuser.

```
SELECT pgx_set_keystore_passphrase('oldPassphrase', 'newPassphrase');
```

After changing the passphrase, you must immediately back up the keystore.

Refer to "[B.2 Transparent Data Encryption Control Functions](#)" for information on the `pgx_set_keystore_passphrase` function.

5.6.3 Enabling Automatic Opening of the Keystore

When using an automatically opening keystore, you do not need to enter the passphrase and you can automatically open the keystore when the instance starts. Execute the `pgx_keystore` command to enable automatic opening of the keystore.

```
> pgx_keystore --enable-auto-open C:\key\store\location\keystore.ks
Enter the passphrase:
Automatic opening of the keystore is now enabled
>
```



Refer to "pgx_keystore" in the Reference for information on `pgx_keystore` command.

When automatic opening is enabled, an automatically opening keystore is created in the same directory as the original keystore. The file name of the automatically opening keystore is `keystore.aks`. The file `keystore.aks` is an obfuscated copy of the decrypted content of the `keystore.ks` file. As long as this file exists, there is no need to enter the passphrase to open the keystore when starting the instance.

Do not delete the original keystore file, `keystore.ks`. It is required for changing the master encryption key and the passphrase. When you change the master encryption key and the passphrase, `keystore.aks` is recreated from the original keystore file, `keystore.ks`.

Protect `keystore.ks`, `keystore.aks`, and the directory that stores the keystore so that only the user who starts the instance can access them.

Configure the permission of the files so that only the user who starts the instance can access the SQL functions and commands that create these files. Accordingly, manually configure the same permission mode if the files are restored.

Set the permission mode in [Properties] in Windows(R) Explorer.



Refer to [Help and Support] in Windows(R) for information on [Properties].

To use WebAdmin for backup and recovery, you must enable automatic opening of the keystore.

An automatically opening keystore will only open on the computer where it was created.

To disable automatic opening of the keystore, delete `keystore.aks`.

5.6.4 Backing Up and Recovering the Keystore

Back up the keystore at the following times in case it is corrupted or lost. Note that you must store the database and the keystore on separate data storage media. Storing both on the same data storage medium risks the danger of the encrypted data being deciphered if the medium is stolen. A passphrase is not required to open an automatically opening keystore, so store this type of keystore in a safe location.

- When the master encryption key is first configured
- When the master encryption key is changed
- When the database is backed up
- When the keystore passphrase is changed



Do not overwrite an old keystore when backing up a keystore. This is because during database recovery, you must restore the keystore to its state at the time of database backup. When the backup data of the database is no longer required, delete the corresponding keystore.

Example

- Back up the database and the keystore on March 1, 2022.

```
> pgx_dmpall -D D:\database\inst1
> copy C:\key\store\location\keystore.ks C:\keybackup\keystore_20220301.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 March 5, 2022.

```
> psql -c "SELECT pgx_set_master_key('passphrase')" postgres
> copy C:\key\store\location\keystore.ks C:\keybackup\keystore_20220305.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 March 5, 2022.

```
> copy C:\keybackup\keystore_20220305.ks C:\key\store\location\keystore.ks
```

- If there is no backup of the keystore containing the latest master encryption key, recover the keystore by restoring the keystore that was backed up along with the database on 1 March 2022.

```
> copy C:\keybackup\keystore_20220301.ks C:\key\store\location\keystore.ks
> pgx_rcvall -B E:\backup\inst1 -D D:\database\inst1 --keystore-passphrase
```

Specify the following in the `pgx_rcvall` command:

- Specify the data storage directory in the `-D` option. If the `-D` option is omitted, the value of the `PGDATA` environment variable is used by default.
- Specify the backup data storage directory in the `-B` option.
- The `--keystore-passphrase` option prompts you to enter the passphrase to open the keystore.

If you have restored the keystore, repeat the process of enabling automatic opening of the keystore. This ensures that the contents of the automatically opening keystore (`keystore.aks`) are identical to the contents of the restored keystore.

It is recommended that you do not back up the automatically opening keystore file, `keystore.aks`. If the database backup medium and the backup medium storing the automatically opening keystore are both stolen, the attacker will be able to read the data even without knowing the passphrase.

If the automatically opening keystore is corrupted or lost, you must again enable automatic opening. The `keystore.aks` file will be recreated from `keystore.ks` at this time.

See

Refer to "`pgx_rcvall`" and "`pgx_dmpall`" in the Reference for information on the `pgx_rcvall` and `pgx_dmpall` commands.

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

Refer to "[B.2 Transparent Data Encryption Control Functions](#)" for information on the `pgx_set_master_key` function.

Refer to "[5.6.3 Enabling Automatic Opening of the Keystore](#)" for information on how to enable automatic opening of the keystore.

5.7 Backing Up and Restoring/Recovering the Database

Fujitsu Enterprise Postgres enables you to use the five backup and recovery methods described below. Regardless of the method you use, you must back up the keystore at the same time.

Note that you must store the database and the keystore on separate data storage media. Storing both on the same data storage medium risks the danger of the encrypted data being deciphered if the medium is stolen.

Backup and recovery using WebAdmin

- Backup

WebAdmin backs up encrypted data.

Back up the key store after backing up the database.

- Recovery

Restore the keystore to its state at the time of database backup. Refer to "[5.6.4 Backing Up and Recovering the Keystore](#)" for details.

Enable automatic opening of the keystore in accordance with the procedure described in "[5.6.3 Enabling Automatic Opening of the Keystore](#)". Then, use WebAdmin to recover the database.

Backup and recovery using the `pgx_dmpall` and `pgx_rcvall` commands

- Backup

The `pgx_dmpall` command backs up the encrypted data.

Back up the key store after backing up the database.

- Recovery

Restore the keystore to its state at the time of the database backup.

Configure automatic opening of the key store as necessary.

If automatic opening of the keystore is not enabled, execute the `pgx_rcvall` command with the `--keystore-passphrase` option specified. This will display the prompt for the passphrase to be entered.



Example

- Back up the database and the keystore on March 1, 2022.

```
> pgx_dmpall -D D:\database\inst1
> copy C:\key\store\location\keystore.ks C:\keybackup\keystore_20220301.ks
```

Specify the following in the `pgx_dmpall` command:

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

- Recover the database and the keystore from the backup taken on March 1, 2022.

```
> copy C:\keybackup\keystore_20220301.ks C:\key\store\location\keystore.ks
> pgx_keystore --enable-auto-open C:\key\store\location\keystore.ks (Execute only when enabling
automatic opening)
> pgx_rcvall -B E:\backup\inst1 -D D:\database\inst1 --keystore-passphrase
```

Specify the following in the `pgx_rcvall` command:

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

- Specify the backup data storage directory in the -B option.
 - The --keystore-passphrase option prompts you to enter the passphrase to open the keystore.
-

Dump and restore using SQL

- Backup

The files output by the `pg_dump` and `pg_dumpall` commands are not encrypted. You should, therefore, encrypt the files using OpenSSL commands or other means before saving them, as described in "[5.8 Importing and Exporting the Database](#)" below.

Back up the key store after backing up the database.

- Restore

If the backup data has been encrypted using, for example Open SSL commands, decrypt that data.

The data generated by the `pg_dumpall` command includes a specification to encrypt tablespaces by default. For this reason, the `psql` command encrypts tablespaces during restoration.

File system level backup and restore

- Backup

Stop the instance and backup the data directory and the tablespace directory using the file copy command of the operating system. The files of encrypted tablespaces are backed up in the encrypted state.

Back up the key store after performing the backup.

- Restore

Restore the keystore to its state at the time of the database backup.

Stop the instance and restore the data directory and the tablespace directory using the file copy command of the operating system.

Continuous archiving and point-in-time recovery

- Backup

The `pg_basebackup` command backs up the encrypted data as is.

Back up the key store after performing the backup.

- Recovery

Restore the keystore to its state at the time of the database backup.

Configure automatic opening of the key store as necessary.

If automatic opening of the keystore is not enabled, refer to "[5.3 Opening the Keystore](#)" for information on starting an instance by specifying `pg_ctl --keystore-passphrase start`.



See

- Refer to "Reference" in the PostgreSQL Documentation for information on the following commands:

- `psql`
- `pg_dump`
- `pg_basebackup`

- Refer to the Reference for information on the following commands:

- `pgx_rcvall`
 - `pgx_dmpall`
 - `pg_dumpall`
-

If you have restored the keystore, repeat the process of enabling automatic opening of the keystore. This ensures that the contents of the automatically opening keystore (keystore.aks) are identical to the contents of the restored keystore.

Refer to "[5.6.3 Enabling Automatic Opening of the Keystore](#)" for information on how to enable automatic opening of the keystore.

5.8 Importing and Exporting the Database

The files output by the COPY TO command are not encrypted. Therefore, when transferring files to other systems, you should encrypt files using OpenSSL commands, or use file transfer software that performs encrypted communication for Windows, to encrypt the data being transferred.

Use a safe method to delete obsolete plain text files.

You can use the following methods to safely delete files:

- fsutil command



Example

```
# Export the contents of the table my_table to a CSV file.
> psql -c "COPY my_table TO 'C:\WINDOWS\Temp\my_table.csv' (FORMAT CSV)" postgres
# Encrypt the exported file.
> C:\OpenSSL-Win32\bin\openssl enc -e -aes256 -in C:\WINDOWS\Temp\my_table.csv -out my_table.csv.enc
(The user is prompted to enter the passphrase to be used for encryption)

# Check the size of plain text files, and delete them after zero padding
> dir C:\WINDOWS\Temp\my_table.csv
> fsutil file setzerodata offset=0 length=7 C:\WINDOWS\Temp\my_table.csv
> del C:\WINDOWS\Temp\my_table.csv

# Decrypt the encrypted files on other systems.
> C:\OpenSSL-Win32\bin\openssl enc -d -aes256 -in my_table.csv.enc -out my_table.csv
(The user is prompted to enter the passphrase to be used for decryption)
```

If you use COPY FROM to import data to tables and indexes in an encrypted tablespace, the imported data is automatically encrypted before being stored.

5.9 Encrypting Existing Data

You cannot encrypt existing unencrypted tablespaces. In addition, you cannot change encrypted tablespaces so that they do not encrypt.

As an alternative, transfer the tables and indexes to other tablespaces. You can use the following SQL commands for this.

```
ALTER TABLE table_name SET TABLESPACE new_tablespace;
ALTER INDEX index_name SET TABLESPACE new_tablespace;
ALTER DATABASE database_name SET TABLESPACE new_tablespace;
```



See

Refer to "SQL Commands" under "Reference" in the PostgreSQL Documentation for information on SQL commands.

5.10 Operations in Cluster Systems

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

5.10.1 HA Clusters that do not Use Database Multiplexing

Take the following points into account when using transparent data encryption in an HA cluster environment that does not use database multiplexing.

Placement and automatic opening of the keystore file

There are two alternatives for placing the keystore file:

- Sharing the keystore file
- Placing a copy of the keystore file

Sharing the keystore file

This involves using the same keystore file on the primary server and the standby server.

As the standby server is not active while the primary server is running, this file would not be accessed simultaneously, and therefore, it can be shared.

To manage the keystore file in a more secure manner, place it on the key management server or the key management storage isolated in a secure location.

Enable the automatic opening of the keystore on both the primary and standby servers.

Placing a copy of the keystore file

This involves placing a copy of the primary server keystore file on the standby server.

You can do this if you cannot prepare a shared server or disk device that can be accessed from both the primary and standby servers.

However, if you change the master encryption key and the passphrase on the primary server, you must copy the keystore file to the standby server again.

To manage the keystore file in a more secure manner, prepare the key management server or the key management storage isolated in a secure location for both the primary and standby servers, and place the keystore files there.

Enable the automatic opening of the keystore on both the primary and standby servers. Note that copying the automatically opening keystore file (keystore.aks) to the standby server does not enable the automatic opening of the keystore.

5.10.2 Database Multiplexing Mode

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

Placing the keystore file

Place a copy of the primary server keystore file on the standby server.

This is required as the keystore file cannot be shared, and both servers may need to access it simultaneously.



.....
To manage the keystore file in a more secure manner, place it on the key management server or the key management storage isolated in a secure location. A keystore used by both the primary and standby servers can be managed on the same key management server or key management storage.

However, create different directories for the keystores to be used by the primary server and the standby server. Then copy the keystore for the primary server to the directory used on the standby server.

.....

Automatically opening the keystore

You must enable automatic opening of the keystore.

To do this, enable automatic opening of the keystore in all servers that make up database multiplexing. The settings for automatic opening of the keystore include information unique to each server, so simply copying the file does not enable it.

Changing the passphrase

Changes to the passphrase are reflected in all servers that make up database multiplexing, so no special operation is required.

Building and starting a standby server

Before using the `pg_basebackup` command or `pgx_rcvall` command to build a standby server, copy the keystore file from the primary server to the standby server. When using an automatically opening keystore, use the copied keystore file to enable automatic opening on the standby server.

Open the keystore each time you start the standby server. This step is necessary for decrypting and restoring encrypted WAL received from the primary server. To open the keystore, specify the `--keystore-passphrase` option in the `pg_ctl` command or `pgx_rcvall` command and enter the passphrase, or use an automatically opening keystore.

If specifying `--keystore-passphrase` in the `pg_ctl` command, refer to "5.3 Opening the Keystore" for details.

Changing the master encryption key and the passphrase

Change the master encryption key and the passphrase on the primary server. You need not copy the keystore from the primary server to the standby server. You need not even restart the standby server or reopen the keystore. Changes to the master encryption key and the passphrase are reflected in the keystore on the standby server.



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

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

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

5.11 Security-Related Notes

- Decrypted data is cached in the database server memory (shared buffer). As a result, unencrypted data is stored in a minidump, which is the process memory dump. You should, therefore, safely delete the memory dump. You can safely delete files by using the following command:
 - `fsutil` command
- Unencrypted data may be written from the database server memory to the operating system's swap area. To prevent leakage of information from the swap area, consider either disabling the use of swap area or encrypting the swap area using a full-disk encryption product.
- The content of the server log file is not encrypted. Therefore, in some cases the value of a constant specified in a SQL statement is output to the server log file. To prevent this, consider setting a parameter such as `log_min_error_statement`.
- When executing an SQL function that opens the keystore and modifies the master encryption key, ensure that the SQL statement containing the passphrase is not output to the server log file. To prevent this, consider setting a parameter such as `log_min_error_statement`. If you are executing this type of SQL function on a different computer from the database server, encrypt the communication between the client and the database server with SSL.
- Starting with FEP 10, logical replication is available, which allows non-backed up clusters to subscribe to databases where transparent data encryption is enabled. Logical replication does not need to have the same encryption strategy between publisher and subscriber.

In this scenario, if the user wants to encrypt the subscribed copy of data as well, then it is the user's responsibility to create encryption policies to the subscribed databases. By default, published encrypted tablespace data will not be encrypted in the subscriber side.

5.12 Tips for Installing Built Applications

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

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

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

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

```
SET tablespace_encryption_algorithm = 'AES256';  
CREATE TABLESPACE crm_tablespace LOCATION 'C:\crm\data';
```

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


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

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

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

Chapter 6 Policy-based Login Security

To apply a policy for login security to a user, define the policy as a profile and assign the profile to the user. The contents of the profile are saved as database objects.



The items that can be set for a profile are described below.

Managing Dormant Users

You can automatically lock users who have not been connected to the database for a long time.

This setting is for users with the LOGIN attribute.

Managing Policies When Using Password Authentication

You can set the following policies for users who use password authentication (password, md5, scram-sha-256).

- Set a password life time
- Restrict password reuse
- Lock accounts that have failed to login continuously
- Allow passwords to be set in encrypted form
- Set the gradual password rollover time

Gradual password rollover is when you change a password and then keep the old password in effect for a while.

This setting specifies the valid period.

This is useful, for example, if it is difficult to make a new password available system-wide instantly.

6.1 Advance Preparation

Ensure that the postgres database exists and that you can connect to it.

If no connections are possible, and if using a database other than the default postgres database, specify the name of the database to which you can connect in the `userprofile_database` parameter in the `postgresql.conf` file.

Refer to "[Parameters for the Policy-based Login Security](#)" for parameters.

6.2 Changing the Contents of the default Profile

When you create a database user with `CREATE ROLE`, the database user is assigned the default profile. When you create the default profile, all password configuration and authentication restrictions and the gradual password rollover feature are disabled, so change the parameter values to suit your policy. Change the value with the `pgx_alter_profile` function.

Users with `CREATEROLE` privilege can modify the contents of a profile.

Refer to "[B.3.1 Profile Management Functions](#)" for functions.

Refer to "6.7 Profile parameters" for more information on parameters.

[Example]

```
SELECT pgx_alter_profile('default',
'{
  "INACTIVE_USER_TIME": 30,
  "PASSWORD_LIFE_TIME": 50,
  "PASSWORD_GRACE_TIME": 10,
  "PASSWORD_REUSE_MAX": 5,
  "PASSWORD_LOCK_TIME": 0.5,
  "PASSWORD_ALLOW_HASHED": true,
  "PASSWORD_ROLLOVER_TIME": 0.125
}');
```

6.3 Creating and Assigning Profiles

If you want to apply a different policy than the default profile, create a new profile and assign it to the user.

The `pgx_create_profile` function creates a profile.

To assign a profile to a user, use the `pgx_assign_profile_to_user` function.

Users with `CREATEROLE` privilege can create and assign profiles.

Refer to "B.3.1 Profile Management Functions" and "B.3.2 User Management Functions" for functions.

6.4 Actions to be Taken in the Event of Deviation from Policy

When a user is locked

Locked in the following cases:

- Restriction by `INACTIVE_USER_TIME`
- Restriction by `FAILED_LOGIN_ATTEMPTS`

If you want to allow a locked user to login again, unlock the user. Unlocking is performed using the `pgx_unlock_user` function. Operations can be performed by users with `CREATEROLE` privilege.

Refer to "B.3.2 User Management Functions" for functions.

[Example]

```
SELECT pgx_unlock_user('user1');
```

To release a lock when a database administrator is locked, users other than the locked user with `CREATEROLE` privilege can do so. If the user does not exist, Fujitsu Enterprise Postgres must be started in single-user mode and the lock released.



Point

.....
You can set `PASSWORD_LOCK_TIME` to automatically release locks due to `FAILED_LOGIN_ATTEMPTS`.
.....

When the password expires

When a password life time is over, the password expires and cannot be used to connect to and operate on the database until the password is changed.

The password can be changed by the expired user itself, or by a user who has the `CREATEROLE` privilege and who is an administrator of that user (who has `ADMIN` privilege for that user).

6.5 Settings in Streaming Replication Configuration

For streaming replication, the policy is enabled on both the primary and standby servers.

Perform changing the contents of the profile, creating and assigning profiles, unlock, and change the password can only be done on the primary server.

This section describes how to setting streaming replication configuration.

1. Grant privilege to users for streaming replication

The standby server connects to the primary server and propagates any state changes it detects to the primary server. The used users for streaming replication (specified in `primary_conninfo`) is used to connect for state change. Grant privilege to this user. Allow the user to SET ROLE directly to the group role `pgx_update_profile_status`.

[Example]

```
# GRANT pgx_update_profile_status TO repluser WITH SET TRUE;
```

This action enables users for streaming replication to lock all other database users or unlock the lock state due to policy deviations. Also, membership in the `pgx_update_profile_status` group role should only be granted to users for streaming replication.

2. Add a record to `pg_hba.conf`

Add a record to accept connections from the standby server. Duplicate the record for streaming replication and specify the database name in the `userprofile_database` parameter as the destination database name. Authentication methods other than password authentication are recommended.

```
host replication repluser standbyServerAddress authenticationMethod # For Streaming
Replication
host postgres repluser standbyServerAddress authenticationMethod # For state change
propagation
```

Information

- If the primary and standby servers are disconnected, profile-based restrictions still apply on each server. However, the standby server itself cannot change its password or unlock it explicitly. After removing the cause of the disconnect, restore the connection to the primary server and business operations, and then change the password or explicitly unlock.
 - Connection information for reflecting status changes to the primary server or standby server is output as "User profile status sender" to `application_name` in the `pg_stat_replication` view. When monitoring the streaming replication status with the `pg_stat_replication` view, exclude the "User profile status sender" line from the monitoring target.
-

6.6 Backup and Recovery

Backup

Profile contents, user and profile assignment status, and password history can be backed up with physical backups and the `pg_dumpall` command.

When dumping with SQL statements using the `pg_dumpall` command, database role information is backed up at the same time it is dumped.

Information

When the `pg_dumpall` command is run without the `--no-role-passwords` option, the database user's password is included in the backup file in the form of an encrypted form. Therefore, when recovering the database user's password, temporarily set this parameter to true in the default profile of the restore destination. If the restore is successful, the contents of the default profile are restored to the contents at the time of the backup.

Recovery

If you recover using older backup data, you may be considered a dormant user at the time of recovery, or your password may have expired. In this case, unlock or change the password after recovery.

6.7 Profile parameters

Details of the parameters set in the profile are explained.

INACTIVE_USER_TIME

Number of days before auto-locking users who have not been connected to the database for a long time (users who cannot see their sessions)

[Supported values]

integer: A INTEGER value greater than or equal to 1

The unit is days. The maximum value is 24855 days.

DEFAULT: The value of the same parameter in the default profile

UNLIMITED: No auto lock

This setting is for users with the LOGIN attribute.

If a session to the database cannot be verified for at least INACTIVE_USER_TIME, the user is locked out and cannot log in.

Sessions are checked into the database every hour and the information is saved.

If the system shuts down before the save, it is assumed that there were no logins between the last save and the shutdown. As a result, the lock period might be shorter than the value specified in this parameter.

When using streaming replication, the session confirmation to the database is performed every hour on each server. The primary server makes a locking decision while the standby server sends session confirmation information to the upstream server. This can cause a time lag of several hours before the session confirmation information from the standby server is communicated to the primary server. This lag can result in locking even when connected to a standby server.

You can determine when each user's session was last seen by the system by looking at the value of the userprlastactivetime column in the pgx_user_profile system catalog.

If the device is locked by the setting of this parameter, use the pgx_unlock_user function to unlock the device. Automatic cancellation by PASSWORD_LOCK_TIME is not performed.

When a logged-in user changes to another role with SET ROLE, the new role is not considered logged in to the database. It is assumed that the old role is still logged in to the database.

PASSWORD_LIFE_TIME

Number of days the same password can be used for authentication

[Supported values]

numeric: A NUMERIC value greater than or equal to 0

The unit is days. Hours and seconds can be specified with decimal places (e.g. 4.5 is equivalent to "4 days and 12 hours"). Precision is 1 second. The maximum value is 24855 days.

DEFAULT: The value of the same parameter in the default profile

UNLIMITED: No life time (same password can be used indefinitely)

The password life time will be over after PASSWORD_LIFE_TIME days from the last time the password was updated for the target user. The timing of updating the profile is not the starting point. Therefore, if you specify an extremely short number of days (such as 1 day), it may already be past the life time at the time of renewal.

It is possible to specify when a password becomes invalid using the VALID UNTIL clause of CREATE ROLE or ALTER ROLE. If you specify both the VALID UNTIL clause and PASSWORD_LIFE_TIME, both values are valid. Note that in this case, you can login only if both constraints are met.

If there is no grace period, the password expires when the password life time is over. If you login using password authentication in this state, you will receive a "password expired" warning and you will not be able to execute commands other than changing your password. You can change the password to resume normal operations. Other than password authentication, you can connect to and work with the database.

For streaming replication, users will not be able to connect to the standby server after the password life time is over. There is no grace period for password expiration on standby servers. Password changes must be made on the primary server.

PASSWORD_GRACE_TIME

The number of days after a password life time is over before the password expires.

[Supported values]

numeric: A NUMERIC value greater than or equal to 0

The unit is days. Hours and seconds can be specified with decimal places (e.g. 4.5 is equivalent to "4 days and 12 hours"). Precision is 1 second. The maximum value is 24855 days.

DEFAULT: The value of the same parameter in the default profile

UNLIMITED: Indefinite period

The password expiration grace period is `PASSWORD_GRACE_TIME` days from the first login time after the password life time is over. You can login with your current password during the grace period, but a warning prompts you to change your password. You can operate normally except that a warning is displayed. If specified as **UNLIMITED**, the grace period is infinite and you will be warned to change your password at every login. Once transitioned to the grace period, even if you change the profile value after that, you cannot go back to before the grace period, and you cannot change the life time. If 0 is specified, there is no grace period and the password expires at the first login after the password life time is over. In this case, you can perform normal operations again by changing the password.

Because the standby server does not have a grace period before the password expires, users cannot connect to the standby server after the password life time is over. Password changes must be made on the primary server.

PASSWORD_REUSE_TIME

Number of days the same password cannot be reused

The password cannot be reused by the same user for this period from the time the password is updated.

This parameter must be set in combination with `PASSWORD_REUSE_MAX`.

[Supported values]

numeric: A NUMERIC value greater than or equal to 0

The unit is days. Hours and seconds can be specified with decimal places (e.g. 4.5 is equivalent to "4 days and 12 hours"). Precision is 1 second. The maximum value is 24855 days.

DEFAULT: The value of the same parameter in the default profile

UNLIMITED: Not reusable (However, if `PASSWORD_REUSE_MAX` is also **UNLIMITED**, the password can be reused without restriction)

PASSWORD_REUSE_MAX

Number of password changes required before the same password can be reused

This parameter must be set in combination with `PASSWORD_REUSE_TIME`.

[Supported values]

integer: An INTEGER value greater than or equal to 0

DEFAULT: The value of the same parameter in the default profile

UNLIMITED: Not reusable (However, if `PASSWORD_REUSE_TIME` is also **UNLIMITED**, the password can be reused without restriction)

Both `PASSWORD_REUSE_TIME` and `PASSWORD_REUSE_MAX` constraints must be met to reuse passwords.

For example, if you specify `PASSWORD_REUSE_TIME = 30`, `PASSWORD_REUSE_MAX = 10`, a certain password can be reused if 30 days have passed since the update time and the password has been updated 10 times or more. When password update fails due to these parameters, the `SQLSTATE` will be "22023: invalid_parameter_value".

Also, if one parameter has a value and the other parameter specifies `UNLIMITED`, the password cannot be reused. If you want to use only one condition for reuse judgment, you need to set the value of the unused parameter to 0. However, if `UNLIMITED` is specified for both, these parameters are ignored and passwords can be reused without restriction.

When changing a password, if the password is specified in `MD5` or `SCRAM` encrypted format, you cannot check that the password is being reused.

PASSWORD_ALLOW_HASHED

Whether to allow passwords to be specified in `MD5` or `SCRAM` encrypted form when changing passwords. Allow if true.

[Supported values]

boolean: true or false

DEFAULT: The value of the same parameter in the default profile

If true, there are no restrictions on how to set or change passwords.

However, note the following when changing passwords in encrypted form:

- Unable to check for password reuse (`PASSWORD_REUSE_TIME`, `PASSWORD_REUSE_MAX`)
- Password rollover disabled (`PASSWORD_ROLLOVER_TIME`)
- Inability to check password complexity using the extension `passwordcheck`

If false, the password can only be set or changed by specifying the password in clear text in the `PASSWORD` clause of a `CREATE ROLE` or `ALTER ROLE` statement. You cannot change the password using the `psql` command `\password` meta-command.

In this case, all password checks that cannot be performed in encrypted form are possible.



Passwords specified in `CREATE ROLE` or `ALTER ROLE` statements may, depending on configuration, be logged in the `psql` command history and the server log.

FAILED_LOGIN_ATTEMPTS

Number of consecutive failed login attempts allowed by the user

[Supported values]

integer: An `INTEGER` value greater than 0

DEFAULT: The value of the same parameter in the default profile

`UNLIMITED`: Indefinite period (Can fail any number of times)

If password authentication fails consecutively for the number of times specified by this parameter, the user is locked and cannot login.

The number of failed login attempts is counted separately on each server.

PASSWORD_LOCK_TIME

Number of days after a user is locked due to consecutive login failures before the user is unlocked

[Supported values]

numeric: A `NUMERIC` value greater than or equal to 0

The unit is days. Hours and seconds can be specified with decimal places (e.g. 4.5 is equivalent to "4 days and 12 hours"). Precision is 1 second. The maximum value is 24855 days.

DEFAULT: The value of the same parameter in the default profile

UNLIMITED: Locked indefinitely

If set to UNLIMITED, the user will be locked indefinitely and will not be automatically unlocked. Unlocking requires an explicit `pgx_unlock_user` function call by a user with CREATEROLE privilege.

PASSWORD_ROLLOVER_TIME

Number of days after password change before old password expires

[Supported values]

numeric: A NUMERIC value greater than or equal to 0

The unit is days. Hours and seconds can be specified with decimal places (e.g. 4.5 is equivalent to "4 days and 12 hours"). Precision is 1 second. The maximum value is 60 days. The minimum value is 0.0416 days (approximately 1 hour).

DEFAULT: The value of the same parameter in the default profile

From the time the user changes the password until PASSWORD_ROLLOVER_TIME has elapsed, the user can login using the old password.

Specify a value equal to or less than the lesser of the PASSWORD_GRACE_TIME (except when 0 is specified) or PASSWORD_LIFE_TIME. If a larger value is specified, the smaller of these parameters is used. For example, if PASSWORD_GRACE_TIME is 10 and PASSWORD_LIFE_TIME is 50, specify 10 or less for this parameter. If 15 is specified in this parameter, 10 is assumed to be specified.

If 0 is specified, the user cannot login with the old password. The default profile has an initial value of 0.

To expire the combinable period immediately before PASSWORD_ROLLOVER_TIME expires, a user or a user with CREATEROLE privilege can call the `pgx_make_password_rollover_expire` function.

If you execute the ALTER ROLE statement with PASSWORD NULL, password authentication is disabled. Therefore, you cannot log in with the old password.

In either of the following cases, the password rollover is disabled and the previous password expires immediately:

- When changing the password, the password is specified in MD5 or SCRAM encrypted format.
 - The password method has changed since the last time the password was set.
- The relevant parameters are:

- password_encryption
- scram_iterations

6.8 Worker Processes

This feature allows up to two background worker processes to reside. This information is displayed in the statistics view `pg_stat_activity` when it needs to be processed by a worker process. The worker processes added by this feature are those with column `backend_type` values of "user profile status writer" and "user profile status sender".

These processes may wait on the wait event "UserProFileWorkerMain" of type "Activity".

Chapter 7 Data Masking

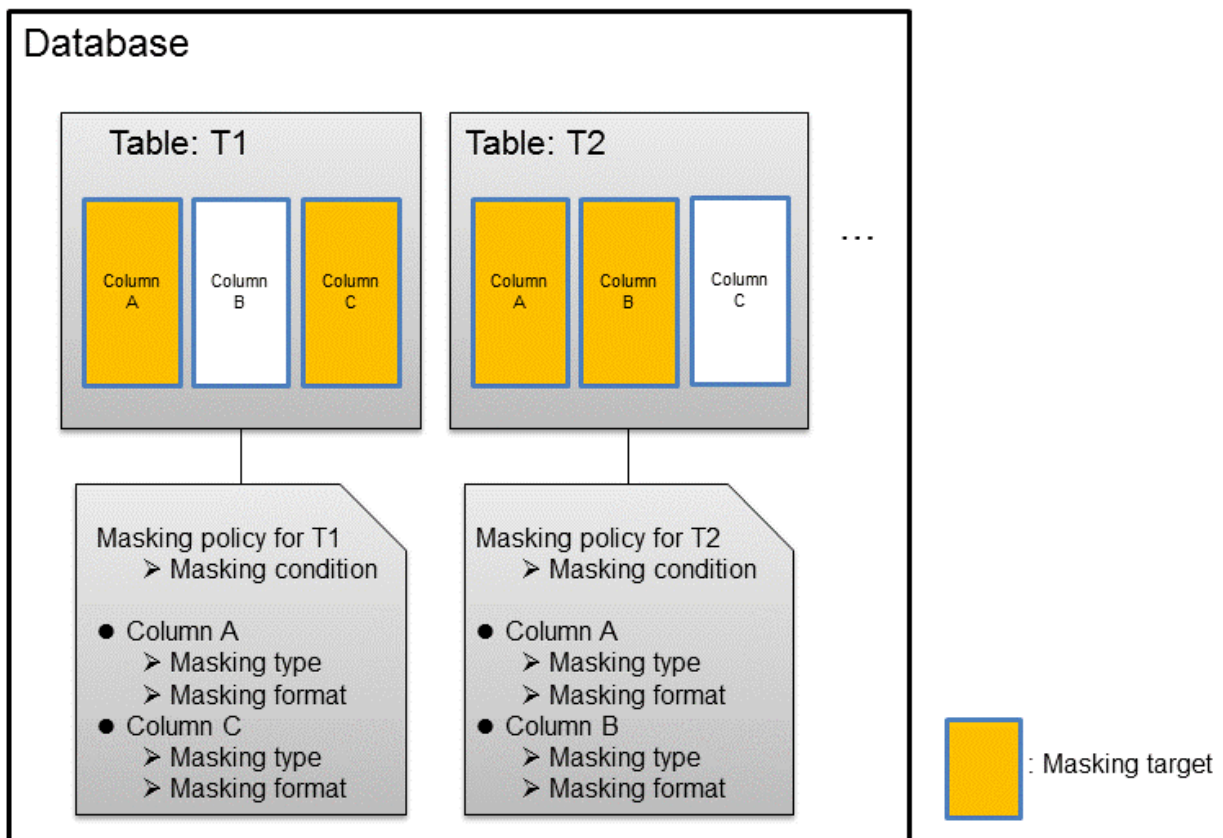
Data masking is a feature that can change the returned data for queries generated by applications, so that it can be referenced by users. For example, for a query of employee data, digits except the last four digits of an eight-digit employee number can be changed to "*" so that it can be used for reference.

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.

7.1 Masking Policy

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

Figure 7.1 Masking policy



Information

When a masking policy is defined, the search performance for the corresponding table may deteriorate.

7.1.1 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

Information

- 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.

7.1.2 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.

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.

The following masking types can be specified and selected depending on the masking target data type. When masking the data that includes multibyte characters, do not specify partial masking for masking type. The result may not be as expected.

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.

7.1.3 Masking Condition

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

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

Information

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

7.1.4 Masking Format

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

Full masking

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




See

Refer to "7.3 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 <code>function_parameters</code>
Numeric type	<p><i>'replacementCharacter, startPosition, endPosition'</i></p> <ul style="list-style-type: none"> - <i>replacementCharacter</i>: Specify the number to display. Specify a value from 0 to 9. - <i>startPosition</i>: 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>. <p> Example</p> <p>Specify as below to change the values from the 1st to 5th digits to 9.</p> <pre>function_parameters := '9, 1, 5'</pre> <p>In this example, if the original data is "123456789", it will be changed to "999996789".</p>
Character type	<p><i>'inputFormat, outputFormat, replacementCharacter, startPosition, endPosition'</i></p> <ul style="list-style-type: none"> - <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. - <i>endPosition</i>: Specify the position of "V" as an end position of masking. When working out the end position, do not include positions of "F". For example, to specify the position of the 11th "V" from the left, specify 11. Specify a positive integer that is greater than <i>startPosition</i>. <p> Example</p> <p>Specify as below to mask a telephone number other than the first three digits using *.</p> <pre>function_parameters := 'VVVFVVVVFVVVV, VVV-VVVV-VVVV, *, 4, 11'</pre> <p>In this example, if the original data is "012-3156-7890", it will be changed to "012-****-****".</p>
Date/timestamp type	'MDYHMS'

Category	Method of specifying function_parameters
	<ul style="list-style-type: none"> - 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. <p> Example</p> <p>Specify as below to mask hour, minute, and second and display 00:00:00.</p> <pre>function_parameters := 'MDYh0m0s0'</pre> <p>In this example, if the original data is "2022-02-10 10:10:10", it will be changed to "2022-02-10 00:00:00".</p>

 **See**

- Refer to "B.4.2 pgx_create_confidential_policy" for information on function_parameters.
- Refer to "7.3 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.

When column data type is `character(n)` or `char(n)` and if the string length after change exceeds `n`, the extra characters will be truncated and only characters up to the `n`th character will be displayed.

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

 **Example**

Specify as below to change all three characters starting from b to X.

```
regexp_pattern := 'b..'
regexp_replacement:= 'X'
regexp_flags := 'g'
```

In this example, if the original data is "foobarbaz", it will be changed to "fooXX".



See

- Refer to "POSIX Regular Expressions" in the PostgreSQL Documentation and check pattern, replacement, and flags for information on the values that can be specified for `regex_pattern`, `regex_replacement`, and `regex_flags`.
- Refer to "7.3 Data Types for Masking" for information on the data types for which masking can be performed.

7.2 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.
 - Be sure to add "pgx_datamasking" to the beginning of the parameter "shared_preload_libraries".
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
```



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. Only database superusers can configure masking policies.

The masking policy setting is explained in the following steps using an example.

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

7.2.1 Creating a Masking Policy

Create a masking policy. 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.

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

1. Create a masking policy

Execute the `pgx_create_confidential_policy` system management function to create a masking policy.

The following values are configured in this example.

- Masking target: Numeric type c1
- Masking type: FULL
- Masking condition: '1=1'

```
postgres=# select pgx_create_confidential_policy(table_name := 't1', policy_name := 'p1',
expression := '1=1', column_name := 'c1', function_type := 'FULL');
pgx_create_confidential_policy
-----
t
(1 row)
```

2. Confirm the displayed data

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

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



See

- Refer to "[B.4.2 pgx_create_confidential_policy](#)" for information on the `pgx_create_confidential_policy` system management function.

7.2.2 Changing a Masking Policy

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

1. Change a masking policy

Execute the `pgx_alter_confidential_policy` system management function to change a masking policy.

The following values are changed in this example.

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

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



```

pgx_enable_confidential_policy
-----
t
(1 row)

```

2. Confirm the displayed data

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

```

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

```

3. Enable a masking policy

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

```

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

```

4. Confirm the displayed data

Confirm that the masking target data has been correctly changed.

```

postgres=# select * from t1;
 c1 |      c2
-----+-----
  0 | 012-****-****
  0 | 012-****-****
  0 | 012-****-****
(3 row)

```



See

- Refer to "[B.4.4 pgx_enable_confidential_policy](#)" for information on the `pgx_enable_confidential_policy` system management function.

7.2.5 Deleting a Masking Policy

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

1. Delete a masking policy

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

```

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

```

2. Confirm the displayed data

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

```

postgres=# select * from t1;
 c1 |      c2
-----+-----

```

1		012-3456-7890
2		012-3456-7891
3		012-3456-7892
(3 row)		



- Refer to "B.4.3 [pgx_drop_confidential_policy](#)" for information on the `gx_drop_confidential_policy` function.

7.3 Data Types for Masking

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

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

Category	Data type	Masking type		
		Full masking	Partial masking	Regular expression masking
Numeric type	smallint	Y	Y	N
	integer	Y	Y	N
	bigint	Y	Y	N
	decimal	Y	Y	N
	numeric	Y	Y	N
	float	Y	Y	N
	real	Y	Y	N
	double precision	Y	Y	N
Character type	character varying(<i>n</i>)	Y	Y	Y
	varchar(<i>n</i>)	Y	Y	Y
	character(<i>n</i>)	Y	Y	Y
	char(<i>n</i>)	Y	Y	Y
Date/timestamp type	date	Y	Y	N
	timestamp	Y	Y	N

7.4 Security Notes

- Starting with FEP 10, logical replication is available, which allows non-backed up clusters to subscribe to databases where data masking policies are enabled. Logical replication allows publisher and subscriber databases to have their own or the same data masking policies.

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

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

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

Chapter 8 Periodic Operations

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

8.1 Configuring and Monitoring the Log

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

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

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

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



See

- Refer to "Error Reporting and Logging" under "Server Administration" in the PostgreSQL Documentation for information on logs.
- Refer to "Configuring Parameters" in the Installation and Setup Guide for Server for information on log settings when operating with WebAdmin.

8.2 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

8.2.1 Monitoring Disk Usage

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

- fsutil volume diskfree command

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

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



Information

If you are using WebAdmin for operations, a warning is displayed when disk usage reaches 80%

8.2.2 Securing Free Disk Space

Secure free disk space by using the following operating system commands to delete unnecessary files, other than the database, from the same disk unit.

- del command

You can also secure disk space by performing the following tasks periodically:

- To secure space on the data storage destination disk:

Execute the REINDEX statement. Refer to "8.5 Reorganizing Indexes" for details.

- To secure space on the backup data storage destination disk:

Execute backup using WebAdmin or the pgx_dmpall command.

8.3 Automatically Closing Connections

If an application stops responding and abnormally terminates for any reason, the connection from the application may remain active on the database server. If this situation continues for an extended period, other applications attempting to connect to the database server may encounter an error, or an error indicating that the tables are unavailable may occur.

It is, therefore, recommended that idle connections be closed automatically at regular intervals.

Set the following parameters in the postgresql.conf file to indicate the time permitted to elapse before a connection is closed.

Parameter name	Setting	Description
tcp_keepalives_idle	Time until keepalive is sent (seconds) If 0, the default value of the system is used.	Sends keepalive to an idle connection at the specified interval in seconds It is recommended to specify 30 seconds.
tcp_keepalives_interval	keepalive send interval (seconds) If 0, the default value of the system is used.	Sends keepalive at the specified interval It is recommended to specify 6 seconds.



Note

The maximum number of connections allowed is 125, unless the desktop heap setting is changed.



See

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

8.4 Monitoring the Connection State of an Application

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

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



See

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

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

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

- [8.4.1 Using the View \(pg_stat_activity\)](#)
- [8.4.2 Using pgAdmin](#)

8.4.1 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 backend_type = 'client backend' and state='idle in
transaction' and current_timestamp > cast(query_start + interval '60 minutes' as timestamp);
-[ RECORD 1 ]-----+-----
datid          | 13003
datname        | db01
pid            | 4638
leader_pid     |
usesysid      | 10
username       | fsep
application_name | apl01
client_addr    | 192.33.44.15
client_hostname |
client_port    | 27500
backend_start  | 2022-02-24 09:09:21.730641+09
xact_start     | 2022-02-24 09:09:23.858727+09
query_start    | 2022-02-24 09:09:23.858727+09
state_change   | 2022-02-24 09:09:23.858834+09
wait_event_type | Client
wait_event     | ClientRead
state          | idle in transaction
backend_xid    |
backend_xmin   |
query_id      |
query         | begin;
backend_type   | client backend
```

See


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

8.4.2 Using pgAdmin

This section describes the procedure for monitoring connections using [Server Status] in pgAdmin.

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

2. In the [Dashboard] tab, identify client connections that have been in the waiting state for an extended period.



8.5 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. 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.

Check the disk usage of the data storage destination using the method described in "8.2 Monitoring Disk Usage and Securing Free Space".



See

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



Point

Typically, reorganize indexes once a month at a suitable time such as when conducting database maintenance. Use SQL statements to check index usage. If this usage is increasing on a daily basis, adjust the frequency of recreating the index as compared to the free disk space.

The following example shows the SQL statements and the output.

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

[SQL statements]

```
SELECT
  nspname AS schema_name,
  relname AS index_name,
  round(100 * pg_relation_size(indexrelid) / pg_relation_size(indrelid)) / 100 AS index_ratio,
  pg_size_pretty(pg_relation_size(indexrelid)) AS index_size,
  pg_size_pretty(pg_relation_size(indrelid)) AS table_size
FROM pg_index I
  LEFT JOIN pg_class C ON (C.oid = I.indexrelid)
  LEFT JOIN pg_namespace N ON (N.oid = C.relnamespace)
WHERE
  C.relkind = 'i' AND
  pg_relation_size(indrelid) > 0
ORDER BY pg_relation_size(indexrelid) DESC, index_ratio DESC;
```

[Output]

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



See

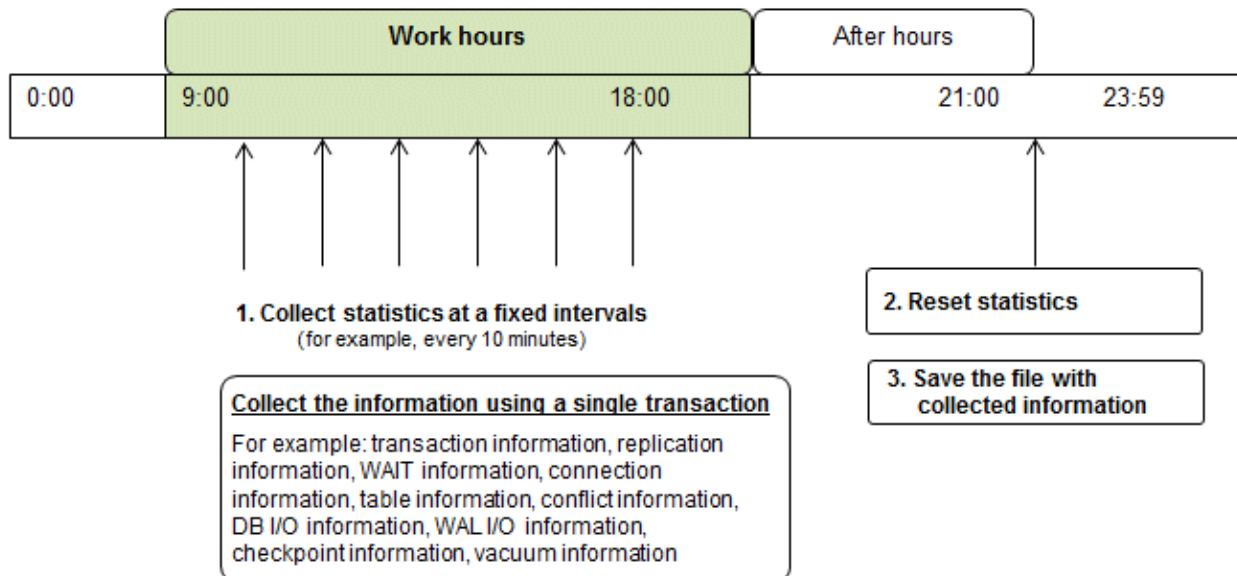
Refer to "Notes on Application Compatibility" in the Application Development Guide for information on maintaining application compatibility.

8.6 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 8.1 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 "8.6.1 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 "8.6.3 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

8.6.1 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 8.1 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 " D.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 " D.3 pgx_stat_latch " for details.
pgx_stat_walwriter	Displays statistics related to WAL writing, in a single line. Refer to " D.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 " D.5 pgx_stat_sql " for details.
pgx_stat_gmc	Displays statistics related to Global Meta Cache hit ration and used memory size. Refer to " D.6 pgx_stat_gmc " for detail. Also refer to Chapter 13 Global Meta Cache for information on the Global Meta Cache.

8.6.2 Collection Configuration

The procedure for configuring collection depends on the information content.

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

Information common to PostgreSQL



See

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

Information added by Fujitsu Enterprise Postgres

Information added by Fujitsu Enterprise Postgres is collected by default.

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

View name	Parameter
<code>pgx_stat_lwlock</code>	<code>track_waits (*1)</code>
<code>pgx_stat_latch</code>	
<code>pgx_stat_sql</code>	<code>track_sql</code>
<code>pgx_stat_gmc</code>	<code>track_gmc</code>

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.

8.6.3 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
<code>pg_stat_reset_shared(text)</code>	void	<p>Reset some cluster-wide statistics counters to zero, depending on the argument (requires superuser privileges).</p> <p>Calling <code>pg_stat_reset_shared('lwlock')</code> will zero all counters shown in <code>pgx_stat_lwlock</code>.</p> <p>Similarly, in the following cases, all values of the pertinent statistics counter are reset:</p> <ul style="list-style-type: none">- If <code>pg_stat_reset_shared('latch')</code> is called: All values displayed in <code>pgx_stat_latch</code>- If <code>pg_stat_reset_shared('walwriter')</code> is called: All values displayed in <code>pgx_stat_walwriter</code>- If <code>pg_stat_reset_shared('sql')</code> is called: All values displayed in <code>pgx_stat_sql</code>

Function	Type of return value	Description
		- If <code>pg_stat_reset_shared('gmc')</code> is called: All values except size column in <code>pgx_stat_gmc</code>

 See

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

8.7 Scheduling of an Aggressive Freeze for Tuples (VACUUM FREEZE)

An architectural limitation of PostgreSQL is that it requires transaction ID reclamation via tuple freezing to avoid the problem of disruptive transaction wraparound. For more information about transaction wraparound issues, refer to "Preventing Transaction ID Wraparound Failures" in the PostgreSQL Documentation.

Tuple freezing works with autovacuum and `VACUUM FREEZE`(aggressive freeze for tuples). However, on a system that consumes transactions quickly, autovacuum alone might be enough to run into wraparound problems. This is because autovacuum is slow to reduce frequency of collisions with applications and increase system load. Such systems must perform an aggressive tuple freeze at the appropriate time.

Fujitsu Enterprise Postgres provides scripts to perform efficient aggressive freeze for tuples.

This section describes how to monitor whether the freezing process is on time in a running system, and how to estimate the running time of the script.

 Point

If you insert or update a large amount of data at once, the next freeze process will take a long time, so be sure to perform aggressive freeze for tuples after such operations.

8.7.1 Monitoring Trends in Transaction ID Usage


Execute the following SQL to monitor the trend in transaction ID usage.

```
# SELECT max(age(datfrozenxid)) FROM pg_database WHERE dataallowconn = true;
```

If autovacuum does not sufficiently reduce transaction ID usage, schedule aggressive freeze for tuples.

Determine whether the amount of waste reduced by automatic vacuum is sufficient

Looking at the entire interval that includes several autovacuum freeze processes, the transaction ID usage will be graphed like a sawtooth, sandwiching the `autovacuum_freeze_max_age` setting value, depending on the start and end of the autovacuum operation. If the usage peaks are gradually getting larger, the reduction may be insufficient. In particular, if it no longer falls below `autovacuum_freeze_max_age`, the reduction is definitely insufficient.



8.7.2 How to Schedule Aggressive Freeze for Tuples

1. Determine a time period with low business load when aggressive freeze for tuples will have minimal impact.
2. Copy the script that implements aggressive freeze for tuples and set the connection information, multiplicity, and timeout value parameters.

For streaming replication configuration, specify multiple connection destinations in PGHOST and PGPORT, and specify "read-write" or "primary" in PGTARGETSESSIONATTRS to connect to the primary server.

The scripts are stored in the following location under the Fujitsu Enterprise Postgres installation directory:

```
<Install directory>\share\execute_freeze.ps1.sample
```

3. Register the script from step 2 in the job scheduler.
Specify the execution timing according to the planned time period.

Example: Use the Task Scheduler to run a script called execute_freeze.ps1 at 2:05 AM

```
schtasks /create /tn "ExecuteVacuumFreeze" /tr "powershell -ExecutionPolicy Bypass -windowstyle hidden -File 'C:\path\to\execute_freeze.ps1'" /sc daily /st 02:05 /f
```

When a schedule review is required

Even after aggressive freeze for tuples is scheduled, continue to monitor the trend in transaction ID usage. If you determine that the freezing process is not keeping up, the schedule must be reviewed. To collect the statistical information required for the review, enable the extended functionality for aggressive freeze for tuples using the following procedure.

1. Modify the shared_preload_libraries parameter in postgresql.conf

```
shared_preload_libraries = 'pgx_stat_vacuum_freeze'
```

2. Execute the following SQL

```
SELECT datname FROM pg_database WHERE dataallowconn = true;
```

Execute the following CREATE EXTENSION for all databases to be output and define the extension.

```
CREATE EXTENSION pgx_stat_vacuum_freeze;
```

Monitor the trends in transaction usage, and if you determine that the freezing process is not keeping up, refer to "[8.7.3 Tuning the Allocation Time for Aggressive Freeze for Tuples](#)" to calculate the time required for freezing tuples and add more allocation time.

8.7.3 Tuning the Allocation Time for Aggressive Freeze for Tuples

The interval between autovacuum runs is treated as one interval, and aggressive freeze for tuples is scheduled to be completed before the next autovacuum run. For example, if your environment is one week where freezing by autovacuum occurs, determine the required time within that interval.

Calculate the time required for aggressive freeze for tuples as follows, based on the time it took to freeze transaction IDs by past autovacuum stored in the [pgx_stat_freeze_results](#) table for each database you are using.

In this example, the interval for autovacuum freezing is set to one week. Perform this for each database you operate and calculate the total.

```
# SELECT sum(total_elapsed_time) - sum(sleep_times) as total_time FROM pgx_stat_freeze_results WHERE
measurement_start_time >= '2024/10/21 00:00:00' AND measurement_end_time < '2024/10/28 00:00:00' AND
autovacuum = true;
```

If you are unable to secure sufficient time when the workload is low, increase the degree of parallelism. You can specify the degree of parallelism in the script.

8.8 Database Monitoring with Amazon CloudWatch

Amazon CloudWatch is a monitoring tool provided by Amazon Web Services (hereinafter referred to as AWS). Fujitsu Enterprise Postgres works with Amazon CloudWatch to collect database-related metrics and logs. The collected metrics can be used to monitor the status of the database and create alert notifications when thresholds are exceeded. In addition, by analyzing the collected logs on CloudWatch, it is possible to detect abnormalities in logins and operations.

8.8.1 Metrics Collection

Metrics collection uses psql commands to retrieve metrics and the AWS CLI to forward metrics to Amazon CloudWatch.

1. Preparation

Refer to the AWS documentation, install the AWS CLI, and configure it to transfer metrics to Amazon CloudWatch.

2. Create a metrics collection script

Sample scripts for collecting metrics are stored in the following location under the Fujitsu Enterprise Postgres installation directory. Create a script for collecting metrics by referring to the sample script. For the metrics collected by the sample script, refer to "[8.8.3 Metrics Collected by the Sample Script](#)".

```
<Install directory>\share\putting-metrics-to-monitoring-tool.psql.sample
```

Edit the following connection information defined in the sample script.

- Amazon CloudWatch connection destination to store metrics

Define the Amazon CloudWatch Namespace to collect metrics in "NAMESPACE".

```
$NAMESPACE="YOUR Cloud Watch NAMESPACE"
```

- Database connection destination for collecting metrics

Edit the variables below to define the PostgreSQL information for retrieving metrics.

PGHOST: Specify the host name to connect to.

PGPORT: Specify the port to connect to.

PGDATABASE: Specify the database name to connect to.

PGUSER: Specify the PostgreSQL user to use for connection. To view all statistical information for all sessions, specify a superuser or a role with the built-in role `pg_read_all_stats` privileges.

3. Setting environment variables

To run the metrics collection script, you must set the path where the Fujitsu Enterprise Postgres server feature and AWS CLI are installed in the environment variable `PATH`.

The `PATH` environment variable can be added to the environment variables of the user running the metrics collection script, or added within the script.

The following is an example configuration when the AWS CLI is installed in `C:\Program Files\Amazon\AWSCLIV2`.

```
set PATH = "C:\Program Files\Fujitsu\fspev<x>server64\bin;C:\Program Files\Amazon\AWSCLIV2;%PATH%"
```

When setting environment variables within a script, make sure to write them at the beginning of the script, before the environment variable for the database connection destination.

```
$env:PATH = "C:\Program Files\Fujitsu\fsepv<x>server64\bin;C:\Program Files\Amazon\AWSCLIV2;" +
$env:PATH
```

4. Define the database connection credentials

If necessary, define the password and the storage location of the SSL communication certificate in an environment variable or configuration file.

5. Periodically run metrics collection scripts

Use a task scheduler or similar to periodically run the metrics collection script as the OS user who can transfer metrics using the aws cloudwatch put-metric-data command.

Below is an example of a task scheduler definition for collecting metrics every 10 minutes.

```
# Task Scheduler screen
Right-click on "Task Scheduler Library"
-> Select Create basic task

# Definition when creating a task
1. Create a basic task
Enter "Name" and "Description" and click "Next"
2. Trigger
Select "Daily" and click "Next"
3. Daily
Start: Any date
Interval: 1 day
Select and click "Next"
4. Operation
Select "Start a program" and click "Next"
5. Start a program
Enter the following and click "Next"
Program/script: C:\Windows\System32\cmd.exe
Add arguments: /c "C:\Program Files\PowerShell\7\pwsh.exe" -ExecutionPolicy Bypass -File
"( Script placement location)\putting-metrics-to-monitoring-tool.ps1" >> "( Log storage location)
\put-metrics.log" 2>&1"
6. Completed
Check "Open the Properties dialog for this task when I click Finish" and click "Finish"

# Properties screen
1. General tab
Check "Run whether user is logged on or not"
Define the user who will run the task with "Change User or Group"
2. Trigger tab
Repeat every: Check the box and select 10 minutes from the dropdown
```

In this example, the script's standard output is output to the log file "put-metrics.log", allowing you to detect whether the script is running normally or abnormally. Note that in this example, the log file continues to be output, so we recommend using a task scheduler or similar to periodically rotate and delete the logs.

The metrics collection interval should be determined based on the purpose of collection. The following are examples of collection intervals.

- Long intervals, such as once every 10 minutes or once an hour, are recommended when monitoring to check operational trends.
- Short intervals, such as once a minute, are recommended when monitoring to quickly respond to increases in load, such as a sudden increase in the number of connections.

If the collection interval is short, the load may increase because the number of queries to the database increases accordingly. Consider the business load when deciding on the execution interval and the queries to be executed for metric collection.

8.8.2 Log Collection

You can collect logs using the CloudWatch agent. For instructions on how to collect logs, see the Amazon CloudWatch manual.

8.8.3 Metrics Collected by the Sample Script

In the sample script, metrics items are created from the following perspectives to regularly check whether the database cluster is operating properly.

- Is the database capacity sufficient?
- Are your SQL performance and connections adequate?
- Are each node operating normally?

Table 8.2 List of metrics collected by the sample script

Metrics name	Metrics
pg_stat_bgwriter_*	Maps displayed in statistics collection
pg_stat_database_*	Maps displayed in statistics collection
pg_stat_database_conflicts_*	Maps displayed in statistics collection
pg_stat_archiver_*	Maps displayed in statistics collection
pg_stat_activity_*	Maps displayed in statistics collection
pg_stat_replication_*	Maps displayed in statistics collection
pg_replication_slots_*	Mapping to the pg_replication_slots system catalog
pg_locks_*	Mapping to the pg_locks system catalog
pg_capacity_connection_*	Connection metrics (e.g. txns running for 1 hour)
pg_capacity_schema_*	Schema disk space metrics
pg_capacity_tblspace_*	Tablespace disk space metrics
pg_capacity_tblvacuum_*	Metrics for tables that haven't been vacuumed in days
pg_capacity_longtx_*	Number of transactions running for more than 5 minutes
pg_capacity_datfrozenxid_*	Trends in transaction ID usage (Used for scheduling aggressive freeze for tuples (VACUUM FREEZE))
pg_performance_locking_detail_*	Blocked process details
pg_performance_locking_*	Number of blocked processes
pg_stat_user_tables_*	Vital statistics from pg_stat_user_tables
pg_statio_user_tables_*	Vital statistics from pg_statio_user_tables
pg_stat_statements_*	Statistics for SQL statements executed by the server
pg_capacity_tblbloat_*	Table fetch bloat

8.9 Monitoring Deferred SQL and Periodically Backing up statistics

After production starts, SQL execution can be delayed due to a variety of factors, including unexpected system load and data bias. You can detect delays early by monitoring for delays that affect operations. Monitoring can also reduce the time required for recovery tasks, such as determining the cause and resolving delays.

One cause of deferred SQL is a change in statistics. Some delays caused by statistics can be resolved by restoring the statistics before the delay occurred. For this reason, back up your statistics periodically to guard against delays.

Operational workflow

The workflow required to monitor delayed SQL and back up statistics is shown below.

Corresponding time	Work item	Work details
During environment construction	(1) Enabling the pg_dbms_stats extension	Enable the pg_dbms_stats extension with the Create Extension statement.
	(2) Logging settings for delayed SQL monitoring	Set the log_min_duration_statement parameter etc. in postgresql.conf.
During system development	(3) Creating and reviewing a script for backing up statistics (*1)	Provide a script to back up statistics according to the statistics and Vacuum operation.
During operation	(4) Periodic backup of statistics	Execute the script created in (3) to back up the statistics.
	(5) Monitoring delayed SQL and adjusting the allowable delay time	Monitor the occurrence of delayed SQL and adjust the allowable delay time to find the optimal settings for operation.
When delayed SQL occurs	(6) Investigating the cause of delayed SQL	If a delayed SQL occurs, we will investigate the cause of the SQL delay and consider countermeasures.
	(7) Restoring statistics (*2)	Delayed SQL is eliminated by restoring and fixing the backed up statistics.

*1: Although it is possible to perform the work during operation, please consider the impact on operation and perform the work before starting operation.

*2: If it is determined that the cause is not statistics, take appropriate measures according to the cause.

The specific work content is shown below.

(1) Enabling the pg_dbms_stats extension

To back up and restore statistics, you need the export and import functions of the pg_dbms_stats extension, so enable the pg_dbms_stats extension.

For instructions on how to enable it, refer to "Setting Up pg_dbms_stats" in the Installation and Setup Guide for Server.

(2) Logging settings for delayed SQL monitoring

Database logs are used to monitor delayed SQL.

Make the following settings in postgresql.conf. Note that the date and time are important when investigating the cause of delayed SQL, set the format and file name to include the date and time.

Parameter	Specified value
logging_collector	on
log_line_prefix	Specify the following format and add information required for investigation, such as the time and executed application. [%t]%u %d %p[%l]
log_filename	Since you will need to check past execution records, set it so that you can see the date and time. postgresql.%Y-%m-%d
log_min_duration_statement	300000

.....

If the value specified for the `log_min_duration_statement` parameter is small, even SQL that does not need to be considered delayed will be output to the log, and if it is too large, delayed SQL cannot be detected. Therefore, if the setting value cannot be estimated, set it small at the start of operation, and adjust it to a larger value as necessary while monitoring during operation.

.....

(3) Creating and reviewing a script for backing up statistics

To back up statistics, use the export function of the `pg_dbms_stats` extension.

The backup is used not only for restoration but also for identifying the cause of delayed SQL, so the time the backup was performed is important. Therefore, be sure to create a script like the one below so that the backup time is known.

We also provide a sample script that uses the export function to export on a database basis. This allows users to respond more flexibly by making modifications themselves, such as making changes on a schema basis, so copy the script to the execution current before using it.

The sample script is stored below. Also, "`<x>`" indicates the product version.

```
c:\Program Files\Fujitsu\fssepv<x>server64\OSS\pg_dbms_stats \doc\extension\export_plain_stats-  
<x>.sql.sample
```

The following is an example of a file created when it is copied to the current execution directory as "export_plain_stats-<x>.sql".

```
echo off  
SET CURRENTDIR=d:\pg_dbms_stats\backup  
  
pushd "%CURRENTDIR%"  
  
Rem make filename  
SET NOWDATE=%date:~0,4%%date:~5,2%%date:~8,2%  
SET TMPTIME=%time: =0%  
SET NOWTIME=%TMPTIME:~0,2%%TMPTIME:~3,2%  
EXECDATE=%NOWDATE%%NOWTIME%  
set FILENAME=pg_dbms_stats_back.%EXECDATE%.dmp  
  
rem export  
psql -d database1 -f export_plain_stats-<x>.sql  
  
rem rename dump file  
ren export_stats.dmp ${FILENAME}  
popd
```

(4) Periodic backup of statistics

Execute the script created in "(3) Creating and reviewing a script for backing up statistics" to back up the statistics information on a regular basis.

Backup execution interval

Perform backup at the following timings according to the statistics update method.

When statistics is automatically updated (autovacuum is enabled)

Statistics are updated when autovacuum is run.

Therefore, use the OS's cron command or schedule function of the task manager to periodically execute the export function and back up the statistics to a file.

Therefore, use the OS's cron command or a scheduler such as task manager to periodically execute the script created in "(3) Creating and reviewing a script for backing up statistics" and back up the statistics information to a file.

Obtain the shortest autovacuum execution interval from the log, and back up the statistics at an interval shorter than the obtained autovacuum execution interval. autovacuum logs can be obtained by setting the `log_autovacuum_min_duration` parameter to 0.

When users update statistics

If the user controls the updating of statistics using the ANALYZE command, etc., back up the statistics at the same time as updating the statistics.

Backup retention period

The backed up statistics will be used as investigation material to determine the cause of delayed SQL if a performance problem occurs. Therefore, keep it for the period of time expected to be required from the occurrence of the operational problem to its resolution.

Point

- The capacity required for backup mainly depends on the number of objects (schemas, tables, etc.) contained in the database. Estimate the required capacity from the backup file when checking the operation of the script and secure this space.
- To restore backed up statistics, the names of objects such as schemas and tables in the backup source must match. If you change the name of an object during operation, change the backup file name or storage location so that you can check whether it matches the object definition when restoring statistics.

(5) Monitoring delayed SQL and adjusting the allowable delay time

Monitors SQL that is delayed beyond the time set in the log_min_duration_statement parameter.

The settings in this section will be output in the following format.

```
[2024-04-30 10:20:11 JST]user1 postgres 3414[1]LOG: duration: 301001.541 ms statement: SELECT
pg_sleep(301)
[2024-04-30 10:26:12 JST]user1 postgres 3414[1]LOG: duration: 302002.321 ms statement: SELECT
pg_sleep(302)
```

When detected SQL is deemed to be delayed more than expected

Refer to "[\(6\) Investigating the cause of delayed SQL](#)" to determine the cause and take appropriate measures.

When SQL with execution time that does not affect operation is detected

The current setting just outputs redundant logs, so make major changes to the log_min_duration_statement parameter and then reload to enable it.

(6) Investigating the cause of delayed SQL

There are many factors that can cause SQL delays, so you must first identify the cause.

Here we will show you how to investigate SQL delays caused by changes in statistics.

Note that in the following cases, delays caused by statistics are unlikely, so recommend investigating from a different perspective (I/O, system load, etc.).

- When the SQL statement does not contain a search condition
- When the SQL statement uses the pg_hint_plan extension

1. Check the delayed SQL from the server log and identify the schema used by that SQL.
2. In a database cluster (such as a development environment) separate from the production environment, create an environment (hereafter referred to as the reproduction environment) with the same database name and table definition as the environment where the problem occurred. The reproduction environment also requires the pg_dbms_stats extension, so enable the pg_dbms_stats extension.
3. Use the pg_dump command from the production environment to obtain the target table definition. Since the data is not needed at this time, use the -s option of the pg_dump command from the production environment and use the extracted table definition.

The following is an example of output to the file ddl_schema.dmp when the schema to be investigated is schema_1 in database database1.

```
pg_dump -d database1 -s schema_1 > ddl_schema.dmp
```

Using `ddl_schema.dmp`, create a research schema named `schema_1` in the database named `database1` in the reproduction environment.

```
psql -d database1 -f ddl_schema.dmp
```

- In order to identify the cause of the SQL delay, use the backed up statistics to check whether there is a cause of the delay in the statistics. Import the statistics into the schema in the reproduction environment using the `pg_dbms_stats` import function (`dbms_stats.import_schema_stats`).

The following is an example of restoring statistics for schema `schema_1` from the backup file `"pg_dbms_stats_back.20240422011223.dmp"`.

```
psql -d database1 -c
"SELECT dbms_stats.import_schema_stats('schema_1', 'pg_dbms_stats_back.20240422011223.dmp')"
```

Then, execute the `EXPLAIN` command for the deferred SQL to obtain the execution result (query plan).

- Restore the statistics by separately importing the statistics that were backed up during the previous time period when deferred SQL was running stably.

For example, if the target SQL is executed daily as part of regular batch processing, use the data from the same time the previous day. For online processing, use the data from a previous time period when a similar business application is expected to be executed.

The following is an example of restoring the statistics from the backup file `"pg_dbms_stats_back.20240421010001.dmp"` from one day ago, assuming the SQL is used in batch processing.

```
psql -d database1 -c
"SELECT dbms_stats.import_schema_stats('schema_1', 'pg_dbms_stats_back.20240421010001.dmp')"
```

It likewise obtains the result of executing the `EXPLAIN` command on deferred SQL (query plan).

- Compare the execution results (query plans) of 4. and 5., and take the following action based on the results.

- If the query plans are the same or the expected execution time of the query plan in 5. is longer

This may be a problem other than statistics (system load, index imbalance, I/O, etc.), so investigate from a different perspective.

- Other than the above

The problem is likely that the optimal query plan was not selected due to updated statistics.

Refer to "[\(7\) Restoring statistics](#)" and import the statistics from the period when the system was operating stably into the production environment to restore the system.

(7) Restoring statistics

Restore the production environment statistics to the state they were in before the delay occurred. The procedure is the same as for verification, but it is as follows:

```
psql -d database1 -c
"SELECT dbms_stats.import_schema_stats('schema_1', 'pg_dbms_stats_back.20240422000000')"
```

8.10 Performance Tuning

8.10.1 Enhanced Query Plan Stability

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

If you know that the statistics will not change significantly, you can control the selection of the same query plan by directly specifying which query plan to select in `pg_hint_plan`, or by fixing the statistics referenced by the planner in `pg_dbms_stats`.

Also, `pg_dbms_stats` from Fujitsu Enterprise Postgres can fix statistics (height of Btree indexes) that OSS `pg_dbms_stats` cannot account for. Refer to "[8.10.1.1 Fixing the Height of a Btree Index](#)".



See

- For information about setting up `pg_hint_plan` and `pg_dbms_stats`, refer to "`pg_hint_plan`" and "`pg_dbms_stats`" in the Installation and Setup Guide for Server.
- For basic usage of `pg_hint_plan` and `pg_dbms_stats`, refer to below.
 - Control execution plans with `pg_hint_plan`
<https://www.postgresql.fastware.com/postgresql-insider-tun-hint-plan>
 - Control execution plans by fixing statistics with `pg_dbms_stats`
<https://www.postgresql.fastware.com/postgresql-insider-tun-dbms-stt>
- For more information about `pg_hint_plan` and `pg_dbms_stats`, refer to the OSS web page.

8.10.1.1 Fixing the Height of a Btree Index

When PostgreSQL creates a query plan, it also references the height of the Btree index to be used as part of the statistics. However, OSS `pg_dbms_stats` does not include the height of the Btree index in its fixed statistics. The height of a Btree index is also fixed in `pg_dbms_stats` for Fujitsu Enterprise Postgres.

Fixing the height of a Btree index covers Fujitsu Enterprise Postgres 16 and later, so to maintain compatibility with earlier versions, specify something like:

- PostgreSQL 9.2 and earlier compatibility

```
pg_dbms_stats.use_tree_height = off
```

- Fujitsu Enterprise Postgres 15 SP2 and earlier compatibility

```
pg_dbms_stats.lock_tree_height = off
```

See the table below for details on each parameter. If you want the parameter value to be valid for all sessions, specify it in `postgresql.conf` and then reload.

Parameter	Specified value	Default
<code>pg_dbms_stats.use_tree_height</code>	Specify whether to include the Btree index height in cost calculation and plan generation when statistics are fixed. on: Includes Btree index height in cost calculations and plan generation. off: Do not include Btree index height in cost calculation or plan generation.	on
<code>pg_dbms_stats.lock_tree_height</code>	Specify whether to also fix the height of the Btree index when you fix the statistics. Note that this is valid only when the specified value of <code>pg_dbms_stats.use_tree_height</code> is "on". on: Fixes the height of the Btree index at the height when locked. off: The height of the Btree index is not fixed and varies depending on the amount of data in the index.	on

When you want to import statistics from your legacy environment

The format of files handled by the import/export functions of this version of pg_dbms_stats is different from the format of pg_dbms_stats in OSS or Fujitsu Enterprise Postgres 15SP2 or earlier. Therefore, statistics exported from pg_dbms_stats for OSS or Fujitsu Enterprise Postgres 15SP2 or earlier should be imported using the dbms_stats.import_<obj>_stats_no_tree_height function.

How to specify the dbms_stats.import_<obj>_stats_no_tree_height function

```
dbms_stats.import_<obj>_stats_no_tree_height('Absolute path of export file')
```

<obj> indicates a database (currently connected), schema, table, or column.

Import using the dbms_stats.import_<obj>_stats function provided by OSS is not possible.

If you use the wrong import function, an error message containing the following HINT will be output. Check the combination of input file and import function type, and use the correct import function.

```
HINT: The import function may be incorrectly combined with the format of the exported data. Please check the documentation for the relationship between the import function and the available data.
```

Chapter 9 Streaming Replication Using WebAdmin

This chapter describes streaming replication using WebAdmin.

A standby instance for streaming replication can be created from a standalone instance, a master instance, or another standby instance. The standby instance connects to the master instance and uses WAL records to replicate data. The standby instance can be used for read-only operations.

WebAdmin cannot back up standby instances.

Additionally, streaming replication can be set to asynchronous or synchronous mode.


Point

- 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].

For characters that cannot be specified in WebAdmin, refer to “[Appendix J WebAdmin Disallow User Inputs Containing Hazardous Characters](#)”.

9.1 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.
 - [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 I WebAdmin Wallet](#)" for information on creating wallet entries.
 - [Instance name]: Name of the standby database instance to create.
The name must meet the conditions below:
 - Maximum of 16 characters
 - The first character must be an ASCII alphabetic character
 - The other characters must be ASCII alphanumeric characters
 - [Instance port]: Port number of the standby database instance.
 - [Host IP address]: The IP address of the server machine where the standby instance is to be created. This information is needed to configure the standby instance to be connected to the master.
 - [Data storage path]: Directory where the database data will be stored
 - [Backup storage path]: Directory where the database backup will be stored
 - [Transaction log path]: Directory where the transaction log will be stored
 - [Encoding]: Database encoding system
 - [Replication mode]: Replication mode of the standby instance to be created ("Asynchronous" or "Synchronous")

- [Application name]: The reference name of the standby instance used to identify it to the master instance.

The name must meet the conditions below:

- Maximum of 16 characters
- The first character must be an ASCII alphabetic character
- The other characters must be ASCII alphanumeric characters

4. Click  to create the standby instance.


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
ERROR: pgx_rcvall: backup of the database has not yet been performed, or an incorrect backup storage
directory was specified
```


9.2 Promoting a Standby Instance

Streaming replication between a master and standby instance can be discontinued using WebAdmin.

Follow the procedure below to promote a standby instance to a standalone instance, thereby discontinuing the streaming replication.

1. In the [Instances] tab, select the standby instance that needs to be promoted.
2. Click .
3. Click [Yes] from the confirmation dialog box.

The standby instance will be promoted and will become a standalone instance, which is not part of a streaming replication cluster.

Once the standby instance is promoted to become a standalone instance, the backup storage status will be "Error". This is because no backups are available when the instance is newly promoted to a standalone instance. The status will be reset if a new backup is performed by clicking [Solution] or .




9.3 Converting an Asynchronous Replication to Synchronous

You can convert an asynchronous standby instance to a synchronous standby instance.

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.

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.



See

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.

9.4 Converting a Synchronous Replication to Asynchronous

Follow the procedure below to convert a Synchronous standby instance to Asynchronous.

1. In the [Instances] tab, select the master instance of the relevant cluster.
2. Click
3. In the [Streaming replication] section, edit the value for [Synchronous standby names].
 - Remove the "Application name" of the standby instance you want to be in Asynchronous mode.
4. Click
5. Select the master instance and click
6. Select the standby instance. [Instance type] will now show the updated status.



See

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.

9.5 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 I WebAdmin Wallet](#)" for information on creating wallet entries. Replication credential (user name and password) should not contain hazardous characters. Refer to "[Appendix J WebAdmin Disallow User Inputs Containing Hazardous Characters](#)".
6. Enter [Host IP address].

Specify the IP address of the node where the standby instance was created.
7. Click to open the [Join replication cluster] dialog box.

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

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

Chapter 10 Installing and Operating the In-memory Feature

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

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

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

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

This feature can only be used in Advanced Edition.

10.1 Installing Vertical Clustered Index (VCI)

This section describes the installation of VCI.

1. [Evaluating whether to Install VCI](#)
2. [Estimating Resources](#)
3. [Setting up](#)

10.1.1 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

Point

.....

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 "10.1.4 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.
-

10.1.2 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.

10.1.3 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](#)

10.1.3.1 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

Parameter name	Setting value	Description	Required
vci.max_parallel_degree	Maximum number of background workers used for parallel scan	Maximum number of background workers used for parallel scan. Add this value to max_worker_processes.	N
vci.smc_directory	Directory name in which a temporary file is created as the dynamic shared memory	Directory name in which a temporary file is created as the dynamic shared memory during a scan using a VCI.	N

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
vci.smc_directory = 'E:\\vci\\work'
```

Information

If performing operations that use VCI, do not delete the shared library names specified to "shared_preload_libraries" and "session_preload_libraries". If operations that access VCI are performed after the VCI was defined, unexpected behavior may occur.

See

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

10.1.3.2 Installing the Extensions

Execute CREATE EXTENSION to install the VCI and pg_prewarm extensions. Both extensions need to be installed for each database. Also, extensions can only be installed on the public schema by a superuser.

- Installing VCI

```
db01=# CREATE EXTENSION vci;
```

- Installing pg_prewarm

```
db01=# CREATE EXTENSION pg_prewarm;
```

Some operations cannot be performed for VCI extensions. Refer to "[10.2.1 Commands that cannot be Used for VCI](#)" for details.

10.1.3.3 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);
```

Point

- Some table types cannot be specified on the ON clause of CREATE INDEX. Refer to "10.1.4.1 Relation Types" for details.
- Some data types cannot be specified on the column specification of CREATE INDEX. Refer to "10.1.4.2 Data Types" for details.
- Some operations cannot be performed for VCI. Refer to "10.2.1 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.

10.1.3.4 Confirming that the VCI has been Created

Execute the SELECT statement to reference the `pg_indexes` catalog, and confirm that the VCI was created for the target columns.

Example

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

10.1.4 Data that can Use VCI

This section describes on which relation types and for which data types VCIs can be created.

10.1.4.1 Relation Types

VCIs cannot be created on some relation types.

The ON clause of CREATE INDEX described in "10.1.3.3 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

10.1.4.2 Data Types

VCI's cannot be created for some data types.

The column specification of CREATE INDEX described in "10.1.3.3 Creating a VCI" cannot specify a column with data type on which VCI's 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

Category	Data type	Supported by VCI?
	inet	N
	macaddr	N
	macaddr8	N
Bit string	bit(<i>n</i>)	Y
	bit varying(<i>n</i>)	Y
Text search	tsvector	N
	tsquery	N
UUID	uuid	Y
XML	xml	N
JSON	json	N
	jsonb	N
Range	int4range	N
	int8range	N
	numrange	N
	tsrange	N
	tstzrange	N
	daterange	N
Object identifier	oid	N
	regproc	N
	regprocedure	N
	regoper	N
	regoperator	N
	regclass	N
	regtype	N
	regconfig	N
	regdictionary	N
pg_lsn type	pg_lsn	N
Array type	-	N
User-defined type (Basic type, enumerated type, composite type, and range type)	-	N

10.2 Operating VCI

This section describes how to operate VCI.

10.2.1 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, except when the public schema is specified. This operation is not required for VCI.

- Operations that cannot be performed on relations containing a VCI

Command	Subcommand	Description
ALTER INDEX	SET	The subcommands on the left cannot be performed if a VCI is specified.
	SET TABLESPACE	
	ALL IN TABLESPACE	If the operation is required, delete the VCI using DROP INDEX, and re-create it using CREATE INDEX after completing the operation.
ALTER OPERATOR CLASS	RENAME TO	The subcommands on the left cannot be performed if a VCI is specified.
	OWNER TO	
	SET SCHEMA	This operation is not supported in VCI.
ALTER OPERATOR FAMILY	ADD	
	DROP	
	RENAME TO	
	OWNER TO	
ALTER TABLE	ALL IN TABLESPACE <i>name</i> [OWNED BY <i>roleName</i>] SET TABLESPACE <i>newTablespace</i>	A tablespace that contains a VCI cannot be specified. If the operation is required, delete the VCI using DROP INDEX, and re-create it using CREATE INDEX after completing the operation.
	DROP [COLUMN] [IF EXISTS] <i>colName</i> [RESTRICT CASCADE]	A column that contains a VCI cannot be specified. If the operation is required, delete the VCI using DROP INDEX, and re-create it using CREATE INDEX after completing the operation.
	ALTER [COLUMN] <i>colName</i> [SET DATA] TYPE <i>dataType</i> [COLLATE <i>collation</i>] [USING <i>expr</i>]	If the operation is required, delete the VCI using DROP INDEX, and re-create it using CREATE INDEX after completing the operation.
	CLUSTER ON <i>indexName</i>	A VCI cannot be specified.
	REPLICA IDENTITY {DEFAULT USING INDEX <i>indexName</i> FULL NOTHING}	This operation is not supported in VCI.
CLUSTER	-	A table that contains a VCI and VCI cannot be specified.

Command	Subcommand	Description
		If the operation is required, delete the VCI using <code>DROP INDEX</code> , and re-create it using <code>CREATE INDEX</code> after completing the operation.
CREATE INDEX	UNIQUE	The subcommands on the left cannot be performed if a VCI is specified. This operation is not supported in VCI.
	CONCURRENTLY	
	[ASC DESC]	
	[NULLS { FIRST LAST }]	
	WITH	
	WHERE	
	INCLUDE	
CREATE OPERATOR CLASS	-	A VCI cannot be specified. This operation is not supported in VCI.
CREATE OPERATOR FAMILY	-	
CREATE TABLE	EXCLUDE	
DROP INDEX	CONCURRENTLY	The subcommands on the left cannot be performed if a VCI is specified. This operation is not supported in VCI.
REINDEX	-	A VCI cannot be specified. This command is not required as VCI uses daemon's automatic maintenance to prevent disk space from increasing.

Client application command

- Operations that cannot be performed on relations containing a VCI

Command	Description
clusterdb	Clustering cannot be performed for tables that contain a VCI.
reindexdb	VCIs cannot be specified on the <code>--index</code> option.

10.2.2 Data Preload Feature

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

```
db01=# SELECT pgx_prewarm_vci('idx_vci');
```



See

Refer to "[B.5 VCI Data Load Control Function](#)" for information on `pgx_prewarm_vci`.

Chapter 11 Parallel Query

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

- CPU load calculation
- Increase of workers during runtime
- Statistics view displays the action state

11.1 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 checks the available CPU resources when determining the number of workers for a parallel query.

When checking the available CPU resources, the system checks the CPU usage (whether there are sufficient free system resources) at the time the SQL statement is executed, and a parallel query is only planned if there is enough free CPU usage for two or more cores on the server.

For example, if a server has four cores and the total free space from core 1 to core 4 is 200% or more, a parallel query will be planned, but if it is less than 200%, a parallel query will not be planned.

11.2 Increase of Workers during Runtime

Parallel query can cause CPU overload because workers cannot be added, for example, because there are not enough workers or no workers at the start of the query.

This Fujitsu Enterprise Postgres enhancement allows systems to allocate additional workers during query execution (if there are workers available at the time). The ability to increase workers during runtime is available only with parallel query. This improves query performance.

11.3 Statistics View Displays the Action State

You can check the action state by using the existing statistics collector. Monitoring the statistics enables you to check the action state of parallel processing through, for example, changes in the system load.

The following table describes the column added to the statistics views:

pg_stat_all_tables

Column	Data type	Description
parallel_query	bigint	Number of parallel query initialized in the table.

The parallel query column has also been added to the following views:

- pg_stat_sys_tables
- pg_stat_user_tables
- pg_stat_xact_all_tables
- pg_stat_xact_sys_tables
- pg_stat_xact_user_tables



See

Refer to "The Statistics Collector" in "Server Administration" in the PostgreSQL Documentation for information on the statistics collector and views.

Values will be set for the pgx_stat_sql view in Fujitsu Enterprise Postgres if parallel query is run.



See

.....
Refer to "D.5 pgx_stat_sql" for information on pgx_stat_sql view.
.....

Chapter 12 High-Speed Data Load

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

This feature is available only in the Advanced Edition.

Information

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.

12.1 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](#)

12.1.1 Deciding whether to Install

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

Database server memory

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

See

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

Frequency of checkpoints

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

Example

```
LOG:  checkpoints are occurring too frequently (19 seconds apart)
HINT:  Consider increasing the configuration parameter "max_wal_size".
```

12.1.2 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.

Information

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 "[12.1.1 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.

12.1.3 Setup

This section describes how to set up high-speed data load.

Setup flow

1. [Setting Parameters](#)
2. [Installing the Extension](#)

12.1.3.1 Setting Parameters

Set the parameters required for high-speed data load in `postgresql.conf`. After that, start or restart the instance.

The table below lists the `postgresql.conf` parameters that must be changed, and the values that must be added to their current values. After editing `postgresql.conf`, start or restart the instance.

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

As shown in the following example, 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.

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
```

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

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

 See

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

12.1.3.2 Installing the Extension

Execute CREATE EXTENSION to install the high-speed data load extension. The extension needs to be installed on each database. Also, extensions can only be installed on the public schema by a superuser.

 Example

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

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

12.2 Using High-Speed Data Load

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

12.2.1 Loading Data

To load data from a file into a Fujitsu Enterprise Postgres table, execute the pgx_loader command in load mode. The data inserted using this feature is dumped as a COPY command by the pg_dump command and the pg_dumpall command.

 Example

The example below loads the file E:\path\to\data.csv (2000 records) into table tbl using a degree of parallelism of 3.

```
> pgx_loader load -j 3 -c "COPY tbl FROM 'E:\path\to\data.csv' WITH CSV"
LOAD 2000
```

 Point

If an external file contains data that violates the format or constraints, the data load may fail partway through, resulting in delays for routine tasks such as nightly batch processing. Therefore, it is recommended to remove the invalid data before executing the data load.



- 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.

12.2.2 Checking Progress

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



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

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

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

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

Refer to "D.7 `pgx_stat_progress_loader`" for information on the `pgx_stat_progress_loader` view.



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

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

		1000														
9313		222		testdb		333		COPY FROM		FILE		0		0		63000
		1000														
9314		222		testdb		333		COPY FROM		FILE		0		0		63000
		1000														

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

12.2.3 Recovering from a Data Load that Ended Abnormally

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

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

Checking for in-doubt transactions

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

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

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

Example

The example below retrieves the global transaction identifier (gid) of in-doubt transactions performed by the database role myrole and that used table tbl. The retrieved global transaction identifiers pgx_loader:9589 and pgx_loader:9590 identify in-doubt transactions.

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

2. Refer to the pg_prepared_xacts system view.

Check if the in-doubt transactions retrieved above exist.

Example

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

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



See

Refer to "G.1 `pgx_loader_state`" for information on the `pgx_loader_state` table.

Resolving in-doubt transactions

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

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



Example

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

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



Point

The recovery mode of the `pgx_loader` command only resolves transactions prepared by high-speed data load. For transactions prepared by an application using distributed transactions other than this feature, follow the procedure described in "16.13 Actions in Response to Error in a Distributed Transaction".

12.3 Removing High-Speed Data Load

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

12.3.1 Removing the Extension

Execute `DROP EXTENSION` to remove the high-speed data load 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.

The extension needs to be removed on each database. Also, extensions can only be removed on the public schema by a superuser.



Example

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

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

Chapter 13 Global Meta Cache

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

This feature can only be used in Advanced Edition.

13.1 Usage

Describes how to use the Global Meta Cache feature.

13.1.1 Deciding Whether to Enable the Global Meta Cache Feature

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

In particular, consider using Global Meta Cache if the total size of the meta cache for each process exceeds the amount of installed memory, or takes up a large portion of that memory, thereby squeezing memory allocations to the database cache or the Operating system file cache. Using Global Meta Cache may increase the time it takes to execute a single SQL to reference a meta cache on shared memory, but you can expect a greater benefit from being able to allocate more memory, such as for the database cache.

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

13.1.2 Estimating Memory for Global Meta Cache

To enable the Global Meta Cache feature, the `pgx_global_metacache` parameter must specify an upper limit on the size of the shared memory (Hereinafter, the GMC area) dedicated to Global Meta Cache. Ideally, this upper limit should be the size estimated in "[Appendix A Parameters](#)". Values lower than this can still work, but refer to "[13.1.3 How the GMC Memory Area Is Used](#)" on using the GMC area to understand the disadvantages.

13.1.3 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.

13.1.4 Enabling the Global Meta Cache Feature

To Enable the Global Meta Cache feature edit the `postgresql.conf` file and set the `pgx_global_metacache` parameter. Restarting the instance after editing the `postgresql.conf` file is required. Refer to "[Appendix A Parameters](#)" for information on the parameters.

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

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

Example

Here is an example postgresql.conf configuration:

```
pgx_global_metacache = 800 MB
```

Wait Events

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

[GMC Feature Wait Events]

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

Point

If GMCSweep is happened frequently, increase the pgx_global_metacache setting.

See

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

13.1.5 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.

13.2 Statistics

Describes the statistics for the Global Meta Cache feature.

13.2.1 System View

You can check the cache hit ratio and size of the GMC area in the system view pgx_stat_gmc. Refer to "D.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 "8.6 Monitoring Database Activity" in the Operations Guide for information on the statistics.

Chapter 14 Local Meta Cache Limit

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

This feature is available only in the Advanced Edition.

14.1 Usage

Describes how to use the Local Meta Cache Limit feature.

14.1.1 Deciding Whether to Enable the Local Meta Cache Limit Feature

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

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

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

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

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

14.1.2 How to Set Parameters for the Local Meta Cache Limit Feature

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

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



Example

Here is an example postgresql.conf configuration:

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

14.1.3 Cache Removal when Local Meta Cache Limit is Enabled

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

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

In PostgreSQL, in order to acquire memory at high speed, a memory block of a certain size is acquired from the OS, and a small memory is cut out from the block and used. The memory for the metacache is cut out in the same way. Therefore, it is possible to return the memory block to the OS by destroying all the meta caches scattered throughout the memory block. When this happens, the next SQL execution will

be slowed down due to the re-creation of the meta cache. Therefore, upper limit of feature should be set to a value larger than the size of the meta cache used by at least one transaction.

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

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

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

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

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

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

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

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

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

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

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

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

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

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



Example

Here is an example postgresql.conf configuration:


```
pgx_cache_hit_log_interval= 30min
```

Chapter 15 Backup/Recovery Using the Copy Command

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

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

This appendix describes backup/recovery using the copy command.



Point


- 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.

15.1 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.




 **Information**

The backup data storage destination cannot be used as these backup areas used by the copy command.

Backup using the backup information file

The copy command must determine the backup destination on each backup, as it is necessary to cycle through the backup areas. Backup can be automated by using the backup information file, which contains information about the backup destination.



Information

The backup information file is prepared in the backup data storage destination by the `pgx_dmpall` command, and contains information that can be read or updated by the copy command. This file is managed by associating it with the latest backup successfully completed by the `pgx_dmpall` command, so the latest backup information relating to the copy command registered by the user can be retrieved. Additionally, the content of the backup information file can be displayed using the `pgx_rcvall` command.

Configuration of the copy command for backup

The `pgx_dmpall` command calls the copy command for backup after execution for the three modes below. It is therefore necessary for the copy command for backup to implement the required processing for each of the modes. The user can use any method to hand over backup information between modes within the copy command, such as creating temporary files.

- 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.

Configuration of the copy command for recovery

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

- restore mode

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

Point

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

Example

Using a batch file

```
if "%1" == "prepare" goto prepare
if "%1" == "backup" goto backup
if "%1" == "finalize" goto finalize
if "%1" == "restore" goto restore
:prepare
    processingRequiredForPrepareMode
    exit /b
:backup
    processingRequiredForBackupMode
    exit /b
:finalize
    processingRequiredForFinalizeMode
    exit /b
```

```
:restore
  processingRequiredForRestoreMode
exit /b
```

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.bat.sample
```

15.2 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 "[15.4 Copy Command Interface](#)" for details.

Performing backup

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

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

Example

```
> pgx_dmpall -D D:\database\inst1 -Y D:\database\command\backup.bat
--exclude-copy-cluster -P D:\database\command\tablespace_list.txt
```

Point

- To exclude up the database cluster from backup using the copy command, specify the `--exclude-copy-cluster` option.
- To back up only some tablespaces using the copy command, use the `-P` option specifying the full path of the file created in step 1 of preparation for backup.

See

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

Checking backup status

Use the `pgx_rcvall` command to check the backup status.

Execute the `pgx_rcvall` command with the `-l` option specified to output backup data information. If backup was performed using the copy command, the resources backed up using the copy command will also be output.

Example

```
> pgx_rcvall -l -D D:\database\inst1
Date                Status  Dir                                Resources backed up by the copy command
2022-03-01 13:30:40 COMPLETE E:/backup/inst1/2022-03-01_13-30-40 pg_data,dbspace,indexspace
```

15.3 Recovery Using the Copy Command

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

Determining the backup area of the latest backup

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

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

Example

```
> pgx_rcvall -D 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 "15.2 Backup Using the Copy Command".

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

Example

```
> pgx_rcvall -D D:\database\inst1 -B E:\backup\inst1 -Y D:\database\command\recovery.bat
```

Point

If the latest backup was performed using the copy command, the `pgx_rcvall` command automatically recognizes which database resources were backed up using the copy command, or whether resources were backed up to the backup data storage destination. Therefore, recovery can be performed by simply executing the `pgx_rcvall` command specifying the copy command for recovery.

See

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

15.4 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.

15.4.1 Copy Command for Backup

Feature

User exit (for the copy command) called from the `pgx_dmpall` command.

Format

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

If the operation mode is "prepare"

```
copyCommandName prepare "pathOfBackupInfoFile" "pathOfBackupTargetListFile"
```

If the operation mode is "backup"

```
copyCommandName backup
```

If the operation mode is "finalize"

```
copyCommandName finalize "pathOfBackupInfoFile"
```

Argument

- Operation mode

Mode	Description
prepare	Implements the preparation process for backing up using the copy command. Called before the PostgreSQL backup mode is started.
backup	Implements the backup process. Called during the PostgreSQL backup mode.
finalize	Implements the backup completion process. Called after the PostgreSQL backup mode is completed.

- Full path of the backup information file

Full path of the backup information file of the latest backup, enclosed in double quotation marks. If a backup has not been performed, specify '-'.

- Full path of the backup target list file

Full path of the file containing the resources to be backed up using the copy command, enclosed in double quotation marks. One of the following is described in each resource name.

Resource	Description
Database cluster	pg_data
Tablespace	Tablespace name



Example

To back up the database cluster and the tablespaces dbspace and indexspace using the copy command, the file should contain the following:

```
pg_data
dbspace
indexspace
```



Information

The encoding of resource names output to the backup target list file by the pgx_dmpall command is the encoding used when this command connects to the database with auto specified for the client_encoding parameter, and is dependent on the locale at the time of command execution.

The number of arguments vary depending on operation mode. The argument of each operation mode is as follows.

Operation mode	First argument	Second argument	Third argument
prepare	Operation mode	Backup information file path name	Backup target list file path name
backup		None	None
finalize		Backup information file path name	

Additionally, the access permissions for the backup information file and backup target list file are different depending on the operation mode. The access permissions of each operation mode are as follows:

Operation mode	Backup information file	Backup target list file
prepare	Can be viewed	Can be viewed
backup	-	-
finalize	Can be viewed and updated	-

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.

15.4.2 Copy Command for Recovery

Feature

User exit (for the copy command) called from the `pgx_rcvall` command.

Format

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

```
copyCommandName restore "pathOfBackupInfoFile" "pathOfBackupTargetListFile"
```

Argument

- Operation mode

Mode	Description
restore	Performs restore.

- Full path of the backup information file

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

- Full path of the backup target list file

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

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

Backup information file	Backup target list file
Can be viewed	Can be viewed

Return value

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

Description

- The copy command operates with the privileges of the operating system user who executed the pgx_rcvall command. Therefore, grant copy command execution privileges to users who will execute the pgx_rcvall command. Additionally, have the copy command change users as necessary.
- The copy command is called once only in restore mode.
- To copy multiple resources simultaneously, have the copy command copy them in parallel.
- The backup information file and backup target list file cannot be deleted. Additionally, the privileges cannot be changed.
- Standard output and standard error output of the copy command are output to the terminal where the pgx_rcvall command was executed.
- If the copy command becomes unresponsive, the pgx_rcvall command will also become unresponsive. If the status of the copy command is deemed to be unresponsive by the operating system, use an operating system command to forcibly stop it.
- Output the copy command execution trace and the result to a temporary file, so that if it terminates in error, the cause can be investigated at a later time.
- Do not execute Fujitsu Enterprise Postgres commands or PostgreSQL applications in the copy command.
- There may be files and directories not required for recovery using the archive log included in the backup, such as postmaster.pid, pg_wal/*subdirectory* and pg_replslot in the database cluster. If such unnecessary files and directories exist, have the copy command delete them after the restore.

Chapter 16 Actions when an Error Occurs

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

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

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

Even if a disk is not defective, the same input-output error messages, as those generated when the disk is defective, may be output. The recovery actions differ for these error messages.

Check the status of the disk, and select one of the following actions:

- If the disk is defective

Refer to "[16.1 Recovering from Disk Failure \(Hardware\)](#)", and take actions accordingly.

- If the disk is not defective

Refer to "[16.14 I/O Errors Other than Disk Failure](#)", and take actions accordingly.

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

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

Determining the cause of an error

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



See

Refer to "Configuring Parameters" in the Installation and Setup Guide for Server for information on server logs.

Approximate recovery time

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

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

- Data storage destination or transaction log storage destination

```
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

16.1 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:

- [16.1.1 Using WebAdmin](#)
- [16.1.2 Using Server Command](#)

Point

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

16.1.1 Using WebAdmin

Recover a failed disk using WebAdmin.

In a streaming replication configuration, the master instance can be recovered, but the standby instance cannot. If disk failure occurs on a standby instance, it may be necessary to delete and re-create the instance. Also, recovering the master instance stops streaming replication to and from all standby instances. In such an event, the standby instances can be promoted to standalone instances or can be deleted and re-created

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

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

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

1. Stop applications
 - Stop applications that are using the database.
2. Stop the instance
 - Stop the instance. Refer to "[2.1.1 Using WebAdmin](#)" for information on how to stop an instance. WebAdmin automatically stops instances if recovery of the database cluster is performed without stopping the instance.
3. Recover the failed disk
 - Replace the disk, and then recover the volume configuration information.
4. Create a tablespace directory
 - If a tablespace was defined after backup, create a directory for it.
5. Recover the keystore, and enable automatic opening of the keystore
 - Do the following if the data in the database has been encrypted:
 - Restore the keystore to its state at the time of the database backup.
 - Enable automatic opening of the keystore.
6. Recover the database cluster
 - Log in to WebAdmin, and in the [Instances] tab, click [Solution] for the error message in the lower-right corner.
7. Run recovery
 - In the [Restore Instance] dialog box, click [Yes].

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

8. Resume applications

Resume applications that are using the database.

 **Point**

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

If this happens, refer to "[16.10.3 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.

 **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.

16.1.2 Using Server Command

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

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

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

1. Stop applications

Stop applications that are using the database.

2. Stop the instance

Stop the instance, refer to "[2.1.2 Using Commands](#)" for details.

If the instance fails to stop, refer to "[16.11 Actions in Response to Failure to Stop an Instance](#)".

3. Recover the failed disk

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

4. Create a storage destination directory

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

In [Properties] in Windows(R) Explorer, set appropriate permissions so that only the instance administrator can access the storage destination directory. (Refer to [Help and Support] in Windows(R) for information on [Properties].)



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

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

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

6. Recover the database cluster

Recover the database cluster using the backup data.

Specify the following in the `pgx_rcvall` command:

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

Example

```
> pgx_rcvall -D D:\database\inst1 -B E:\backup\inst1
```

If recovery fails

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

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



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

```
FATAL: the database system is starting up
```

7. Start the instance

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

8. Resume applications

Resume applications that are using the database.

If failure occurred on the backup data storage disk

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

There are two methods of taking action:

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

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

No	Step	Instance stopped	
		No	Yes
1	Confirm that transaction log mirroring has stopped	Y	N
2	Stop output of archive logs	Y	N
3	Stop applications	N	Y
4	Stop the instance	N	Y
5	Recover the failed disk	Y	Y
6	Create a backup data storage destination directory	Y	Y
7	Resume output of archive logs	Y	N
8	Resume transaction log mirroring	Y	N
9	Start the instance	N	Y
10	Run backup	Y	Y
11	Resume applications	N	Y

Y: Required

N: Not required

The procedure is as follows:

If an instance has not been stopped

1. Confirm that transaction log mirroring has stopped

Use the following SQL function to confirm that transaction log mirroring has stopped.

```
postgres=# SELECT pgx_is_wal_multiplexing_paused();
pgx_is_wal_multiplexing_paused
-----
t
(1 row)
```

If transaction log mirroring has not stopped, then stop it using the following SQL function.

```
postgres=# SELECT pgx_pause_wal_multiplexing();
LOG:  multiplexing of transaction log files has been stopped
pgx_pause_wal_multiplexing
-----
(1 row)
```

2. Stop output of archive logs

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

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

- Changing archive_command

Specify a command that will surely complete normally, so that archive logs will be regarded as having been output.

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

- Reload the configuration file

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

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

3. Recover the failed disk

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

4. Create a backup data storage destination

Create a backup data storage destination.

In [Properties] in Windows(R) Explorer, set appropriate permissions so that only the instance administrator can access the backup data storage destination directory. (Refer to [Help and Support] in Windows(R) for information on [Properties].)

Refer to "[3.2.2 Using Server Commands](#)" for information on how to create a backup data storage destination.

5. Resume output of archive logs

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

6. Resume transaction log mirroring

Execute the `pgx_resume_wal_multiplexing` SQL function.

Example

```
SELECT pgx_resume_wal_multiplexing()
```

7. Run backup

Use the `pgx_dmpall` command to back up the database cluster.

Specify the following value in the `pgx_dmpall` command:

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

Example

```
> pgx_dmpall -D D:\database\inst1
```

If an instance has been stopped

1. Stop applications

Stop applications that are using the database.

2. Stop the instance

Stop the instance. Refer to "[2.1.2 Using Commands](#)" for details.

If the instance fails to stop, refer to "[16.11 Actions in Response to Failure to Stop an Instance](#)".

3. Recover the failed disk

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

4. Create a backup data storage destination

Create a backup data storage destination.

In [Properties] in Windows(R) Explorer, set appropriate permissions so that only the instance administrator can access the backup data storage destination directory. (Refer to [Help and Support] in Windows(R) for information on [Properties].)

Refer to "[3.2.2 Using Server Commands](#)" for details.

5. Start the instance

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

6. Run backup

Use the pgx_dmpall command to back up the database cluster.

Specify the following value in the pgx_dmpall command:

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

Example

```
> pgx_dmpall -D D:\database\inst1
```

7. Resume applications

Resume applications that are using the database.



See

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

16.2 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:

- [16.2.1 Using WebAdmin](#)
- [16.2.2 Using the pgx_rcvall Command](#)



Point

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

16.2.1 Using WebAdmin

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

Refer to "16.1.1 Using WebAdmin" for details.

16.2.2 Using the pgx_rcvall Command

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

Follow the procedure below to recover the data storage disk.

1. Stop applications

Stop applications that are using the database.

2. Stop the instance

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

If the instance fails to stop, refer to "16.11 Actions in Response to Failure to Stop an Instance".

3. Confirm the backup date and time

Pinpoint a date and time prior to the data corruption based on the content of the job log or event log.

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

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

5. Recover the database cluster

Use the `pgx_rcvall` command to recover the database cluster.

Specify the following values in the `pgx_rcvall` command:

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

Example

In the following examples, "March 20, 2022 10:00:00" is specified as the recovery time.

```
> pgx_rcvall -D D:\database\inst1 -B E:\backup\inst1 -e "2022-03-20 10:00:00"
```

If recovery fails

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.

Information

The following message displayed during recovery is output as part of normal operation of `pgx_rcvall` command (therefore the user does not need not be concerned).

```
FATAL: the database system is starting up
```

6. Start the instance

Start the instance. Refer to "2.1.2 Using Commands" for information on how to start an instance.

If necessary, re-execute transaction processing from the specified recovery time, and then resume database operations.

7. Resume applications

Resume applications that are using the database.

See

Refer to "pgx_rcvall" in the Reference for information on the `pgx_rcvall` command.

16.3 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:

- [16.3.1 Using WebAdmin](#)
- [16.3.2 Using the `pgx_revall` Command](#)

Point

- Back up the database cluster after recovering it. Backup deletes obsolete archive logs (transaction logs copied to the backup data storage destination), freeing up disk space and reducing the recovery time.
 - 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.
-

16.3.1 Using WebAdmin

Recover data to a backup point using WebAdmin.

In a streaming replication configuration, the master instance can be recovered, but the standby instance cannot. If disk failure occurs on a standby instance, it may be necessary to delete and re-create the instance. Also, recovering the master instance stops streaming replication to and from all standby instances. In such an event, the standby instances can be promoted to standalone instances or can be deleted and re-created.

Follow the procedure below to recover the data in the data storage disk.

1. Stop applications

Stop applications that are using the database.

2. Stop the instance


Stop the instance. Refer to "[2.1.1 Using WebAdmin](#)" for information on how to stop an instance.

3. Recover the keystore, and enable automatic opening of the keystore

Do the following if the data in the database has been encrypted:

- Restore the keystore to its state at the time of the database backup.
- Enable automatic opening of the keystore.

4. Recover the database cluster

Log in to WebAdmin, and in the [Instances] tab, select the instance to be recovered and click .

5. Recover to the backup point

In the [Restore Instance] dialog box, click [Yes].

Recovery is performed. An instance is automatically started when recovery is successful.

6. Resume database operations

If necessary, re-execute transaction processing from the backup point to when an erroneous operation was performed, and then resume database operations.

16.3.2 Using the pgx_rcvall Command

The `pgx_rcvall` command recovers database clusters to the restore point created with the server command. Refer to "Setting a restore point" in "3.2.2 Using Server Commands" for information on how to create a restore point.

Follow the procedure below to recover the data in the data storage disk.

1. Stop applications

Stop applications that are using the database.

2. Stop the instance

Stop the instance. Refer to "2.1.2 Using Commands" for information on how to stop an instance.

If the instance fails to stop, refer to "16.11 Actions in Response to Failure to Stop an Instance".

3. Confirm the restore point

Use a restore point recorded in an arbitrary file, as explained in "3.2.2 Using Server Commands", to determine a restore point prior to the erroneous operation.

4. Recover the keystore, and enable automatic opening of the keystore

When the data in the database has been encrypted, restore the keystore to its state at the time of the database backup. Configure automatic opening of the keystore as necessary.

5. Recover the database cluster

Use the `pgx_rcvall` command to recover the database cluster.

Specify the following values in the `pgx_rcvall` command:

- Specify the data storage destination in the `-D` option. If the `-D` option is omitted, the value of the `PGDATA` environment variable is used by default.
- Specify the backup data storage destination in the `-B` option.
- The `-n` option recovers the data to the specified restore point.

Example

The following example executes the `pgx_rcvall` command with the restore point "batch_20220303_1".

```
> pgx_rcvall -D D:\database\inst1 -B E:\backup\inst1 -n batch_20220303_1
```

If recovery fails

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.



Information

The following message displayed during recovery is output as part of normal operation of `pgx_rcvall` command (therefore the user does not need not be concerned).

```
FATAL: the database system is starting up
```

6. Start the instance

Start the instance.

Refer to "2.1.2 Using Commands" for information on how to start an instance.

7. Restart operation of the database

If necessary, re-execute transaction processing from the specified recovery time to the point when an erroneous operation was performed, and then resume database operations.



Refer to "pgx_rcvall" in the Reference for information on the pgx_rcvall command.

16.4 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 "16.4.1 When using the view (pg_stat_activity)")
- pgAdmin (refer to "16.4.2 Using pgAdmin")

Use the system management function (pg_terminate_backend) to disconnect connections.

16.4.1 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 (<x>) (*1)
Type "help" for help.
```

*1: <x> indicates the PostgreSQL version on which this product is based.

2. Close connections from clients that have been in the waiting state for an extended period.

Use pg_terminate_backend() to close connections that have been trying to connect for an extended period.

However, when considering continued compatibility of applications, do not reference or use system catalogs and functions directly in SQL statements. Refer to "Notes on Application Compatibility" in the Application Development Guide for details.

Example

The following example closes connections where the client has been in the waiting state for at least 60 minutes.

```
select pid,username,application_name,client_addr,pg_terminate_backend(pid) from pg_stat_activity
where backend_type = 'client backend' and state='idle in transaction' and current_timestamp >
cast(query_start + interval '60 minutes' as timestamp);
-[ RECORD 1 ]-----+-----
pid                | 4684
username           | fsepuser
application_name   | ap11
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.

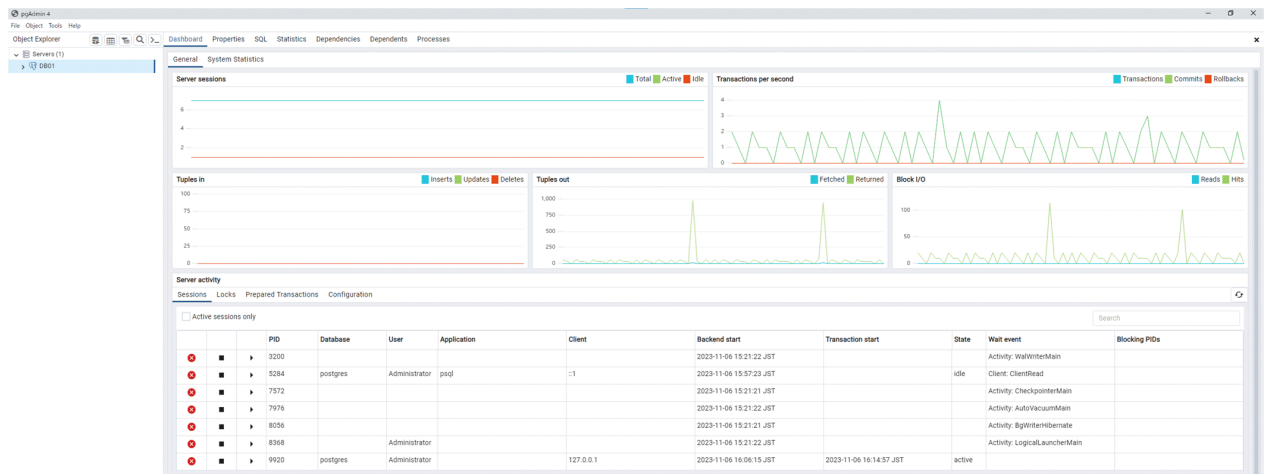
- Refer to "Notes on Application Compatibility" in the Application Development Guide for information on how to maintain application compatibility.

16.4.2 Using pgAdmin

If using pgAdmin, follow the procedure below to close connections.

1. In the [Browser] pane, click the relevant database server.
2. Close the client connections that have been in a wait state for an extended period.

Click the [Dashboard] tab. In the [Server activity] section, select the connections that have been in an "idle" or "idle in transaction" state for an extended period. For each of these connections, click  to close the session.



16.5 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.

16.6 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:

- [16.6.1 Using a Tablespace](#)
- [16.6.2 Replacing the Disk with a Larger Capacity Disk](#)

16.6.1 Using a Tablespace

Fujitsu Enterprise Postgres enables you to use a tablespace to change the storage destination of database objects, such as tables and indexes, to a different disk.

The procedure is as follows:

1. Create a tablespace

Use the CREATE TABLESPACE command to create a new tablespace in the prepared disk.

2. Modify the tablespace

Use the ALTER TABLE command to modify tables for the newly defined tablespace.



See

Refer to "SQL Commands" under "Reference" in the PostgreSQL Documentation for information on the CREATE TABLESPACE command and ALTER TABLE command.

16.6.2 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:

- [16.6.2.1 Using WebAdmin](#)
- [16.6.2.2 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.



Point

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.

16.6.2.1 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 "16.1.1 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.

16.6.2.2 Using Server Commands

Follow the procedure below to replace the disk and migrate resources at the data storage destination by using server commands.

1. Back up files

If the disk at the data storage destination contains any required files, back up the files. It is not necessary to back up the data storage destination.

2. Stop applications

Stop applications that are using the database.

3. Back up the database cluster

Back up the latest resources at the data storage destination. Refer to "3.2.2 Using Server Commands" for details.

4. Stop the instance

After backup is complete, stop the instance. Refer to "2.1.2 Using Commands" for information on how to stop an instance.

If the instance fails to stop, refer to "16.11 Actions in Response to Failure to Stop an Instance".

5. Replace with a larger capacity disk

Replace the disk. Then, recover the volume configuration information.

6. Create a data storage destination

Create a data storage destination. If a tablespace was defined, also create a directory for it.

In [Properties] in Windows(R) Explorer, set appropriate permissions so that only the instance administrator can access the data storage destination directory. (Refer to [Help and Support] in Windows(R) for information on [Properties].)

7. Recover the keystore, and enable automatic opening of the keystore

When the data in the database has been encrypted, restore the keystore to its state at the time of the database backup. Configure automatic opening of the keystore as necessary.

8. Recover the database cluster

Use the `pgx_rcvall` command to recover the database cluster.

- Specify the data storage destination in the `-D` option. If the `-D` option is omitted, the value of the `PGDATA` environment variable is used by default.
- Specify the backup storage directory in the `-B` option.

Example

```
> pgx_rcvall -D D:\database\inst1 -B E:\backup\inst1
```

If recovery fails

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.

Information

The following message displayed during recovery is output as part of normal operation of `pgx_rcvall` command (therefore the user does not need not be concerned).

```
FATAL: the database system is starting up
```

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 Commands](#)" for information on how to start an instance.

10. Resume applications

Resume applications that are using the database.

11. Restore files

Restore the files backed up in step 1.

16.7 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:

- [16.7.1 Temporarily Saving Backup Data](#)
- [16.7.2 Replacing the Disk with a Larger Capacity Disk](#)

16.7.1 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:

- [16.7.1.1 Using WebAdmin](#)

16.7.1.1 Using WebAdmin

Follow the procedure below to recover the backup data storage disk.

1. Temporarily save backup data

Move backup data to a different directory and temporarily save it, and secure space in the backup data storage destination directory.

The reason for saving the backup data is so that the data in the data storage destination can be recovered even if it is corrupted before you perform recovery. If there is no disk at the save destination and you consider that there is no risk of corruption at the data storage destination, delete the backup data.

The following example saves backup data from the backup data storage destination directory (E:\backup\inst1) under F:\mnt\usb\backup.

Example

```
> mkdir F:\mnt\usb\backup
> move E:\backup\inst1\* F:\mnt\usb\backup
```

Note: Place the temporary backup destination directory in a location where it will not impact on operating system resources or Fujitsu Enterprise Postgres resources.

2. Back up the database cluster

Back up the latest resources at the data storage destination. Refer to "3.2.1 Using WebAdmin" for details.

3. Delete temporarily saved backup data

If backup completes normally, the temporarily saved backup data becomes unnecessary and is deleted.

The following example deletes backup data that was temporarily saved in F:\mnt\usb.

Example

```
> rmdir /S /Q F:\mnt\usb\backup
```

16.7.1.2 Using Server Commands

The following describes the procedure for recovering the backup storage disk.

There are two methods of taking action:

- Performing recovery while the instance is active
- Stopping the instance before performing recovery

The following table shows the different steps to be performed depending on whether you stop the instance.

No	Step	Instance stopped	
		No	Yes
1	Stop transaction log mirroring	Y	N
2	Stop output of archive logs	Y	N
3	Stop applications	N	Y
4	Stop the instance	N	Y
5	Temporarily save backup data	Y	Y
6	Resume output of archive logs	Y	N
7	Resume transaction log mirroring	Y	N
8	Start an instance	N	Y

No	Step	Instance stopped	
		No	Yes
9	Run backup	Y	Y
10	Resume applications	N	Y
11	Delete temporarily saved backup data	Y	Y

Y: Required

N: Not required

The procedure is as follows:

Performing recovery while the instance is active

1. Stop transaction log mirroring

Stop transaction log mirroring.

```
postgres=# SELECT pgx_pause_wal_multiplexing();
LOG:  multiplexing of transaction log files has been stopped
pgx_pause_wal_multiplexing
-----
(1 row)
```

2. Stop output of archive logs

Transaction logs may accumulate during replacement of backup storage disk, and if the data storage disk or the transaction log storage disk becomes full, there is a risk that operations may not be able to continue.

To prevent this, use the following methods to stop output of archive logs.

- Changing the archive_command parameter

Specify a command that will surely complete normally, so that archive logs will be regarded as having been output.

If you specify echo, a message is output to the server log, so it may be used as a reference when you conduct investigations.

- Reloading the configuration file

Run the pg_ctl reload command or the pg_reload_conf SQL function.

If you simply want to stop output of errors without the risk that operations will not be able to continue, specify an empty string (") in archive_command and reload the configuration file.

3. Temporarily save backup data

Move backup data to a different directory and temporarily save it, and secure space in the backup data storage destination directory.

The reason for saving the backup data is so that the data in the data storage destination can be recovered even if it is corrupted before you perform the next step. If there is no disk at the save destination and you consider that there is no risk of corruption at the data storage destination, delete the backup data.

The following example saves backup data from the backup data storage destination directory (E:\backup\inst1) under F:\mnt\usb\backup.

Example

```
> mkdir F:\mnt\usb\backup
> move E:\backup\inst1\* F:\mnt\usb\backup
```

Note: Place the temporary backup destination directory in a location where it will not impact on operating system resources or Fujitsu Enterprise Postgres resources.

4. Resume output of archive logs

Return the archive_command setting to its original value, and reload the configuration file.

5. Resume transaction log mirroring

Execute the pgx_resume_wal_multiplexing SQL function.

Example

```
SELECT pgx_resume_wal_multiplexing()
```

6. Run backup

Use the pgx_dmpall command to back up the database cluster.

Specify the following option in the pgx_dmpall command:

- Specify the directory of the data storage destination in the -D option. If the -D option is omitted, the value of the PGDATA environment variable is used by default.

Example

```
> pgx_dmpall -D D:\database\inst1
```

7. Delete temporarily saved backup data

If backup completes normally, the temporarily saved backup data becomes unnecessary and is deleted.

The following example deletes backup data that was temporarily saved in F:\mnt\usb.

Example

```
> rmdir /S /Q F:\mnt\usb\backup
```

If an instance has been stopped

1. Stop applications

Stop applications that are using the database.

2. Stop the instance

Stop the instance. Refer to ["2.1.2 Using Commands"](#) for details.

If the instance fails to stop, refer to ["16.11 Actions in Response to Failure to Stop an Instance"](#).

3. Temporarily save backup data

Move backup data to a different directory and temporarily save it, and secure space in the backup data storage destination directory.

The reason for saving the backup data is so that the data in the data storage destination can be recovered even if it is corrupted before you perform recovery. If there is no disk at the save destination and you consider that there is no risk of corruption at the data storage destination, delete the backup data.

The following example saves backup data from the backup data storage destination directory (E:\backup\inst1) under F:\mnt\usb\backup.

Example

```
> mkdir F:\mnt\usb\backup  
> move E:\backup\inst1\* F:\mnt\usb\backup
```

Note: Place the temporary backup destination directory in a location where it will not impact on operating system resources or Fujitsu Enterprise Postgres resources.

4. Start the instance

Start the instance. Refer to ["2.1.2 Using Commands"](#) for information on how to start an instance.

5. Run backup

Use the `pgx_dmpall` command to back up the database cluster.

Specify the following value in the `pgx_dmpall` command:

- Specify the data storage destination in the `-D` option. If the `-D` option is omitted, the value of the `PGDATA` environment variable is used by default.

Example

```
> pgx_dmpall -D D:\database\inst1
```

6. Resume applications

Resume applications that are using the database.

7. Delete temporarily saved backup data

If backup completes normally, the temporarily saved backup data becomes unnecessary and is deleted.

The following example deletes backup data that was temporarily saved in `F:\mnt\usb`.

Example

```
> rmdir /S /Q F:\mnt\usb\backup
```



See

- Refer to "pgx_rcvall" and "pgx_dmpall" in the Reference for information on the `pgx_rcvall` command and `pgx_dmpall` command.
- Refer to "Write Ahead Log" under "Server Administration" in the PostgreSQL Documentation for information on `archive_command`.
- Refer to "[B.1 WAL Mirroring Control Functions](#)" for information on the `pgx_is_wal_multiplexing_paused` and `pgx_resume_wal_multiplexing`.

16.7.2 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:

- [16.7.2.1 Using WebAdmin](#)
- [16.7.2.2 Using Server Commands](#)

16.7.2.1 Using WebAdmin

Follow the procedure below to recover the backup storage disk.

1. Back up files

If the disk at the backup data storage destination contains any required files, back up the files. It is not necessary to back up the backup data storage destination.

2. Temporarily save backup data

Save the backup data to a different directory.

The reason for saving the backup data is so that the data in the data storage destination can be recovered even if it is corrupted before you perform the next step. If there is no disk at the save destination and you consider that there is no risk of corruption at the data storage destination, delete the backup data.

The following example saves backup data from the backup data storage destination directory (`E:\backup\inst1`) under `F:\mnt\usb\backup`.

Example

```
> mkdir F:\mnt\usb\backup
> move E:\backup\inst1\* F:\mnt\usb\backup
```

Note: Place the temporary backup destination directory in a location where it will not impact on operating system resources or Fujitsu Enterprise Postgres resources.

3. Stop applications

Stop applications that are using the database.

4. Replace with a larger capacity disk

Replace the disk. Then, recover the volume configuration information.

5. 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 "16.1.1 Using WebAdmin".

6. Restore files

Restore the files backed up in step 1.

7. Delete temporarily saved backup data

If backup completes normally, the temporarily saved backup data becomes unnecessary and is deleted.

The following example deletes backup data that was temporarily saved in F:\mnt\usb.

Example

```
> rmdir /S /Q F:\mnt\usb\backup
```

16.7.2.2 Using Server Commands

The procedure for recovering the backup data storage disk is described below.

There are two methods of taking action:

- Performing recovery while the instance is active
- Stopping the instance before performing recovery

The following table shows the different steps to be performed depending on whether you stop the instance.

No	Step	Instance stopped	
		No	Yes
1	Back up files	Y	Y
2	Temporarily save backup data	Y	Y
3	Confirm that transaction log mirroring has stopped	Y	N
4	Stop output of archive logs	Y	N
5	Stop applications	N	Y
6	Stop the instance	N	Y
7	Replace with a larger capacity disk	Y	Y
8	Create a backup storage directory	Y	Y
9	Resume output of archive logs	Y	N
10	Resume transaction log mirroring	Y	N
11	Start the instance	N	Y
12	Run backup	Y	Y

No	Step	Instance stopped	
		No	Yes
13	Resume applications	N	Y
14	Restore files	Y	Y
15	Delete temporarily saved backup data	Y	Y

Y: Required

N: Not required

The procedure is as follows:

If an instance has not been stopped

1. Back up files

If the disk at the backup data storage destination contains any required files, back up the files. It is not necessary to back up the backup data storage destination.

2. Temporarily save backup data

Save the backup data to a different directory.

The reason for saving the backup data is so that the data in the data storage destination can be recovered even if it is corrupted before you perform the next step. If there is no disk at the save destination and you consider that there is no risk of corruption at the data storage destination, delete the backup data.

The following example saves backup data from the backup data storage destination directory (E:\backup\inst1) under F:\mnt\usb\backup.

Example

```
> mkdir F:\mnt\usb\backup
> move E:\backup\inst1\* F:\mnt\usb\backup
```

3. Confirm that transaction log mirroring has stopped

Use the following SQL function to confirm that transaction log mirroring has stopped.

```
postgres=# SELECT pgx_is_wal_multiplexing_paused();
pgx_is_wal_multiplexing_paused
-----
t
(1 row)
```

If transaction log mirroring has not stopped, then stop it using the following SQL function.

```
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, so that archive logs will be regarded as having been output.

If you specify echo, a message is output to the server log, so it may be used as a reference when you conduct investigations.

- Reloading the configuration file

Run the `pg_ctl reload` command or the `pg_reload_conf` SQL function.

If you simply want to stop output of errors without the risk that operations will not be able to continue, specify an empty string ("") in `archive_command` and reload the configuration file.

5. Replace with a larger capacity disk

Replace the disk. Then, recover the volume configuration information.

6. Create a backup data storage destination

Create a backup data storage destination.

In [Properties] in Windows(R) Explorer, set appropriate permissions so that only the instance administrator can access the backup data storage destination directory. (Refer to [Help and Support] in Windows(R) for information on [Properties].)

Refer to "[3.2.2 Using Server Commands](#)" for details.

7. Resume output of archive logs

Return the `archive_command` setting to its original value, and reload the configuration file.

8. Resume transaction log mirroring

Execute the `pgx_resume_wal_multiplexing` SQL function.

Example

```
SELECT pgx_resume_wal_multiplexing()
```

9. Run backup

Use the `pgx_dmpall` command to back up the database cluster.

Specify the following value in the `pgx_dmpall` command:

- Specify the data storage destination in the `-D` option. If the `-D` option is omitted, the value of the `PGDATA` environment variable is used by default.

Example

```
> pgx_dmpall -D D:\database\inst1
```

10. Restore files

Restore the files backed up in step 1.

11. Delete temporarily saved backup data

If backup completes normally, the temporarily saved backup data becomes unnecessary and is deleted.

The following example deletes backup data that was temporarily saved in `F:\mnt\usb`.

Example

```
> rmdir /S /Q F:\mnt\usb\backup
```

If an instance has been stopped

1. Back up files

If the disk at the backup data storage destination contains any required files, back up the files. It is not necessary to back up the backup data storage destination.

2. Temporarily save backup data

Save the backup data to a different directory.

The reason for saving the backup data is so that the data in the data storage destination can be recovered even if it is corrupted before you perform the next step. If there is no disk at the save destination and you consider that there is no risk of corruption at the data storage destination, delete the backup data.

The following example saves backup data from the backup data storage destination directory (E:\backup\inst1) under F:\mnt\usb\backup.

Example

```
> mkdir F:\mnt\usb\backup
> move E:\backup\inst1\* F:\mnt\usb\backup
```

Note: Place the temporary backup destination directory in a location where it will not impact on operating system resources or Fujitsu Enterprise Postgres resources.

3. Stop applications

Stop applications that are using the database.

4. Stop the instance

Stop the instance. Refer to "[2.1.2 Using Commands](#)" for information on how to stop an instance.

If the instance fails to stop, refer to "[16.11 Actions in Response to Failure to Stop an Instance](#)".

5. Replace with a larger capacity disk

Replace the disk. Then, recover the volume configuration information.

6. Create a backup data storage destination

Create a backup data storage destination.

In [Properties] in Windows(R) Explorer, set appropriate permissions so that only the instance administrator can access the backup data storage destination directory. (Refer to [Help and Support] in Windows(R) for information on [Properties].)

Refer to "[3.2.2 Using Server Commands](#)" for details.

7. Start the instance

Start the instance. Refer to "[2.1.2 Using Commands](#)" for information on how to start an instance.

8. Run backup

Use the `pgx_dmpall` command to back up the database cluster.

Specify the following value in the `pgx_dmpall` command:

- Specify the data storage destination in the `-D` option. If the `-D` option is omitted, the value of the `PGDATA` environment variable is used by default.

Example

```
> pgx_dmpall -D D:\database\inst1
```

9. Resume applications

Resume applications that are using the database.

10. Restore files

Restore the files backed up in step 1.

11. Delete temporarily saved backup data

If backup completes normally, the temporarily saved backup data becomes unnecessary and is deleted.

The following example deletes backup data that was temporarily saved in F:\mnt\usb.

Example

```
> rmdir /S /Q F:\mnt\usb\backup
```



See

- Refer to "pgx_rcvall" and "pgx_dmpall" in the Reference for information on the pgx_rcvall command and pgx_dmpall command.
- Refer to "Write Ahead Log" under "Server Administration" in the PostgreSQL Documentation for information on archive_command.
- Refer to "[B.1 WAL Mirroring Control Functions](#)" for information on the pgx_is_wal_multiplexing_paused and pgx_resume_wal_multiplexing.

16.8 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.

16.8.1 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:

- [16.8.1.1 Using WebAdmin](#)
- [16.8.1.2 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.



Point

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.

16.8.1.1 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 "[16.1.1 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.

16.8.1.2 Using Server Commands

Follow the procedure below to replace the disk and migrate resources at the transaction log storage destination by using server commands.

1. Back up files

If the disk at the transaction log storage destination contains any required files, back up the files. It is not necessary to back up the transaction log storage destination.

2. Back up the database cluster

Use server commands to back up the latest data storage destination resources and transaction log storage destination resources. Refer to "[3.2.2 Using Server Commands](#)" for information on how to perform backup.

3. Stop applications

Stop applications that are using the database.

4. Stop the instance

After backup is complete, stop the instance. Refer to "[2.1.2 Using Commands](#)" for information on how to stop an instance.

If the instance fails to stop, refer to "[16.11 Actions in Response to Failure to Stop an Instance](#)".

5. Replace with a larger capacity disk

Replace the disk. Then, recover the volume configuration information.

6. Create a transaction log storage destination

Create a transaction log storage destination. If a tablespace was defined, also create a directory for it.

In [Properties] in Windows(R) Explorer, set appropriate permissions so that only the instance administrator can access the transaction log destination directory. (Refer to [Help and Support] in Windows(R) for information on [Properties].)

7. Recover the keystore, and enable automatic opening of the keystore

When the data in the database has been encrypted, restore the keystore to its state at the time of the database backup. Configure automatic opening of the keystore as necessary.

8. Recover the database cluster

Use the `pgx_rcvall` command to recover the database cluster.

- Specify the data storage destination in the -D option. If the -D option is omitted, the value of the PGDATA environment variable is used by default.
- Specify the backup storage directory in the -B option.

Example

```
> pgx_rcvall -D D:\database\inst1 -B E:\backup\inst1
```

If recovery fails

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.

 **Information**

The following message displayed during recovery is output as part of normal operation of pgx_rcvall command (therefore the user does not need not be concerned).

```
FATAL: the database system is starting up
```

 **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 Commands" for information on how to start an instance.

10. Resume applications

Resume applications that are using the database.

11. Restore files

Restore the files backed up in step 1.

16.9 Errors in More Than One Storage Disk

If an error occurs in the storage destination disks or resources are corrupted, determine the cause of the error from event logs and server logs and remove the cause.

If errors occur in either of the following combinations, you cannot recover the database.

Recreate the instance, and rebuild the runtime environment.

Data storage disk	Transaction log storage disk	Backup data storage disk
Error	-	Error
-	Error	Error



See

Refer to "Setup" in the Installation and Setup Guide for Server for information on how to create an instance and build the runtime environment.

16.10 Actions in Response to Instance Startup Failure

If an instance fails to start, refer to the event log and the server log, and determine the cause of the failure.

If using WebAdmin, remove the cause of the error. Then, click [Solution] and [Recheck the status] and confirm that the instance is in the normal state.

The following sections describe common causes of errors and the actions to take.

16.10.1 Errors in the Configuration File

If you have directly edited the configuration file using a text editor or changed the settings using WebAdmin, refer to the event log and the server log, confirm that no messages relating to the files below have been output.

- postgresql.conf
- pg_hba.conf



See

Refer to the following for information on the parameters in the configuration file:

- "Configuring Parameters" in the Installation and Setup Guide for Server
- "Appendix A Parameters"
- "Server Configuration" and "Client Authentication" under "Server Administration" in the PostgreSQL Documentation

16.10.2 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 "[16.14.2 Errors Caused by Power Failure or Mounting Issues](#)", and take actions accordingly.

16.10.3 Other Errors

This section describes the recovery procedure to be used if you cannot take any action or the instance cannot start even after you have referred to the event log and the server log.

There are two methods of recovery:

- [16.10.3.1 Using WebAdmin](#)
- [16.10.3.2 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.

16.10.3.1 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 "[16.2.1 Using WebAdmin](#)" for details.

16.10.3.2 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 "[16.2.2 Using the pgx_rcvall Command](#)" for details.

16.11 Actions in Response to Failure to Stop an Instance


If an instance fails to stop, refer to the event log and the server log, and determine the cause of the failure.

If the instance cannot stop despite taking action, perform the following operation to stop the instance.

There are two methods of recovery:

- [16.11.1 Using WebAdmin](#)
- [16.11.2 Using Server Commands](#)

16.11.1 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.

16.11.2 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.

16.11.2.1 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 "[16.11.2.2 Stopping the Instance Using the Immediate Mode](#)" or "[16.11.2.3 Forcibly Stopping the Server Process](#)".



Example

```
> pg_ctl stop -D D:\database\inst1 -m fast
```

16.11.2.2 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 "16.11.2.3 Forcibly Stopping the Server Process".



Example

```
> pg_ctl stop -D D:\database\inst1 -m immediate
```

16.11.2.3 Forcibly Stopping the Server Process

If both the Fast mode and the Immediate mode fail to stop the instance, use the kill parameter of the pg_ctl command to forcibly stop the server process.

The procedure is as follows:

1. Execute the wmic command to identify the process ID of the server process.

```
c:\>wmic
wmic:root\cli>process where "name = \"postgres.exe\"" get CommandLine,Name,ProcessId
CommandLine                                     Name
ProcessId
:
"C:\Program Files\Fujitsu\fserv<x>server64\bin\postgres.exe" -D "D:\database\inst1"
postgres.exe 896
:
```

The postgres.exe process ID(896) that indicates the data storage destination directory of the applicable instance in the -D option becomes the server process.

2. Forcibly stop the server process

As instance manager, forcibly stop the server process using the pg_ctl command.

```
c:\>pg_ctl kill QUIT 896
```

16.12 Actions in Response to Failure to Create a Streaming Replication Standby Instance

When creating a streaming replication standby instance using WebAdmin, if the instance creation fails, refer to the event log and the server log, and determine the cause of the failure.

When an error occurs in the creation of the standby instance using WebAdmin, it is unlikely that the partially created standby instance can be resumed to complete the operation.

In such a scenario, fix the cause of the error, delete the partially created standby instance, and then create a new standby instance. This recommendation is based on the following assumptions:

- As the instance is yet to be created completely, there are no applications connecting to the database.
- The standby instance is in error state and is not running.
- There are no backups for the standby instance and as a result, it cannot be recovered.



See

Refer to "Deleting Instances" in the Installation and Setup Guide for details on how to delete an instance.

16.13 Actions in Response to Error in a Distributed Transaction

If a system failure (such as server failure) occurs in an application that uses distributed transactions (such as .NET TransactionScope), then transactions may be changed to the in-doubt state.

At that point, resources accessed by the transaction will be locked, and rendered unusable by other transactions.

The following describes how to check for in-doubt transactions, and how to resolve them.

How to check for in-doubt transactions

The following shows how to check for them:

If the server fails

1. An in-doubt transaction will have occurred if a message similar to the one below is output to the log when the server is restarted.

Example

```
LOG: Restoring prepared transaction 2103.
```

2. Refer to system view `pg_prepared_xacts` to obtain information about the prepared transaction.

If the transaction identifier of the prepared transaction in the list (in the `transaction` column of `pg_prepared_xacts`) is the same as the identifier of the in-doubt transaction obtained from the log output when the server was restarted, then that row is the information about the in-doubt transaction.

Example

```
postgres=# select * from pg_prepared_xacts;
 transaction | gid          | prepared | owner      | database
-----+-----+-----+-----+-----
 2103 | 374cc221-f6dc-4b73-9d62-d4fec9b430cd | 2022-03-06 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.



Point

As you can see from the explanations in this section, there is no one way to definitively determine in-doubt transactions.

Consider collecting other supplementary information (for example, logging on the client) or performing other operations (for example, allocating database users per job).

How to resolve in-doubt transactions

From the system view `pg_prepared_xacts` mentioned above, obtain the global transaction identifier (in the `gid` column of `pg_prepared_xacts`) for the in-doubt transaction, and issue either a `ROLLBACK PREPARED` statement or `COMMIT PREPARED` statement to resolve the in-doubt transaction.

Example

- Rolling back in-doubt transactions

```
postgres=# rollback prepared '374cc221-f6dc-4b73-9d62-d4fec9b430cd';  
ROLLBACK PREPARED
```

- Committing in-doubt transactions

```
postgres=# commit prepared '374cc221-f6dc-4b73-9d62-d4fec9b430cd';  
COMMIT PREPARED
```

16.14 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.

- [16.14.1 Network Error with an External Disk](#)
- [16.14.2 Errors Caused by Power Failure or Mounting Issues](#)

16.14.1 Network Error with an External Disk

This is an error that occurs in the network path to/from an external disk.

Determine the cause of the error by checking the information in the event log and the server log, the disk access LED, network wiring, and network card status. Take appropriate action to remove the cause of the error, for example, replace problematic devices.

16.14.2 Errors Caused by Power Failure or Mounting Issues

These are errors that occur when the disk device is not turned on, automatic mounting of the disk was not set, or mounting was accidentally cancelled.

In this case, check the information in the event log and the server log, the disk access LED, and whether the disk is mounted correctly. If problems are detected, take appropriate action.

If mounting has been cancelled, it is possible that mounting was accidentally cancelled, or the existing setting (automatic mounting at the time of starting the operating system) has been changed so that mounting is not performed automatically. In this case, set the mounting to be performed automatically.

16.15 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.

16.15.1 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

Critical errors encountered during anomaly resolution will be displayed, however, rollback of the instance to its previous state is not supported.

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

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.

16.15.2 Mirroring Controller Anomalies

The following conditions will cause a Mirroring Controller anomaly:

- The Mirroring Controller management folder or configuration files have been deleted
- The permissions to the Mirroring Controller management folder or configuration files have been changed such that:
 - The instance administrator's access to Mirroring Controller configuration is denied

- Users other than an instance administrator have access privileges to Mirroring Controller configuration files

WebAdmin checks for anomalies when Mirroring Controller status check is performed.

The following occurs when a Mirroring Controller anomaly is detected:

- All Mirroring Controller functionality is disabled for the replication cluster, except for "Delete Mirroring Controller"
- [Mirroring Controller status] is set to "Error"
- Either of the following messages is displayed in the [Message] section

"Failed to access the Mirroring Controller management folder or configuration files '*path*'. Mirroring Controller functionality has been disabled. Consider deleting Mirroring Controller and adding it again."

"Failed to find the Mirroring Controller management folder or configuration files '*path*'. Mirroring Controller functionality has been disabled. Consider deleting Mirroring Controller and adding it again."

Appendix A Parameters

This appendix describes the parameters to be set in the postgresql.conf file of Fujitsu Enterprise Postgres.

The postgresql.conf file is located in the data storage destination.



Information

The maximum value that can be expressed as a 4-byte signed integer changes according to the operating system. Follow the definition of the operating system in use.

Specify Path

Note the following when specifying the path:

- Specify `\\` as the path delimiter.
- Enclose the path in double quotes (") if it contains spaces.
- `core_directory` (string)

This parameter specifies the directory where the corefile is to be output. If this parameter is omitted, the data storage destination is used by default. This parameter can only be set when specified on starting an instance. It cannot be changed dynamically, while an instance is active.

- `core_contents` (string)

This parameter specifies the contents to be included in the corefile.

- `full`: Outputs all contents of the server process memory to the corefile.
- `none`: Does not output a corefile.
- `minimum`: Outputs only non-shared memory server processes to the corefile. This reduces the size of the corefile. However, in some cases, this file may not contain sufficient information for examining the factor that caused the corefile to be output.

If this parameter is omitted, "minimum" is used by default. This parameter can only be set when specified on starting an instance. It cannot be changed dynamically, while an instance is active.

- `keystore_location` (string)

This parameter specifies the directory that stores the keystore file. Specify a different location from other database clusters. This parameter can only be set when specified on starting an instance. It cannot be changed dynamically, while an instance is active.

- `tablespace_encryption_algorithm` (string)

This parameter specifies the encryption algorithm for tablespaces that will be created. Valid values are "AES128", "AES256", and "none". If you specify "none", encryption is not performed. The default value is "none". To perform encryption, it is recommended that you specify "AES256". Only superusers can change this setting.

- `backup_destination` (string)

This parameter specifies the absolute path of the directory where `pgx_dmpall` will store the backup data. Specify a different location from other database clusters. This parameter can only be set when specified on starting an instance. It cannot be changed dynamically, while an instance is active.

Place this directory on a different disk from the data directory to be backed up and the tablespace directory. Ensure that users do not store arbitrary files in this directory, because the contents of this directory are managed by the database system.

- `search_path` (string)

When using the `SUBSTR` function compatible with Oracle databases, set "oracle" and "pg_catalog" in the `search_path` parameter. You must specify "oracle" before "pg_catalog".

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 / 1024

If this parameter is omitted or a value outside this range is specified, 256 MB will be used.

- vci.max_local_ros (numerical value)

This parameter specifies the maximum memory size used for VCI scan.

- Minimum value: 64 MB
- Maximum value: Maximum value that can be expressed as a 4-byte signed integer

If this parameter is omitted or a value outside this range is specified, 64 MB will be used.

- vci.max_parallel_degree (numerical value)

This parameter specifies the maximum number of background workers used for parallel scan. The number of workers for the entire instance is limited by max_worker_processes, so add the value specified here to max_worker_processes.

A value from -8388607 to 8388607 can be specified.

- Integer (1 or greater): Parallel scan is performed using the specified degree of parallelism.
- 0: Stops the parallel scan process.
- Negative number: The specified value minus the maximum number of CPUs obtained from the environment is used as the degree of parallelism and parallel scan is performed.

If this parameter is omitted or a value outside this range is specified, 0 will be used.

- vci.shared_work_mem (numerical value)

This parameter specifies the maximum memory size used for VCI parallel scan.

- Minimum value: 32 MB
- Maximum value: Maximum value that can be expressed as a 4-byte signed integer

If this parameter is omitted or a value outside this range is specified, 1 GB will be used.

- vci.smc_directory (string)

This parameter specifies a directory name in which a temporary file is created as the dynamic shared memory during a scan using a VCI.

If this parameter is omitted, a directory (*dataStorageDir*\base\pgsql_tmp) under the data storage directory will be used.

Parameters for the Global Meta Cache feature

- `pgx_global_metacache` (numerical value)

Specifies the memory size of the GMC area.

Specify a value calculated by the formula below.

A value lower than the calculated value will still work, but the meta cache may not be able to fit into the GMC area.

In this case, the system will discard the meta cache it thinks it is no longer needed, but if it is needed again, the meta cache will need to be expanded and will not perform well.

If the value is less than 10 MB and is set to a nonzero value that disables the feature, the database startup fails because the Global Meta Cache feature cannot operate.

A setting of 0 disables the Global Meta Cache feature. The default is 0.

Changing this setting requires restarting the database.

```
Size of GMC area
= Max(10MB,
      (All user table x 0.4 KB
       + All user Indexes x 0.3 KB
       + All user columns x 0.8 KB) x 1.5 (*1) )
```

*1) Safety Factor (1.5)

This value takes into account the case where both GMC before and after the change temporarily exist at the same time in shared memory when the table definition is changed or the row of the system catalog is changed.

- `track_gmc` (string)

This parameter enables collection of statistics for `pgx_stat_gmc`.

- `on`: Enables collection of statistics.
- `off`: Disables collection of statistics.

If this parameter is omitted, "on" is used.

Only superusers can change this setting.

Parameters for the Local Meta Cache Limit feature

- `pgx_catalog_cache_max_size`(numerical value)

Specifies the maximum amount of memory that the backend process should use as the catalog cache.

You can enable catalog cache deletion by setting it to 8 KB or more.

A setting of 0 disables the catalog cache removal. The default is 0.

If no units are specified, they are treated as KB.

- Minimum value: 8KB
- Maximum value: Maximum value that can be expressed as a 4-byte signed integer

When calculating the parameter settings, the factors that determine the cache size are calculated as the number of tables, the number of indexes, and the number of columns. What is kept as a catalog cache or relation cache also includes objects such as databases, roles, or procedures, but these are small compared to the above factors and do not need to be factored into them. It also includes a calculation method for `pgx_relation_cache_max_size` because the given memory is distributed between the catalog cache and the relation cache.

The calculation method here assumes that all backends have similar access and that the transaction also has access to a similar number of resources. If you have a small number of singular backends or transactions, consider excluding them as errors.

1. Determine how much memory a backend process can use. Decide by subtracting the memory size required by the entire system such as the database cache from the installed memory and dividing the rest by the number of connections.
2. For best performance, use the following formula to calculate the total memory size of the catalog cache when the backend holds the catalog cache for all resources accessed during its lifetime.

The amount of memory varies depending on whether Global Meta Cache is enabled or disabled. Enabling Global Meta Cache reduces the amount of memory required because most of the cache is located on shared memory.

When Global Meta Cache is enabled:

$$(\text{Number of tables to access} + \text{Number of indexes to access} + \text{Number of columns to access}) \times 0.1\text{KB} \times 1.5 \text{ (*1)}$$

When Global Meta Cache is disabled:

$$\{ \text{Number of tables to access} \times 0.5\text{KB}(\text{pg_class tuple size}) + \text{Number of indexes to access} \times 0.5\text{KB}(\text{pg_index tuple size}) + \text{Number of columns to access} \times 1.0\text{KB}(\text{pg_statistic tuple size}) \} \times 1.5 \text{ (*1)}$$

*1) Safety Factor (1.5)

The system catalog contains columns with variable-length types. For example, the tuple size in pg_class is a constant value multiplied by the number of tables, while relname in pg_class is variable length data.

It is not practical to calculate every definition in detail, so we added 50% to the above formula.

3. In the same way as in 2., calculate the relation cache using the following formula.

$$(1.4\text{KB} \times \text{Number of tables to access} + 2.4\text{KB} \times \text{Number of indexes to access}) \times 1.5 \text{ (*1)}$$

*1) Safety Factor (1.5)

The relation cache is structured to facilitate the use of table and index definitions, holds pointers to various objects, and is sized to include them. It is variable length because the type of object allocated by the table definition and its size change. Since it is not realistic to calculate for all definitions, 50% is added.

4. If the value of 1. the value of 2. + the value of 3., the backend process can keep all caches to the extent allowed, so there is no need to limit the caches. If you want to cap for safety, set the value of 2. to pgx_catalog_cache_max_size and the value of 3. to pgx_relation_cache_max_size.
5. If the value of 1. < the value of 2. + the value of 3. then you need to limit the cache. However, this parameter does not limit the size of the cache used by a transaction. Therefore, take the following steps.
6. Calculate the catalog cache used by a transaction using the formula in 2.
7. Calculate the relation cache used by a transaction using the formula in 3.
8. If the value of 1. < the value of 6. + the value of 7., then the value of 1. needs to be increased. In other words, in some cases, it may be necessary to increase the installed memory or reduce the number of connections.
9. If the value of 1. the value of 6. + the value of 7., the condition of 1. can be satisfied by limiting the cache with this parameter. Divide the value of 1. by the ratio of 2. and 3. and set it as a parameter. Set the value distributed to 2. to pgx_catalog_cache_max_size and the value distributed to 3. to pgx_relation_cache_max_size.
10. The value calculated in 9. is a provisional value. If you cannot meet your target performance, first try to shift the focus of allocation to the relation cache. This is because when executing SQL, the relation cache generated based on the catalog cache is mainly referenced, so it is advantageous to leave a large amount of relation cache. If the performance is still not satisfied, adjust the parameters by referring to "[14.1.4 Performance Impact and Parameter Tuning of the Local Meta Cache Limit Feature](#)".

For Partitioning

Be careful when partitioning the table.

The cached definition changes depending on whether the parent table is specified in the SQL statement or the child table is specified. In particular, note that if you specify a parent table, the definitions of all child tables are cached. This is because when you specify a parent table in an SQL statement, you need to know the definitions of all the child tables in order to determine which child table will contain the desired data. Note that the column information of the parent table is not cached.

When specifying the parent table:

$$\begin{aligned} \text{Number of tables to access} &= \text{Number of parent tables to access} + \text{Number of defined child tables} \\ \text{Number of columns} &= \text{Number of defined columns} \times \text{number of defined child tables} \end{aligned}$$

When specifying the child table directly:

Number of tables to access = Number of child tables actually accessed
Number of columns = Number of defined columns x number of child tables actually accessed

Example)

Suppose the parent table T (1 index, 3 columns) is split from child tables T1 to T5 (1 index, 3 columns, respectively). If the parent table T is specified in SQL, when the child tables that contain the data to be queried are limited to T1 and T2, and when accessing the data using the indexes defined by T1 and T2, calculate as follows.

Number of tables = 1(parent table) + 5(child table) = 6
Number of indexes = 2 (index to access)
Number of columns = 3 (number of columns) x 5 (child table) = 15

If you specify child tables T1 and T2 in SQL and use the indexes defined on T1 and T2 when accessing data, the calculation is as follows.

Number of tables = 2(child table)
Number of indexes = 2 (index to access)
Number of columns = 3 (number of columns) x 2 (child table) = 6

- `pgx_relation_cache_max_size`(numerical value)

Specifies the maximum amount of memory that the backend process should use as the relation cache.

You can enable catalog cache deletion by setting it to 8 KB or more.

A setting of 0 disables the relation cache removal. The default is 0.

If no units are specified, they are treated as KB.

- Minimum value: 8KB
- Maximum value: Maximum value that can be expressed as a 4-byte signed integer

For the calculation method for parameter setting, refer to the calculation method of `pgx_catalog_cache_max_size`.

- `pgx_cache_hit_log_interval`(numerical value)

Specifies the time interval to output a message indicating the cache reference status for each backend process.

When the transaction ends, if the time set in this parameter has elapsed since the previous message was output, the message is output.

If set to 0, a message will be output each time the transaction ends.

Setting -1 disables the output. The default value is 10min.

If no units are specified, they are treated as ms.

Even if `pgx_catalog_cache_max_size` and `pgx_relation_cache_max_size` are disabled, the message output of the corresponding cache will be invalid.

Immediately after connecting to the server, a small transaction occurs before the request from the user application, such as for user authentication. Since it is meaningless to know the hit rate for these, a message will be output at the end of the transaction that started after the time set in this parameter has elapsed after connecting to the server.

For the same reason, setting a small value such as 0 may result in a message being printed at the end of such a small transaction.

You can check which transaction the message corresponds to from the information output at the beginning.

This information depends on the setting of the parameter `log_line_prefix`.

- Minimum value: 0
- Maximum value: 2147483647ms



See

Refer to "Server Configuration" under "Server Administration" in the PostgreSQL Documentation for information on other postgresql.conf parameters.

Parameters for the Policy-based Login Security

- userprofile_database(string)

Specifies the name of the database to which the background worker process connects. This parameter must specify a connectable database name. If you omit this parameter, then it is assumed to be "postgres".

The pseudo database "replication" cannot be specified for streaming replication.

If you change this parameter, reload the configuration file.

Parameters for using OpenSSL legacy algorithms

- openssl_conf(string)

Specifies the OpenSSL configuration file. Specifying a valid configuration file makes legacy algorithms available. Use the example below to prepare the configuration file in any directory.

If this parameter is not specified, the empty string is assumed.

This parameter can only be set by specifying the parameter at instance startup.

You cannot make dynamic changes during instance startup.



Example

```
openssl_conf = 'openssl.conf'
```

```
[OpenSSL Configuration File (openssl.conf) Example]
```

```
=====  
openssl_conf = openssl_init  
  
[openssl_init]  
providers = provider_sect  
  
[provider_sect]  
default = default_sect  
legacy = legacy_sect  
  
[default_sect]  
activate = 1  
  
[legacy_sect]  
activate = 1  
=====
```

- openssl_modules(string)

Specifies the directory that contains additional OpenSSL modules.

Legacy algorithms are available by specifying 'server installation directory\lib\ossl-modules'.

If this parameter is not specified, the empty string is assumed.

This parameter can only be set by specifying the parameter at instance startup.

You cannot make dynamic changes during instance startup.

This parameter sets the OPENSSL_MODULES environment variable to be applied to the server process. Do not set the OPENSSL_MODULES environment variable in any way other than setting this parameter, as this may cause abnormal behavior.



Example

openssl_modules = 'c:\Program Files\Fujitsu\fsepv<x>server64\lib\ossl-modules'

Note that "<x>" indicates the product version.



Information

Legacy algorithms include the following encryption algorithms:

- BF
- CAST5
- DES-ECB
- DES-CBC
- MD4
- Whirlpool

Parameters for scheduling aggressive freeze for tuples

- pgx_stat_vacuum_freeze.track_freeze_info (boolean)

Collects and maintains information about aggressive freeze for tuples. The default is on.

- on: Enables information collection.
- off: Disables information collection.

- pgx_stat_vacuum_freeze.monitoring_interval (numerical value)

Aggregate statistics for aggressive freeze for tuples in this interval, so that statistics for each interval can be obtained in the pgx_stat_freeze_results table. Note that an aggressive freeze for tuples across this interval will still account for the overall tuple freeze at completion time. The default is 1hour.

Set the aggressive freeze for tuples script to run at weekly intervals of days, daily intervals of hours, and hourly intervals of minutes.

- Minimum value: 1min
- Maximum value: 1d

If a value out of range is specified, 1h is set.

If this parameter is omitted, 1h is assumed.

- pgx_stat_vacuum_freeze.log_retention_period (numerical value)

Specify how long statistics are kept. The default is 90 days.

Set the period according to the work cycle. For example, if a monthly batch is to be processed every 1 month, set it to 90 days, which is approximately 3 times that period. to capture changes in the aggressive freeze for tuples over the duration of the cycle.

- Minimum value: 1d
- Maximum value: 180d

If an out-of-range value is specified, 90d is set.

If this parameter is omitted, 90d is assumed.

Appendix B System Administration Functions

This appendix describes the system administration functions of Fujitsu Enterprise Postgres.



Refer to "System Administration Functions" under "The SQL Language" in the PostgreSQL Documentation for information on other system administration functions.

B.1 WAL Mirroring Control Functions

The following table lists the functions that can be used for backup and recovery based on WAL mirroring.

Table B.1 WAL mirroring control functions

Name	Return type	Description
<code>pgx_pause_wal_multiplexing()</code>	void	Stops WAL multiplexing
<code>pgx_resume_wal_multiplexing()</code>	void	Resumes WAL multiplexing
<code>pgx_is_wal_multiplexing_paused()</code>	boolean	Returns true if WAL multiplexing has stopped

If WAL multiplexing has not been configured, these functions return an error. Setting the `backup_destination` parameter in `postgresql.conf` configures WAL multiplexing.

Only superusers can execute these functions.

B.2 Transparent Data Encryption Control Functions

The following table lists the functions that can be used for transparent data encryption.

Table B.2 Transparent data encryption control functions

Name	Return type	Description
<code>pgx_open_keystore(<i>passphrase</i>)</code>	void	Opens the keystore
<code>pgx_set_master_key(<i>passphrase</i>)</code>	void	Sets the master encryption key
<code>pgx_set_keystore_passphrase(<i>oldPassphrase</i>, <i>newPassphrase</i>)</code>	void	Changes the keystore passphrase

B.2.1 `pgx_open_keystore`

The `pgx_open_keystore` function uses the specified passphrase to open the keystore. When the keystore is opened, the master encryption key is loaded into the database server memory. In this way, you can access the encrypted data and create encrypted tablespaces. If the keystore is already open, this function returns an error.

Only superusers can execute this function. Also, this function cannot be executed within a transaction block.

B.2.2 `pgx_set_master_key`

The `pgx_set_master_key` function generates a master encryption key and stores it in keystore.

If the keystore does not exist, this function creates a keystore. If the keystore already exists, this function modifies the master encryption key. If the keystore has not been opened, this function opens it.

The passphrase is a string of 8 to 200 bytes.

Only superusers can execute this function. Also, this function cannot be executed within a transaction block. Processing is not affected by whether the keystore is open.

B.2.3 pgx_set_keystore_passphrase

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 Profile Management Functions and User Management Functions

The table below lists the profile management functions and user management functions used in policy-based login security.

B.3.1 Profile Management Functions

Name	Return type	Description
<code>pgx_create_profile(profile_name name, password_parameter json)</code>	void	<p>Create a new profile.</p> <p>For <code>profile_name</code>, specify the profile name. An error will occur if an existing profile name is specified.</p> <p>For <code>password_parameter</code>, specify parameters and values in key-value format as follows.</p> <pre>{ "INACTIVE_USER_TIME": 15, "PASSWORD_LIFE_TIME": 30, "PASSWORD_GRACE_TIME": "UNLIMITED", "PASSWORD_REUSE_TIME": 10, "PASSWORD_REUSE_MAX": 5, "PASSWORD_LOCK_TIME": 0.5, "FAILED_LOGIN_ATTEMPTS": "DEFAULT", "PASSWORD_ALLOW_HASHED": true, "PASSWORD_ROLLOVER_TIME": 0.125 }</pre> <p>The json value accepts integer, numeric, boolean depending on the parameter, but only accepts strings "DEFAULT" and "UNLIMITED" as special values.</p> <p>For parameters that are omitted in json or that specify null for value, they follow the values in the default profile. Also, <code>password_parameter</code> can be omitted, and if omitted, a profile with all parameters conforming to the default profile will be created.</p>
<code>pgx_alter_profile(profile_name name, alter_parameter json)</code>	void	<p>Update the contents of an existing profile.</p> <p>For <code>profile_name</code>, specify the name of the profile to update. If you specify a profile name that does not exist, an error will occur.</p> <p>In <code>alter_parameter</code>, specify the variable name and value you want to change in key-value format as follows. Does not change the profile contents for parameters that are omitted or for which value is null.</p>

Name	Return type	Description
		<pre>'{ "name": "new_name", "PASSWORD_LIFE_TIME": 50, "PASSWORD_GRACE_TIME": 10 }';</pre> <p>name: Specify the modified profile name. null cannot be specified.</p> <p>If the default profile is specified in profile_name, the following changes cannot be made.</p> <ul style="list-style-type: none"> - Change name - Change each password_parameter to "DEFAULT"
pgx_drop_profile(profile_name name, if_exists boolean, cascade boolean)	void	<p>Delete an existing profile.</p> <p>Specify the profile name to be deleted in profile_name.</p> <p>Specifying the default profile will result in an error.</p> <p>if_exists specifies whether an error occurs when a nonexistent profile name is specified in profile_name. If true, no error. if_exists is optional. Default value is false.</p> <p>cascade specifies whether an error occurs when a profile assigned to a user is specified in profile_name. If false, an error will occur. If true, unassign all of the profiles before deleting them, and any unassigned users will be assigned the default profile instead. cascade is optional. Default value is false.</p> <p>[Example when specifying true] Specify arguments explicitly. (Example of named notation)</p> <pre>SELECT pgx_drop_profile('test_profile', cascade => true);</pre>

B.3.2 User Management Functions

Name	Return type	Description
pgx_assign_profile_to_user(user_name name, profile_name name)	void	<p>Assigns an existing profile to a user.</p> <p>Specify the assignee user in user_name. If you specify a user name that does not exist, an error will occur.</p> <p>Specify the profile name to be assigned in profile_name. If you specify a profile name that does not exist, an error will occur.</p>
pgx_lock_user(user_name name)	void	<p>Lock the user explicitly.</p> <p>Specify the user name in user_name. If you specify a user name that does not exist, an error will occur.</p> <p>If locked by this function, the lock period is indefinite. If you want to unlock it, you need to unlock it with the pgx_unlock_user function.</p> <p>If you lock a user that is already locked, the lock state is overwritten and becomes an indefinite lock even if it was previously a finite lock.</p>
pgx_unlock_user(user_name name)	void	<p>Explicitly unlock the user. The type of lock (profile reference lock or permanent lock) does not matter.</p> <p>Specify the user name in user_name. If you specify a user name that does not exist, an error will occur.</p>

Name	Return type	Description
		Specifying an unlocked user does not result in an error.
<code>pgx_make_password_expire(user_name name, expireat timestampz)</code>	void	Expires the specified user's password immediately. Alternatively, specify the time to expire. Specify the user name in <code>user_name</code> . If you specify a user name that does not exist, an error will occur. <code>expireat</code> can be a time that expires. <code>expireat</code> is optional. If omitted, execution time is specified.
<code>pgx_make_password_rollover_expire(user_name name)</code>	void	Expires immediately the period during which the specified user's old and new passwords can be used together so that the user cannot log in with the old password. Specify the user name in <code>user_name</code> . If you specify a user name that does not exist, an error will occur. If this command is executed outside the period for using the old and new passwords, it is ignored.

B.4 Data Masking Control Functions

The table below lists the functions that can be used for data masking.

Table B.3 Data masking control functions

Name	Return type	Description
pgx_alter_confidential_policy	boolean	Changes masking policies
pgx_create_confidential_policy	boolean	Creates masking policies
pgx_drop_confidential_policy	boolean	Deletes masking policies
pgx_enable_confidential_policy	boolean	Enables or disables masking policies
pgx_update_confidential_values	boolean	Changes replacement characters when full masking is specified for masking type

B.4.1 `pgx_alter_confidential_policy`

Description

Changes masking policies

Format

The format varies depending on the content to be changed. The format is shown below.

- Common format

```
common_arg:
[schema_name      := 'schemaName', ]


```

- Add a masking target to a masking policy

```
pgx_alter_confidential_policy(
commonArg,
[action           := 'ADD_COLUMN', ]
column_name      := 'colName'
```

```
[, function_type := 'FULL'] |
[, function_type := 'PARTIAL', partialOpt] |
[, function_type := 'REGEXP', regexpOpt]
)
```

```
partialOpt:
function_parameters := 'maskingFmt'
```

```
regexpOpt:
regexp_pattern := 'regexpPattern',
regexp_replacement := 'regexpReplacementChar',
[, regexp_flags := 'regexpFlags']
```

- Delete a masking target from a masking policy

```
pgx_alter_confidential_policy(
commonArg,
action := 'DROP_COLUMN',
column_name := 'colName'
)
```

- Change the masking condition

```
pgx_alter_confidential_policy(
commonArg,
action := 'MODIFY_EXPRESSION',
expression := 'expr'
)
```

- Change the content of a masking policy set for a masking target

```
pgx_alter_confidential_policy(
commonArg,
action := 'MODIFY_COLUMN',
column_name := 'colName'
[, function_type := 'FULL'] |
[, function_type := 'PARTIAL', partialOpt] |
[, function_type := 'REGEXP', regexpOpt]
)
```

```
partialOpt:
function_parameters := 'maskingFmt'
```

```
regexpOpt:
regexp_pattern := 'regexpPattern',
regexp_replacement := 'regexpReplacementChar',
[, regexp_flags := 'regexpFlags']
```

- Change the masking policy description

```
pgx_alter_confidential_policy(
commonArg,
action := 'SET_POLICY_DESCRIPTION',
policy_description := 'policyDesc'
)
```

- Change the masking target description

```
pgx_alter_confidential_policy(
commonArg,
action := 'SET_COLUMN_DESCRIPTION',
```

```

column_name      := 'colName',
column_description := 'colDesc'
)

```

Argument

The argument varies depending on the content to be changed. Details are as follows.

- Common arguments

Masking type for which an argument can be specified	Argument	Data type	Description	Default value
All	schema_name	varchar(63)	Schema name of table for which a masking policy is applied	'public'
	table_name	varchar(63)	Name of table for which a masking policy is applied	Mandatory
	policy_name	varchar(63)	Masking policy name	Mandatory

- Add a masking target to a masking policy

Masking type for which an argument can be specified	Argument	Data type	Description	Default value
All	action	varchar(63)	'ADD_COLUMN'	'ADD_COLUMN'
	column_name	varchar(63)	Masking target name	Mandatory
	function_type	varchar(63)	Masking type - 'FULL': Full masking - 'PARTIAL': Partial masking - 'REGEXP': Regular expression masking	'FULL'
Partial masking	function_parameters	varchar(1024)	Masking format for partial masking	Mandatory
Regular expression masking	regexp_pattern	varchar(1024)	Search pattern for regular expression masking	Mandatory
	regexp_replacement	varchar(1024)	Replacement character/string for regular expression masking	Mandatory
	regexp_flags	varchar(20)	Regular expression flags	NULL

- Delete a masking target from a masking policy

Masking type for which an argument can be specified	Argument	Data type	Description	Default value
All	action	varchar(63)	'DROP_COLUMN'	Mandatory
	column_name	varchar(63)	Masking target name	Mandatory

- Change the masking condition

Masking type for which an argument can be specified	Argument	Data type	Description	Default value
All	action	varchar(63)	'MODIFY_EXPRESSION'	Mandatory
	expression	varchar(1024)	Masking condition to be changed	Mandatory

- Change the content of a masking policy set for a masking target

Masking type for which an argument can be specified	Argument	Data type	Description	Default value
All	action	varchar(63)	'MODIFY_COLUMN'	Mandatory
	column_name	varchar(63)	Masking target name	Mandatory
	function_type	varchar(63)	Masking type - 'FULL': Full masking - 'PARTIAL': Partial masking - 'REGEXP': Regular expression masking	'FULL'
Partial masking	function_parameters	varchar(1024)	Masking format for partial masking	Mandatory
Regular expression masking	regexp_pattern	varchar(1024)	Search pattern for regular expression masking	Mandatory
	regexp_replacement	varchar(1024)	Replacement character/string for regular expression masking	Mandatory
	regexp_flags	varchar(20)	Regular expression flags	NULL

- Change the masking policy description

Masking type for which an argument can be specified	Argument	Data type	Description	Default value
All	action	varchar(63)	'SET_POLICY_DESCRIPTION'	Mandatory
	policy_description	varchar(1024)	Masking policy description	Mandatory

- Change the masking target description

Masking type for which an argument can be specified	Argument	Data type	Description	Default value
All	action	varchar(63)	'SET_COLUMN_DESCRIPTION'	Mandatory
	column_name	varchar(63)	Masking target name	Mandatory
	column_description	varchar(1024)	Masking target description	Mandatory

Details about whether arguments can be omitted are as follows.

Argument	Mandatory or optional									
	ADD_COLUMN			DROP_COLUMN	MODIFY_EXPRESSION	MODIFY_COLUMN			SET_POLICY_DESCRIPTION	SET_COLUMN_DESCRIPTION
	Full masking	Partial masking	Regular expression masking			Full masking	Partial masking	Regular expression masking		
schema_name	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
table_name	N	N	N	N	N	N	N	N	N	N
policy_name	N	N	N	N	N	N	N	N	N	N
action	Y	Y	Y	N	N	N	N	N	N	N
column_name	N	N	N	N	-	N	N	N	-	N
function_type	Y	N	N	-	-	Y	N	N	-	-
expression	-	-	-	-	N	-	-	-	-	-
policy_description	-	-	-	-	-	-	-	-	N	-
column_description	-	-	-	-	-	-	-	-	-	N
function_parameters	-	N	-	-	-	-	N	-	-	-
regexp_pattern	-	-	N	-	-	-	-	N	-	-
regexp_replacement	-	-	N	-	-	-	-	N	-	-
regexp_flags	-	-	Y	-	-	-	-	Y	-	-

Y: Can be omitted; N: Cannot be omitted; -: Ignored when specified

Return value

Return value	Description
TRUE	Ended normally
FALSE	Ended abnormally

Execution example 1

Adding masking policy p1 to masking target c2

```
postgres=# select pgx_alter_confidential_policy(table_name := 't1', policy_name := 'p1', action :=
'ADD_COLUMN', column_name := 'c2', function_type := 'PARTIAL', function_parameters := 'VVV-FVVV-FVVV, *, 4, 11');
pgx_alter_confidential_policy
-----
t
(1 row)
```

Execution example 2

Deleting masking target c1 from masking policy p1

```
postgres=# select pgx_alter_confidential_policy(table_name := 't1', policy_name := 'p1', action :=
'DROP_COLUMN', column_name := 'c1');
pgx_alter_confidential_policy
-----
```

```
t
(1 row)
```

Execution example 3

Changing the masking condition for masking policy p1

```
postgres=# select pgx_alter_confidential_policy(table_name := 't1', policy_name := 'p1', action :=
'MODIFY_EXPRESSION', expression := 'false');
pgx_alter_confidential_policy
-----
t
(1 row)
```

Execution example 4

Changing the content of masking policy p1 set for masking target c2

```
postgres=# select pgx_alter_confidential_policy(table_name := 't1', policy_name := 'p1', action :=
'MODIFY_COLUMN', column_name := 'c2', function_type := 'FULL');
pgx_alter_confidential_policy
-----
t
(1 row)
```

Execution example 5

Changing the description of masking policy p1

```
postgres=# select pgx_alter_confidential_policy(table_name := 't1', policy_name := 'p1', action :=
'SET_POLICY_DESCRIPTION', policy_description := 'this policy is an example. ');
pgx_alter_confidential_policy
-----
t
(1 row)
```

Execution example 6

Changing the description of masking target c2

```
postgres=# select pgx_alter_confidential_policy(table_name := 't1', policy_name := 'p1', action :=
'SET_COLUMN_DESCRIPTION', column_name := 'c2', column_description := 'c2 column is FULL. ');
pgx_alter_confidential_policy
-----
t
(1 row)
```

Description

- The arguments for the `pgx_alter_confidential_policy` system management function can be specified in any order.
- The action parameters below can be specified. When action parameters are omitted, `ADD_COLUMN` is applied.

Parameter	Description
<code>ADD_COLUMN</code>	Adds a masking target to a masking policy.
<code>DROP_COLUMN</code>	Deletes a masking target from a masking policy.
<code>MODIFY_EXPRESSION</code>	Changes expression.
<code>MODIFY_COLUMN</code>	Changes the content of a masking policy set for a masking target.
<code>SET_POLICY_DESCRIPTION</code>	Changes <code>policy_description</code> .

Parameter	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.4.2 pgx_create_confidential_policy

Description

Creates masking policies

Format

The format varies depending on the masking type. The format is shown below.

```
pgx_create_confidential_policy(
[schema_name      := 'schemaName',]
table_name       := 'tableName',
policy_name      := 'policyName',
expression       := 'expr'
[, enable        := 'policyStatus']
[, policy_description := 'policyDesc']
[, column_name    := 'colName'
  [, function_type := 'FULL'] |
  [, function_type := 'PARTIAL', partialOpt] |
  [, function_type := 'REGEXP', regexpOpt]
[, column_description := 'colDesc']
])
```

```
partialOpt:
function_parameters := 'maskingFmt'
```

```
regexpOpt:
regexp_pattern      := 'regexpPattern',
regexp_replacement  := 'regexpReplacementChar',
[, regexp_flags     := 'regexpFlags']
```

Argument

Details are as follows.

Masking type for which an argument can be specified	Argument	Data type	Description	Default value
All	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
	policy_name	varchar(63)	Masking policy name	Mandatory
	expression	varchar(1024)	Masking condition	Mandatory
	enable	boolean	Masking policy status - 't': Enabled - 'f': Disabled	't'
	policy_description	varchar(1024)	Masking policy description	NULL
	column_name	varchar(63)	Masking target name	NULL
	function_type	varchar(63)	Masking type - 'FULL': Full masking - 'PARTIAL': Partial masking - 'REGEXP': Regular expression masking	'FULL'
	column_description	varchar(1024)	Masking target description	NULL
Partial masking	function_parameters	varchar(1024)	Masking format for partial masking	Mandatory
Regular expression masking	regexp_pattern	varchar(1024)	Search pattern for regular expression masking	Mandatory
	regexp_replacement	varchar(1024)	Replacement character/string for regular expression masking	Mandatory
	regexp_flags	varchar(20)	Regular expression flags	NULL

Details about whether arguments can be omitted are as follows.

Argument	Mandatory or optional		
	Full masking	Partial masking	Regular expression masking
schema_name	Y	Y	Y
table_name	N	N	N
policy_name	N	N	N
expression	N	N	N
enable	Y	Y	Y
policy_description	Y	Y	Y
column_name	Y	Y	Y
function_type	Y	Y	Y
column_description	Y	Y	Y
function_parameters	-	N	-
regexp_pattern	-	-	N
regexp_replacement	-	-	N

Argument	Mandatory or optional		
	Full masking	Partial masking	Regular expression masking
regexp_flags	-	-	Y

Y: Can be omitted; N: Cannot be omitted; -: Ignored when specified

Return value

Return value	Description
TRUE	Ended normally
FALSE	Ended abnormally

Execution example 1

Creating masking policy p1 that does not contain a masking target

```
postgres=# select pgx_create_confidential_policy(table_name := 't1', policy_name := 'p1',
expression := 'l=1');
pgx_create_confidential_policy
-----
t
(1 row)
```

Execution example 2

Creating masking policy p1 that contains masking target c1 of which the masking type is full masking

```
postgres=# select pgx_create_confidential_policy(schema_name := 'public', table_name := 't1',
policy_name := 'p1', expression := 'l=1', enable := 't', policy_description := 'this policy is an
example.', column_name := 'c1', function_type := 'FULL', column_description := 'c1 column is FULL.');
```

```
pgx_create_confidential_policy
-----
t
(1 row)
```

Execution example 3

Creating masking policy p1 that contains masking target c2 of which the masking type is partial masking

```
postgres=# select pgx_create_confidential_policy( table_name := 't1', policy_name := 'p1',
expression := 'l=1', column_name := 'c2', function_type := 'PARTIAL', function_parameters :=
'VVVFVVVFVVVV, VVV-VVVV-VVVV, *, 4, 11');
```

```
pgx_create_confidential_policy
-----
t
(1 row)
```

Execution example 4

Creating masking policy p1 that contains masking target c3 of which the masking type is regular expression masking

```
postgres=# select pgx_create_confidential_policy( table_name := 't1', policy_name := 'p1',
expression := 'l=1', column_name := 'c3', function_type := 'REGEXP', regexp_pattern := '(.*)(@.*)',
regexp_replacement := 'xxx\2', regexp_flags := 'g');
```

```
pgx_create_confidential_policy
-----
t
(1 row)
```

Description

- The arguments for the `pgx_create_confidential_policy` system management function can be specified in any order.
- If `column_name` is omitted, only masking policies that do not contain masking target will be created.
- One masking policy can be created for each table. Use the `pgx_alter_confidential_policy` system management function to add a masking target to a masking policy.
- The `function_parameters` argument is enabled when the `function_type` is `PARTIAL`. If the `function_type` is other than `PARTIAL`, it will be ignored.
- The arguments below are enabled when the `function_type` is `REGEXP`. If the `function_type` is other than `REGEXP`, these arguments will be ignored.
 - `regexp_pattern`
 - `regexp_replacement`
 - `regexp_flags`
- If a table for which a masking policy is to be applied is deleted, delete the masking policy as well.



- Refer to "String Constants" in the PostgreSQL Documentation for information on the strings to specify for arguments.
- 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.4.3 `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
<code>schema_name</code>	<code>varchar(63)</code>	Schema name of table for which a masking policy is deleted	'public'
<code>table_name</code>	<code>varchar(63)</code>	Name of table for which a masking policy is deleted	Mandatory
<code>policy_name</code>	<code>varchar(63)</code>	Masking policy name	Mandatory

Details about whether arguments can be omitted are as follows.

Argument	Mandatory or optional
<code>schema_name</code>	Y
<code>table_name</code>	N
<code>policy_name</code>	N

Y: Can be omitted; N: Cannot be omitted

Return value

Return value	Description
TRUE	Ended normally
FALSE	Ended abnormally

Execution example

Deleting masking policy p1

```
postgres=# select pgx_drop_confidential_policy(table_name := 't1', policy_name := 'p1');
pgx_drop_confidential_policy
-----
t
(1 row)
```

Description

- The arguments for the `pgx_drop_confidential_policy` system management function can be specified in any order.
- If a table for which a masking policy is to be applied is deleted, delete the masking policy as well.



See

Refer to "String Constants" in the PostgreSQL Documentation for information on the strings to specify for arguments.

B.4.4 `pgx_enable_confidential_policy`

Description

Enables or disables masking policies

Format

```
pgx_enable_confidential_policy(
[schema_name      := 'schemaName', ]
table_name       := 'tableName',
policy_name      := 'policyName',
enable           := 'policyStatus'
)
```

Argument

Details are as follows.

Argument	Data type	Description	Default value
<code>schema_name</code>	<code>varchar(63)</code>	Schema name of table for which a masking policy is enabled or disabled	'public'
<code>table_name</code>	<code>varchar(63)</code>	Name of table for which a masking policy is enabled or disabled	Mandatory
<code>policy_name</code>	<code>varchar(63)</code>	Masking policy name	Mandatory
<code>enable</code>	boolean	Masking policy status - 't': Enabled - 'f': Disabled	Mandatory

Details about whether arguments can be omitted are as follows.

Argument	Mandatory or optional
schema_name	Y
table_name	N
policy_name	N
enable	N

Y: Can be omitted; N: Cannot be omitted

Return value

Return value	Description
TRUE	Ended normally
FALSE	Ended abnormally

Execution example

Enabling masking policy p1

```
postgres=# select pgx_enable_confidential_policy(table_name := 't1', policy_name := 'p1', enable :=
't');
pgx_enable_confidential_policy
-----
t
(1 row)
```

Description

The arguments for the pgx_enable_confidential_policy system management function can be specified in any order.



See

Refer to "String Constants" in the PostgreSQL Documentation for information on the strings to specify for arguments.

B.4.5 pgx_update_confidential_values

Description

Changes replacement characters when full masking is specified for masking type

Format

```
pgx_update_confidential_values(
[number_value := 'numberValue']
[, char_value := 'charValue']
[, varchar_value := 'varcharValue']
[, date_value := 'dateValue']
[, ts_value := 'tsValue']
)
```

Argument

Details are as follows.

Argument	Data type	Description
number_value	integer	Replacement character in numeric type
char_value	varchar(1)	Replacement character in char type
varchar_value	varchar(1)	Replacement character in varchar type
date_value	date	Replacement character in date type
ts_value	timestamp	Replacement character in timestamp type

Return value

Return value	Description
TRUE	Ended normally
FALSE	Ended abnormally

Execution example

Using '*' as a replacement character in char type and varchar type

```
postgres=# select pgx_update_confidential_values(char_value := '*', varchar_value := '*');
pgx_update_confidential_values
-----
t
(1 row)
```

Description

- The arguments for the pgx_update_confidential_values system management function can be specified in any order.
- Specify one or more arguments for the pgx_update_confidential_values system management function. A replacement character is not changed for an omitted argument.



See

Refer to "String Constants" in the PostgreSQL Documentation for information on the strings to specify for arguments.

B.5 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.6 High-Speed Data Load Control Functions

The table below lists the functions that can be used for high-speed data load.

Table B.5 High-speed data load control functions

Name	Return type	Description
pgx_loader	bigint	Creates dynamic shared memory, starts parallel workers and loads data
pgx_loader_recovery	smallint	Resolves in-doubt transactions

The pgx_loader command executes the above functions internally.

Appendix C System Catalogs

Describes the system catalog of Fujitsu Enterprise Postgres.



See

Refer to "System Catalogs" under "Internals" in the PostgreSQL Documentation for information on other system catalogs.

C.1 pgx_profile

A system catalog for managing profile information. -1 is stored if DEFAULT is specified for each parameter in the profile, and -2 is stored if UNLIMITED is specified.

Column	Type	Description
oid	oid	Profile identifier (Primary key constraint)
prfname	name	Profile name (Unique constraint, NOT NULL constraint)
prfpasswordlifetime	integer	Value of PASSWORD_LIFE_TIME (seconds) (NOT NULL constraint)
prfpasswordgracetime	integer	Value of PASSWORD_GRACE_TIME (seconds) (NOT NULL constraint)
prfpasswordreusetime	integer	Value of PASSWORD_REUSE_TIME (seconds) (NOT NULL constraint)
prfpasswordreusemax	integer	Value of PASSWORD_REUSE_MAX (seconds) (NOT NULL constraint)
prfpasswordlocktime	integer	Value of PASSWORD_LOCK_TIME (seconds) (NOT NULL constraint)
prffailedloginattempts	integer	Value of FAILED_LOGIN_ATTEMPTS (seconds) (NOT NULL constraint)
prfpasswordallowhashed	integer	Value of PASSWORD_ALLOW_HASHED (NOT NULL constraint)
prfinactiveusertime	integer	Value of INACTIVE_USER_TIME (seconds) (NOT NULL constraint)
prfpasswordrollovertime	integer	Value of PASSWORD_ROLLOVER_TIME (seconds) (NOT NULL constraint)

C.2 pgx_user_profile

A system catalog that manages profile-related information associated with users. When DROP ROLE, the corresponding row is automatically deleted.

Column	Type	References	Description
userprfroleid	oid	pg_authid.oid	User's identifier (Primary key constraint)
userprfprfid	oid	pgx_profile.oid	Identifier of the profile assigned to the user
userprfaccountlock	smallint		User lock state (NOT NULL constraint) 0: Not locked

Column	Type	References	Description
			1: Lock (Refer to PASSWORD_LOCK_TIME in profile for duration) 2: Lock (Unlimited)
userprfpasswordstatus	char		Password status Updated if there are any changes at login (NOT NULL constraint) o: Current valid password g: In grace period e: Expired
userprflockdate	timestamp with time zone		Time the user was locked
userprfpasswordsetat	timestamp with time zone		Time you updated your current password
userprfpasswordexpire	timestamp with time zone		Password expiration time
userprflastactivetime	timestamp with time zone		Time the user's session was last checked
userprfpasswordrolloverexpire	timestamp with time zone		Time at which the period in which old and new passwords can be used ends

C.3 pgx_auth_password

The system catalog for storing user-associated password information. Used for password rollover.

Since this catalog contains passwords, it must not be readable by third parties.

Column	Type	References	Description
authpwdroleid	oid	pg_authid.oid	User's identifier (Primary key constraint)
authpwdoldpassword	text		Hashed password Used for authentication when password rollover is enabled.
authpwdoldpasswordsetat	timestamp with time zone		Time at which authpwdoldpassword was set
authpwdnewpassword	text		Hashed password Used for authentication when password rollover is enabled.
authpwdnewpasswordsetat	timestamp with time zone		Time at which authpwdnewpassword was set

C.4 pgx_password_history

A system catalog for managing password history. Used for password reuse confirmation and password rollover. When DROP ROLE is executed, the corresponding row is automatically deleted.

Since this catalog contains passwords, it must not be readable by third parties.

This system catalog is updated when the password is updated, and if the password reuse limit is finite, unnecessary history deletion is also performed. Old passwords that have been removed from the catalog are treated as having never been used.

passhistpassword is saved in the format specified by the password_encryption parameter.

Column	Type	Description
passhistroleid	oid	Identifier of the user whose password was set Set unique constraint with (passhistroleid, passhistpassword).
passhistpassword	text	Encrypted password Save rolpassword value in pg_authid as history (NOT NULL constraint)
passhistpasswordsetat	timestamp with time zone	Time when the password was updated (NOT NULL constraint)

Appendix D 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.

D.1 pgx_tablespaces

The `pgx_tablespaces` view provides information related to the encryption of tablespaces.

Table D.1 `pgx_tablespaces` view

Column	Type	References	Description
<code>spctablespace</code>	<code>oid</code>	<code>pg_tablespace.oid</code>	Tablespace OID
<code>spcencalgo</code>	<code>text</code>		Tablespace encryption algorithm

The `spcencalgo` string displays one of the following values:

- none: Tablespace is not encrypted
- AES128: AES with key length of 128 bits
- AES256: AES with key length of 256 bits

D.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 D.2 `pgx_stat_lwlock` view

Column	Type	Description
<code>lwlock_name</code>	<code>name</code>	Name of the lightweight lock
<code>total_waits</code>	<code>bigint</code>	Number of waits caused by the lightweight lock
<code>total_wait_time</code>	<code>double precision</code>	Number of milliseconds spent in waits caused by the lightweight lock
<code>stats_reset</code>	<code>timestamp with timezone</code>	Last time at which this statistics was reset

D.3 pgx_stat_latch

The `pgx_stat_latch` view displays statistics related to latches, with each type of wait information within Fujitsu Enterprise Postgres displayed on a separate line.

Table D.3 `pgx_stat_latch` view

Column	Type	Description
<code>latch_name</code>	<code>name</code>	Name of the latch
<code>total_waits</code>	<code>bigint</code>	Number of waits caused a wait
<code>total_wait_time</code>	<code>double precision</code>	Number of milliseconds spent in waits caused by the latch
<code>stats_reset</code>	<code>timestamp with timezone</code>	Last time at which this statistic was reset

D.4 pgx_stat_walwriter

The `pgx_stat_walwriter` view displays statistics related to WAL writing, in a single line.

Table D.4 `pgx_stat_walwriter` view

Column	Type	Description
<code>dirty_writes</code>	<code>bigint</code>	Number of times old WAL buffers were written to the disk because the WAL buffer was full when WAL records were added
<code>writes</code>	<code>bigint</code>	Number of WAL writes
<code>write_blocks</code>	<code>bigint</code>	Number of WAL write blocks
<code>total_write_time</code>	<code>double precision</code>	Number of milliseconds spent on WAL writing
<code>stats_reset</code>	<code>timestamp with timezone</code>	Last time at which this statistic was reset

D.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 D.5 `pgx_stat_sql` view

Column	Type	Description
<code>selects</code>	<code>bigint</code>	Number of SELECT statements executed In database multiplexing mode, this number includes the SELECT statements executed in Mirroring Controller. Mirroring Controller executes the SELECT statement using the interval specified for the <code>heartbeat_interval</code> of the server definition file (milliseconds).
<code>inserts</code>	<code>bigint</code>	Number of INSERT statements executed
<code>deletes</code>	<code>bigint</code>	Number of DELETE statements executed
<code>updates</code>	<code>bigint</code>	Number of UPDATE statements executed
<code>selects_with_parallelism</code>	<code>bigint</code>	Number of times parallel scan was used in SELECT statements
<code>inserts_with_parallelism</code>	<code>bigint</code>	Not used
<code>deletes_with_parallelism</code>	<code>bigint</code>	Not used
<code>updates_with_parallelism</code>	<code>bigint</code>	Not used
<code>copies_with_parallelism</code>	<code>bigint</code>	Not used
<code>declares</code>	<code>bigint</code>	Number of DECLARE statements executed (number of cursor OPENS)
<code>fetches</code>	<code>bigint</code>	Number of FETCH statements executed
<code>checkpoints</code>	<code>bigint</code>	Number of CHECKPOINT statements executed
<code>clusters</code>	<code>bigint</code>	Number of CLUSTER statements executed
<code>copies</code>	<code>bigint</code>	Number of COPY statements executed
<code>reindexes</code>	<code>bigint</code>	Number of REINDEX statements executed
<code>truncates</code>	<code>bigint</code>	Number of TRUNCATE statements executed
<code>locks</code>	<code>bigint</code>	Number of times a lock occurred
<code>stats_reset</code>	<code>timestamp with timezone</code>	Last time at which this statistic was reset

D.6 pgx_stat_gmc

The `pgx_stat_gmc` view provides information about the GMC areas.

Table D.6 `pgx_stat_gmc` view

Column	Type	Description
<code>searches</code>	<code>bigint</code>	Number of times the cache table is searched.
<code>hits</code>	<code>bigint</code>	Number of times the cache table is hit.
<code>size</code>	<code>bigint</code>	The current amount of memory (bytes) used in the GMC area.
<code>stats_reset</code>	<code>timestamp with timezone</code>	Last time these statistics were reset.

D.7 pgx_stat_progress_loader

The `pgx_stat_progress_loader` view provides overall progress information for `pgx_loader` command.

The `pgx_stat_progress_loader` view displays the sum of the progress information of the back-end processes and the number of parallel worker processes when `pgx_loader` runs.

Table D.7 `pgx_stat_progress_loader` view

Column	Type	Description
<code>pid</code>	<code>integer</code>	Process ID of the backend.
<code>datid</code>	<code>Oid</code>	Oid of the database to connect to.
<code>datname</code>	<code>name</code>	Name of the database to connect to.
<code>relid</code>	<code>Oid</code>	Oid of the table to load.
<code>command</code>	<code>text</code>	Command executes the load process. (Always "COPY FROM" for <code>pgx_loader</code>)
<code>type</code>	<code>text</code>	Type of data source for the load operation.
<code>bytes_processed</code>	<code>bigint</code>	Size of the data at the end of the load. (Backend and worker totals)
<code>bytes_total</code>	<code>bigint</code>	Size of the data to load. (Backend and worker totals)
<code>tuples_processed</code>	<code>bigint</code>	Number of rows that have completed loading. (Backend and worker totals)
<code>tuples_excluded</code>	<code>bigint</code>	Number of rows skipped during the load process. (Backend and worker totals)

Column	Type	Description
table_name	varchar(63)	Name of table for which a masking policy is applied
policy_name	varchar(63)	Masking policy name
expression	varchar(1024)	Masking condition
enable	boolean	Masking policy status - 't': Enabled - 'f': Disabled
policy_description	varchar(1024)	Masking policy description

Execution example

```
postgres=# select * from pgx_confidential_policies;
 schema_name | table_name | policy_name | expression | enable | policy_description
-----+-----+-----+-----+-----+-----
 public      | t1         | p1          | 1=1        | t      |
(1 row)
```

E.3 pgx_confidential_values

This table provides information on replacement characters when full masking is specified for masking type.

Column	Data type	Description	Default value
number_value	integer	Numeric	0
char_value	varchar(1)	char type	Spaces
varchar_value	varchar(1)	varchar type	Spaces
date_value	date	date type	'1970-01-01'
timestamp_value	timestamp	timestamp type	'1970-01-01 00:00:00'

Execution example

```
postgres=# select * from pgx_confidential_values;
 number_value | char_value | varchar_value | date_value | ts_value
-----+-----+-----+-----+-----
          0 |           |               | 1970-01-01 | 1970-01-01 00:00:00
(1 row)
```

Appendix F Tables Used by Aggressive Freeze for Tuples

This appendix explains tables used by the aggressive freeze for tuples feature.

F.1 pgx_stat_freeze_results

The table `pgx_stat_freeze_results` logs statistics about aggressive freeze for tuples. If an aggressive freeze for tuples is performed, that information will be added as an accumulation. The cumulative information is summarized for the time period specified by the `pgx_stat_vacuum_freeze.monitoring_interval` parameter, and the information for the time period specified by the `pgx_stat_vacuum_freeze.log_retention_period` parameter is kept. For the information about the `pgx_stat_vacuum_freeze.monitoring_interval` and `pgx_stat_vacuum_freeze.log_retention_period` parameters, refer to "[Parameters for scheduling aggressive freeze for tuples](#)".

You can gather information to determine how long the aggressive freeze for tuples process is taking, and help estimate the aggressive freeze for tuples schedule.

Column	Type	Description
<code>measurement_start_time</code>	timestamp with time zone	Start time of the measurement
<code>measurement_end_time</code>	timestamp with time zone	End time of the measurement
<code>vacuum_count</code>	bigint	Number of times the freeze operation was performed
<code>autovacuum</code>	boolean	true for autovacuum driven vacuum
<code>scan_pages</code>	bigint	Accumulation of pages scanned during freezing
<code>frozen_pages</code>	bigint	Accumulation of frozen pages
<code>frozen_tuples</code>	bigint	Accumulation of frozen tuples
<code>total_frozen_wal_records</code>	bigint	Accumulation of the number of WAL records produced by freezing
<code>frozen_full_page_image_wal_records</code>	bigint	Accumulation of the number of WAL full-page images produced by freezing
<code>wal_total_bytes</code>	numeric	Accumulation, in bytes, of WAL generated by freezing
<code>total_elapsed_time</code>	double precision	Accumulation of freezing time. Units are milliseconds
<code>sleep_times</code>	double precision	Accumulation of time spent sleeping during freezing. Units are milliseconds

Appendix G Tables Used by High-Speed Data Load

This appendix describes the tables used by high-speed data load.

G.1 pgx_loader_state

The pgx_loader_state table provides information about transactions prepared by high-speed data load.

Column	Type	Description
id	serial	Unique identifier. This value is assigned from the pgx_loader_state_id_seq sequence.
gid	text	Global transaction identifier assigned to a transaction.
state	text	State of the transaction. The value can be one of the following: <ul style="list-style-type: none">- commit: The prepared transaction has been committed.- rollback: The prepared transaction is in in-doubt state.
leader_pid	integer	Process ID of the backend process (leader process) that executed the pgx_loader control function.
role_oid	integer	Role identifier (OID). A prepared transaction can only be completed by the same user who executed the original transaction or by a superuser.
relation_oid	integer	Object identifier (OID).

Information

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 H 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.



See

Refer to "Installing WebAdmin in a Multiserver Configuration" in the Installation and Setup Guide for Server for information on multiserver installation.

H.1 Starting the Web Server Feature of WebAdmin

Follow the procedure below to start the Web server feature of WebAdmin:

1. Display the [Services] window

In the [Administrative Tools], click [Services].

2. Start a service

Select the displayed name "Fujitsu Enterprise Postgres WebAdmin *version*", and then click [Start Service].

You can also start a service by specifying the service name of the Web server feature of WebAdmin in the net start command or sc start command.

H.2 Stopping the Web Server Feature of WebAdmin

This section describes how to stop the Web server feature of WebAdmin.

Follow the procedure below to stop the Web server feature of WebAdmin:

1. Display the [Services] window

In the [Administrative Tools], click [Services].

2. Stop a service

Select the displayed name "Fujitsu Enterprise Postgres WebAdmin *version*", and then click [Stop Service].

You can also stop a service by specifying the service name of the Web server feature of WebAdmin in the net stop command or sc stop command.

Appendix I 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.

I.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.

Credential name, User name and Password should not contain hazardous characters. Refer to “[Appendix J WebAdmin Disallow User Inputs Containing Hazardous Characters](#)”.

- [Credential name]: Name of the credential

The name must meet the conditions below:

- Maximum of 16 characters
- The first character must be an ASCII alphabetic character
- The other characters must be ASCII alphanumeric characters

- [User name]: The operating system user name or database instance user name that will be used later

- [Password]: Password for the user

- [Confirm password]: Reenter the password.

3. Click  to store the credential.

I.2 Using a Credential

Once a credential is created in the Wallet, it can be used during remote instance creation or standby instance creation.

If you select the credential saved in "I.1 Creating a Credential" in [Operating system credential], the user name and password will be automatically populated.

Appendix J 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 K Collecting Failure Investigation Data

For information on collecting failure investigation data, refer to "Collecting Failure Investigation Data" in the Installation and Setup Guide for Server.

Index

	[A]	
About using pgAdmin.....		3
About using WebAdmin.....		2
Actions in Response to Instance Startup Failure.....		132
All user data within the specified tablespace.....		25
Approximate backup time.....		16
Approximate recovery time.....		106
Automatically opening the keystore.....		35
	[B]	
Backing Up and Recovering the Keystore.....		30
Backing Up and Restoring/Recovering the Database.....		32
Backup/Recovery Using the Copy Command.....		97
Backup and recovery using the pgx_dmpall and pgx_rcvall commands.....		32
backup cycle.....		18
Backup data.....		25
Backup operation.....		18
Backup operation (file backup).....		20
Backup status.....		18,20
Backup using the backup information file.....		98
Backup Using the Copy Command.....		100
backup_destination (string).....		139
Building and starting a standby server.....		36
	[C]	
Changing a Masking Policy.....		50
Changing the Keystore Passphrase.....		29
Changing the Master Encryption Key.....		29
Changing the master encryption key and the passphrase.....		36
Checking an Encrypted Tablespace.....		28
Checking backup status.....		101
Checking the operating state of an instance.....		14
Checking the operating status of an instance.....		13
Collecting Failure Investigation Data		177
Configuration of the Copy Command.....		97
Configuration of the copy command for backup.....		99
Configuration of the copy command for recovery.....		99
Confirming a Masking Policy.....		51
Continuous archiving and point-in-time recovery.....		33
Copy Command for Backup.....		102
Copy Command for Recovery.....		104
Copy Command Interface.....		102
core_contents (string).....		139
core_directory (string).....		139
Creating a Masking Policy.....		50
Cyclic usage of the backup area.....		97
	[D]	
Data Masking.....		45
Data Types for Masking.....		53
Deleting a Masking Policy.....		52
Determining the backup area of the latest backup.....		101
	[E]	
Enabling and Disabling a Masking Policy.....		51
Enabling Automatic Opening of the Keystore.....		30
Encrypting a Tablespace.....		27
Encrypting Existing Data.....		34
Encryption mechanisms.....		25
Errors in More Than One Storage Disk.....		131
	[F]	
Faster encryption and decryption based on hardware.....		25
File system level backup and restore.....		33
	[H]	
High-Speed Data Load.....		86
	[I]	
If failure occurred in the data storage disk or the transaction log storage disk.....		107
If failure occurred on the backup data storage disk.....		108,109
If failure occurred on the data storage disk or the transaction log storage directory.....		108
Importing and Exporting the Database.....		34
Installing and Operating the In-memory Feature.....		76
	[K]	
keystore_location (string).....		139
	[L]	
Logging in to WebAdmin.....		3
log in.....		3
	[M]	
Managing the Keystore.....		29
Masking Condition.....		46
Masking Format.....		46
Masking Policy.....		45
Masking Target.....		45
Masking Type.....		46
Monitoring Database Activity.....		58
	[O]	
Opening the Keystore.....		27
openssl_conf(string).....		145
openssl_modules(string).....		145
Operating Fujitsu Enterprise Postgres.....		1
	[P]	
Parallel Query.....		84
Performing backup.....		100
Perform recovery.....		101
Periodic Backup.....		18
pgx_global_metacache (numerical value).....		142
pgx_stat_gmc view.....		169
pgx_stat_latch view.....		167
pgx_stat_lwlock view.....		167
pgx_stat_progress_loader view.....		169
pgx_stat_sql view.....		168
pgx_stat_vacuum_freeze.log_retention_period (numerical value).....		146

pgx_stat_vacuum_freeze.monitoring_interval (numerical value)	146	WebAdmin Wallet.....	175
pgx_stat_vacuum_freeze.track_freeze_info (boolean).....	146		
pgx_stat_walwriter view.....	168		
pgx_tablespace.....	167		
pgx_tablespace view.....	167		
Placement and automatic opening of the keystore file.....	35		
Placing the keystore file.....	35		
Preparing for backup.....	100		
[R]			
Recovery Using the Copy Command.....	101		
reserve_buffer_ratio (numerical value).....	140		
[S]			
Scope of encryption.....	25		
search_path (string).....	139		
Security-Related Notes.....	36		
Security Notes.....	53		
Setting a restore point.....	20		
Setting the Master Encryption Key.....	26		
Starting and Stopping the Web Server Feature of WebAdmin			
.....	174		
Starting an instance.....	12,13		
Starting pgAdmin.....	3		
Startup URL for WebAdmin.....	3		
Stopping an instance.....	12,14		
Streaming replication support.....	26		
Streaming Replication Using WebAdmin.....	72		
Strong encryption algorithms.....	25		
System Administration Functions.....	147		
System Views.....	167		
[T]			
tablespace_encryption_algorithm (string).....	139		
Tables Used by Data Masking	170		
Tips for Installing Built Applications.....	36		
track_gmc (string).....	142		
track_sql (string).....	140		
track_waits (string).....	140		
Transparent Data Encryption Control Functions.....	147		
Two-layer encryption key and the keystore.....	25		
[U]			
userprofile_database(string).....	145		
[V]			
vci.control_max_workers (numerical value).....	140		
vci.cost_threshold (numeric).....	140		
vci.enable (string).....	141		
vci.log_query (string).....	141		
vci.maintenance_work_mem (numerical value).....	141		
vci.max_local_ros (numerical value).....	141		
vci.max_parallel_degree (numerical value).....	141		
vci.shared_work_mem (numerical value).....	141		
[W]			
WAL and temporary files.....	25		
WAL Mirroring Control Functions.....	147		