


Fujitsu Enterprise Postgres 18 SP2  
for x86

# Security Operation Guide

Windows/Linux

The background of the entire page is a photograph of several red ants on a piece of tree bark. One ant is carrying a small green leaf fragment. In the background, there are two large, stylized arches in shades of orange and yellow, set against a blurred green background of foliage.

J2UL-3058-03ENZ0(00)

June 2026

# Preface

---

## Purpose of this document

This document describes security when building and operating a Fujitsu Software Enterprise Postgres database system.

## Intended readers

This document is intended for those who are:

- Considering installing Fujitsu Enterprise Postgres
- Designing, building, and operating the security operating environment in Fujitsu Enterprise Postgres
- Accessing Fujitsu Enterprise Postgres database systems

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

- Business operations
- Fujitsu Enterprise Postgres
- Linux

## Structure of this document

This document is structured as follows:

### [Chapter 1 Overview of Security](#)

Provides an overview of the security system, and explains the security features provided by Fujitsu Enterprise Postgres.

### [Chapter 2 Overview of Security Operation](#)

Provides an overview of security operation.

### [Chapter 3 Tasks of the Manager](#)

Explains the tasks for security measures to be implemented by the manager.

### [Chapter 4 Tasks of Administrators](#)

Explains the tasks for security measures to be implemented by administrators.

### [Chapter 5 Tasks of Users](#)

Explains the tasks for security measures to be implemented by users.

### [Chapter 6 Privileged Access Management](#)

Explains the privileged access management provided by Fujitsu Enterprise Postgres.

### [Chapter 7 Audit Log Feature](#)

Explains the audit log feature provided by Fujitsu Enterprise Postgres.

### [Chapter 8 Confidentiality Management](#)

Explains the confidentiality management feature provided by Fujitsu Enterprise Postgres.

### [Appendix A Tables Used by Confidentiality Management Feature](#)

Explains the tables used by the confidentiality management feature.

### [Appendix B System Management Functions Used by Confidentiality Management Feature](#)

Explains the functions used by the confidentiality management feature.

## References

This document contains abstracts from the following document:

- Database Security Guideline Version 2.0  
(Database Security Consortium (DBSC))

### **Export restrictions**

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

### **Issue date and version**

Edition 3.0: June 2026 Edition 2.0: March 2026 Edition 1.0: December 2025
---

### **Copyright**

Copyright 2018-2026 Fujitsu Limited

# Contents

---

Chapter 1 Overview of Security.....	1
1.1 What is Security?.....	1
1.2 Security Requirements.....	1
1.3 Security Threats.....	2
1.4 Security Scope.....	5
1.5 Security Provided by Fujitsu Enterprise Postgres.....	6
1.5.1 Roles Targeted For Security.....	6
1.5.2 Security Features.....	7
Chapter 2 Overview of Security Operation.....	9
2.1 Security Operation Flow.....	9
Chapter 3 Tasks of the Manager.....	11
3.1 Defining Important Information and Risk Analysis.....	11
3.2 Formulating Account Management Policies.....	11
3.3 Formulating Log Retrieval Policies.....	11
3.4 Formulating Rules.....	12
3.5 Implementing Training.....	13
3.6 Checking the Database Management Operations.....	13
3.7 Periodic Diagnosis of the Status of Security Measures.....	13
Chapter 4 Tasks of Administrators.....	14
4.1 Receiving Training.....	14
4.2 Initial Setup.....	14
4.3 Authentication.....	15
4.3.1 Managing Accounts.....	15
4.3.2 Managing Passwords.....	16
4.3.3 Configuring Connections and Authentication.....	17
4.4 Access Control.....	17
4.5 Encryption.....	18
4.6 Controlling Use of External Media.....	18
4.7 Security Measures for Servers/Applications.....	19
4.8 Log Management.....	19
4.8.1 Retrieving Logs.....	19
4.8.2 Maintaining Logs.....	19
4.9 Detecting Unauthorized Access.....	20
4.10 Analyzing Logs.....	20
Chapter 5 Tasks of Users.....	21
5.1 Receiving Training.....	21
5.2 Managing Accounts/Passwords.....	21
Chapter 6 Privileged Access Management.....	22
6.1 Controlling Elevation via set_user.....	22
6.1.1 Setup for set_user.....	22
6.1.2 Constraints.....	23
Chapter 7 Audit Log Feature.....	24
7.1 Audit Log Output Modes.....	24
7.2 Setup.....	24
7.3 Setting Up the Scalable Audit Log Feature.....	27
7.4 pgaudit Configuration File.....	29
7.5 Session Audit Logging.....	35
7.6 Object Audit Logging.....	41
7.7 Database Multiplexing.....	43
7.7.1 Setup.....	43
7.7.2 Configuring Audit Log Retrieval.....	44

7.8 Failover.....	44
7.8.1 Configuring Audit Log Retrieval.....	44
7.9 When Used with set_user.....	45
7.10 Analyzing Audit Logs in SQL.....	46
7.11 Removing Setup.....	48
<b>Chapter 8 Confidentiality Management.....</b>	<b>50</b>
8.1 Setup.....	50
8.2 Designing Confidentiality Management.....	50
8.2.1 Designing a Confidentiality Matrix.....	50
8.2.1.1 Defining Confidentiality Levels.....	51
8.2.1.2 Defining Confidentiality Groups.....	52
8.2.1.3 Defining Confidentiality Privilege.....	53
8.2.2 Determining Confidentiality Management Roles.....	54
8.2.3 Classify Confidentiality Objects According to the Definition of Confidentiality Level.....	55
8.2.3.1 Defining Confidentiality Objects.....	55
8.2.3.2 Classify Confidentiality Objects.....	55
8.2.4 Classify Roles According to Confidentiality Group Definitions.....	56
8.3 How to Use Confidentiality Management Feature (Definition).....	56
8.3.1 Creating a Confidentiality Management Role.....	59
8.3.2 Creating a Confidentiality Matrix.....	59
8.3.3 Adding Confidentiality Levels to the Confidentiality Matrix.....	60
8.3.4 Adding Confidentiality Groups to the Confidentiality Matrix.....	61
8.3.5 Granting Confidentiality Privileges to Confidentiality Groups.....	61
8.3.6 Adding Confidentiality Objects to Confidentiality Level.....	61
8.3.7 Adding Roles to Confidentiality Groups.....	62
8.4 How to Use Confidentiality Management Feature (Change and Deletion).....	62
8.4.1 Renaming Confidentiality Objects.....	62
8.4.2 Renaming Roles.....	62
8.4.3 Deleting Roles.....	62
8.4.4 Changing Confidentiality Matrix.....	63
8.4.5 Deleting Confidentiality Matrix.....	63
8.4.6 Changing Confidentiality Level.....	63
8.4.7 Deleting Confidentiality Level.....	63
8.4.8 Changing Confidentiality Group.....	63
8.4.9 Deleting Confidentiality Group.....	63
8.4.10 Revoking Confidentiality Privileges.....	64
8.4.11 Removing Confidentiality Objects from Confidentiality Level.....	64
8.4.12 Removing Roles from Confidentiality Groups.....	64
8.5 Suggestions for Monitoring Methods.....	64
8.5.1 How to Detect Privilege Changes without Using Confidentiality Management feature.....	65
8.5.2 How to Check Confidentiality Objects and Roles.....	65
8.6 Backup/Restore.....	66
8.7 Removing Setup.....	67
8.8 Usage Example of Confidentiality Management.....	67
<b>Appendix A Tables Used by Confidentiality Management Feature.....</b>	<b>70</b>
A.1 pgx_confidential_matrix.....	70
A.2 pgx_confidential_level.....	70
A.3 pgx_confidential_group.....	71
A.4 pgx_confidential_privilege.....	71
A.5 pgx_confidential_object.....	72
A.6 pgx_confidential_role.....	73
A.7 pgx_confidential_policy.....	73
<b>Appendix B System Management Functions Used by Confidentiality Management Feature.....</b>	<b>75</b>
B.1 Confidentiality Matrix Manipulation Functions.....	75
B.2 Confidentiality Level Manipulation Functions.....	76


B.3 Confidentiality Group Manipulation Functions..... 78  
B.4 Confidentiality Privilege Manipulation Functions..... 80  
B.5 Confidentiality Object Manipulation Functions..... 81  
B.6 Role Manipulation Functions..... 84  
B.7 Functions that Support Definition Referencing and Comparison with System Catalogs..... 85

# Chapter 1 Overview of Security

## 1.1 What is Security?

Computer security is the protection of information systems and data from risks such as leakage or tampering of information, attacks, intrusions, eavesdropping from external sources, and interference with information services. Security measures are essential for the advance prevention of security threats in order for information systems to gain trust as social infrastructure.

Figure 1.1 Security threats



The security measures in information systems can be classified as follows:

- Network
- Web
- Application
- Database
- PC

This document focuses on database security measures when using Fujitsu Enterprise Postgres.

## 1.2 Security Requirements

Below are the necessary security requirements for information systems.

### Maintenance of security policies

A security policy clarifies the approach the company should take in relation to information assets, and the actions employees should take.

It is necessary to undertake security of information systems while maintaining security policies.

### Integrated security management

Security has the aspects below. It is necessary to manage information in an integrated manner based on these aspects.

### Confidentiality

Access to the information is restricted to prevent leakage of information outside of the company

Example measures: Prevention of information leakage or setup of access privileges

### Integrity

Integrity is guaranteed, ensuring information does not become corrupted or tampered with

Example measures: Prevention or detection of tampering

### Availability

Failure is prevented and normal operation is maintained so that information can be used when needed

Example measures: Power supply measures, system mirroring

## 1.3 Security Threats

---

A security threat is defined as something that threatens the confidentiality, integrity, and availability indicated in "[1.2 Security Requirements](#)" in respect to information assets. This includes technical threats such as accessing a database, but does not include physical destruction.

Threats are considered to be a combination of type of user who is the source of the threat, information assets that need to be protected, techniques, and unauthorized actions. For example, a threat might be a general user exploiting a database vulnerability to obtain database management information, and then tampering with that information.

When considering security measures, it is firstly necessary to clarify what kind of threats there are. A list of possible threats is shown in the table below. Refer to "[Types of user](#)" and "[Information assets](#)" for details on the definition of each type of user and information assets that should be protected.

### Possible threats

Type of user	Information asset	Technique	Unauthorized action
General user	Database management information	Eavesdropping of packets	Unauthorized acquisition (viewing) of information  Unauthorized tampering or destruction (updating) of information
Internal user		Dictionary attack of passwords	
System manager		Unauthorized acquisition of IDs/passwords through social engineering	
System developer		Unauthorized acquisition of information through misuse of settings	
System administrator		Unauthorized acquisition of information through exploiting a database vulnerability	
System operator		Acquisition by an unauthorized route	
General user	General database information	Acquisition by a normal route	Misuse of information that can be acquired normally (taking data outside of the company)
Internal user		SQL issued with the aim of obstructing a job	Obstructing a job (resource depletion)

Type of user	Information asset	Technique	Unauthorized action
General user Internal user	General database information	Eavesdropping of packets	Unauthorized tampering or destruction (updating) of information
		Dictionary attack of passwords	
		Unauthorized acquisition of IDs/passwords through social engineering	
		Unauthorized acquisition of information through exploiting configuration errors	
		Unauthorized acquisition of information through exploiting a database vulnerability	
		Acquisition by an unauthorized route	
System manager System developer System administrator System operator	General database information	Eavesdropping of packets	Unauthorized acquisition (viewing) of information Unauthorized tampering or destruction (updating) of information
		Dictionary attack of passwords	
		Unauthorized acquisition of IDs/passwords through social engineering	
		Unauthorized acquisition of information through exploiting configuration errors	
		Unauthorized acquisition of information through exploiting a database vulnerability	
		Acquisition by an unauthorized route	
System developer	Database management information	Creation of a backdoor	Unauthorized acquisition (viewing) of information Unauthorized tampering or destruction (updating) of information
	General database information		
System manager System administrator	Database management information	Unauthorized acquisition of information by creating an unauthorized database administrator account	Unauthorized acquisition (viewing) of information Unauthorized tampering or destruction (updating) of information
	General database information		
System manager System operator	Database management information	Unauthorized acquisition of information by tampering with database-related files	Unauthorized acquisition (viewing) of information

Type of user	Information asset	Technique	Unauthorized action
	General database information	(definition file, physical file, and so on)	Unauthorized tampering or destruction (updating) of information
Database administrator	Database management information	Misuse of information (taking information outside of the company) after obtaining it through the normal route	Misuse of information that can be acquired normally (taking information outside of the company)
		Unauthorized use of IDs/ passwords from the management information	Tampering with or destroying information that can be acquired
		Unauthorized acquisition of information by tampering with management information	
		SQL issued with the aim of obstructing a job	Obstructing a job (resource depletion)
	General database information	Eavesdropping of packets	Unauthorized acquisition (viewing) of information
		Misuse of information (taking information outside of the company) after obtaining it through an unauthorized route	Unauthorized tampering or destruction (updating) of information
Database operator	Database management information	Eavesdropping of packets	Unauthorized acquisition (viewing) of information
		Dictionary attack of passwords	
		Unauthorized acquisition of IDs/passwords through social engineering	Unauthorized tampering or destruction (updating) of information
		Unauthorized acquisition of information by exploiting configuration errors	
		Unauthorized acquisition of information through exploiting a database vulnerability	
		Acquisition by an unauthorized route	
	General database information	Acquisition by a normal route	Misuse of information that can be acquired normally (taking data outside of the company)
		SQL issued with the aim of obstructing a job	Obstructing a job (resource depletion)

### Types of user

In database security, the persons involved with databases and their roles are defined below.

Type of user	Role
System manager	Manages developers, administrators, and operators
System developer	Builds the network around the database server Builds the database server
System administrator	Operates devices of the surrounding database network Operates the database server
System operator	Operates the surrounding database network
Database administrator	Builds the database system Operates the database system
Database operator	Performs business operations
Internal user	End user inside the company
General user	End user outside the company

### Information assets

In database security, it is necessary to protect the information assets to be stored on the database server.

Such assets are defined below.

#### Database management information

- Database configuration information (system catalog, user ID/password, and so on)
- Database logs (such as access logs)

#### General database information


- Job data
- Applications

## 1.4 Security Scope

---

In database systems, both the database server and the surrounding database network are part of the security scope. It is necessary to clarify the extent of the security scope that each type of user is involved with, and consider security measures for the same.

The relationship of the security scope and the types of user is shown below.



## 1.5 Security Provided by Fujitsu Enterprise Postgres

Fujitsu Enterprise Postgres provides security features that satisfy the security requirements indicated in "1.2 Security Requirements".

This section describes security provided by Fujitsu Enterprise Postgres.

### 1.5.1 Roles Targeted For Security

In Fujitsu Enterprise Postgres database systems, the roles targeted in relation to security are "Manager", "Administrator", and "User". In order to build a robust security system, it is necessary to put security measures in place for each role.

The roles targeted for security and the mapping of **Types of user** indicated in "1.3 Security Threats" are shown in the table below.

Role targeted for security	Type of user
Manager	System manager
Administrator	System developer
	System administrator
	System operator
	Database administrator
	Database operator
User	General user
	Internal user

#### Manager

The manager establishes a security policy and decides on an operations policy for the organization as a whole.

Refer to "[Chapter 3 Tasks of the Manager](#)" for details.

## Administrator

Administrators design, build and operate a system. While doing this, the administrators must implement the security measures in accordance with the security policy established by the manager.

Refer to "[Chapter 4 Tasks of Administrators](#)" for details.

## User

A user is a person other than the manager or an administrator who accesses a database. There may be any number of users. It is necessary for users to be registered in the database system, and that access to the database is restricted according to the access privileges.

Refer to "[Chapter 5 Tasks of Users](#)" for details.

## 1.5.2 Security Features

---

Fujitsu Enterprise Postgres provides the following security features:

- Authentication
- Access control
- Encryption
- Audit log
- Data masking

This section describes each of these features.

### Authentication

The databases that can be accessed can be restricted by authenticating the database users who access the database. Additionally, authentication of the server can be performed to prevent spoofing of the database server.

Refer to "Client Authentication" in "Server Administration" in the PostgreSQL Documentation for details on authentication.

Refer to "Secure TCP/IP Connections with SSL" in "Server Setup and Operation" in the PostgreSQL Documentation for details on server authentication.

In Fujitsu Enterprise Postgres, when password authentication is used as client authentication to connect to a database, the database administrator can force database users to use passwords based on predefined security policies. For more information, refer to "Policy-based Login Security" in the Operations Guide.

### Access control

Database objects can only be used by the object creator or database user who was specified as the owner when the object was created (both persons are hereinafter referred to as "owner"), or superuser, when objects are in their initial state. By having the object owner or superuser control access privileges for database users, it is possible to control what kind of tables the database users who connect to the database can access, and what kind of operations they can perform.

Fujitsu Enterprise Postgres provides security management support features that support the design and operation of access control. For details, refer to "[Chapter 8 Confidentiality Management](#)".

Refer to "Privileges" in "The SQL Language" in the PostgreSQL Documentation for details on object access control.

### Encryption

Fujitsu Enterprise Postgres provides a transparent data encryption feature that satisfies the requirements below.

- Confidential information can be changed into an unidentifiable state.
- The encryption key and data are managed separately.
- The encryption key is replaced at regular intervals.

Also, confidential data should not be operated without encryption. Fujitsu Enterprise Postgres provides security management support features to help prevent this. For details, refer to "Security Management Support".

PostgreSQL provides an encryption feature called "pgcrypto" that can also be used in Fujitsu Enterprise Postgres, however, it is recommended to use the transparent data encryption features because it will otherwise be necessary to modify the applications that consider encryption. Refer to "Protecting Storage Data Using Transparent Data Encryption" in the Operation Guide for details.

Additionally, if communication data transferred between a client and a server contains confidential information, it is necessary to encrypt the communication data to protect it against threats, such as eavesdropping on the network.

Refer to "Configuring Secure Communication Using Secure Sockets Layer" in the Operation Guide for details on encryption of communication data.

## **Audit log**

A feature that addresses threats such as misuse of administrator privileges, unauthorized access to a database by a user, and other such threats. Information for tracing the processing of administrators and users is retrieved and stored as an audit log.

By periodically viewing and monitoring audit logs, the administrators can detect events that are impacting on the system in some way, or are depleting system resources as a result of incorrect operations by users, and can take appropriate measures to prevent information leakages or system failures in advance.

Refer to "[Chapter 7 Audit Log Feature](#)" for details.

## **Data masking**

A feature that changes part of the data to make it available for reference in response to queries issued by an application.

For example, for a query of employee data, digits except the last four digits of an eight-digit employee number can be changed to "\*" so that it can be used for reference without exposing the actual data.

Specifically, the data changed by the data masking feature can be transferred to a test database so that users who perform testing or development can reference the data. During testing, it is desirable to use the data that will be used on a production environment database. However, actual production data should not be used as is for testing because of the risk of leakage of confidential data. This feature enables data that is similar to actual production data to be safely used in test and development environments.

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

# Chapter 2 Overview of Security Operation


## 2.1 Security Operation Flow

---

This section shows the flow of work when building a security environment and performing security operation in Fujitsu Enterprise Postgres.

When performing security operation, there are technical measures to be implemented to address security threats by equipping the system with security features, and manual work, such as the implementation of security guidelines, a training system, and the establishment of usage rules.

Figure 2.1 Security operation flow



## Chapter 3 Tasks of the Manager

The manager formulates security policies, which become guidelines for security measures.

### 3.1 Defining Important Information and Risk Analysis

Before formulating security policies, define important information and perform risk analysis. Based on the importance of the information and the result of risk analysis, decide what kind of security measures to put in place.

In defining the important information, identify what should be protected and classify it by importance in order to effectively implement the security measures. Information that should be protected includes "database management information" and "general database information", as indicated in "[Information assets](#)". Examples of information classifications are "personal information" and "confidential information".

In the risk analysis, refer to "[Possible threats](#)" to identify threats that may arise, and analyze the risks in respect to such threats.

Additionally, by performing a risk analysis once annually as a guide, it is possible to identify threats that may adversely impact the business and related vulnerabilities.

### 3.2 Formulating Account Management Policies

In formulating an account management policy, implement the following and document the formulated policy.

Organize system users and roles

Identify the necessary roles of the relevant system based on "[Types of user](#)". Additionally, organize personnel for each role.

Organize accounts

Organize accounts with the appropriate privileges for each role, and decide on account policies.

- Database administrator account
  - Organize separate accounts for database administrators and database operators
  - Ensure that the database administrator account can only be used by specific persons
  - Perform tasks that do not require database administrator privileges using a separate account without database administrator privileges
- General account

Create an account for general users by application usage.

Review account management policy

Review the accounts in order to effectively implement security measures.

- Regularly check the accounts mentioned above and their privileges, and determine if they are still appropriate
- If there have been system or operational changes, review the accounts and privileges
- If unsuitable accounts and privileges are discovered, modify them as required

### 3.3 Formulating Log Retrieval Policies

In formulating a log retrieval policy, implement the following and document the formulated policy.

Organize the purpose of log retrieval

To clarify what logs will be retrieved for, define their reason for retrieval.

Examples of the purpose might include, "To use for investigation in the event of unauthorized access", and "To submit to investigating authorities as evidence if any issues arise".

#### Decide on the types of logs to be retrieved

In order to retrieve appropriate logs, organize the types of logs that can be retrieved in the target system, and decide on the logs to be retrieved.

Examples of log types are "operating system logs", "application run logs", and "database audit logs".

#### Organize log retrieval target access

In order to decide on access for log retrieval targets, organize what kind of access will take place.

For example, the following access is possible:

- Access related to important information
  - Access to personal information, confidential information, and database management information
  - Access outside of business hours
  - Login
  - Specific SQL
- Access suspected to be unauthorized
  - Large amount of search access
  - Access from different locations
  - Access outside of business hours

#### Decide on the log retrieval content

In order to effectively use retrieved logs, organize the required content as a log, and decide on the retrieval content.

For example, the following output content is possible:

- When (time)
- Who (database account, application user)
- What (object ID, table name)
- Where from (machine name, IP address)
- How (SQL type, SQL statement)
- Execution result (success/fail)

#### Formulate log maintenance policy

In order to use the logs as purposed, formulate the log maintenance policy.

For each log, define its location, storage medium, retention period, access control, and so on.

## 3.4 Formulating Rules

---

Formulate the rules that will become the standard for security measures of the target system. Additionally, prescribe penalties for security violations. For example, formulate rules and penalties as below:

- Rules
  - Applying security patches and update programs
  - Prohibiting unauthorized acquisition of information from the database
  - Prohibiting the saving of acquired information to media that is not permitted for use
- Penalties
  - Prescribe penalties in the company's employment policies and procedures

- Set fines

### **3.5 Implementing Training**

---

In order to have administrators and users recognize the importance and necessity of information security, and to prevent unauthorized access due to operational omissions and mistakes, implement and promote security-related training for administrators and users.

For example, implement promotion of security policies, formulation of training schedules, and formulation of training materials.

### **3.6 Checking the Database Management Operations**

---

In order to prevent operational errors and unauthorized actions by administrators, implement the measures below:

- Always collect the latest information on security incidents and vulnerabilities related to databases
- Implement management operations only after providing advance notice
- Retain records of management operations

### **3.7 Periodic Diagnosis of the Status of Security Measures**

---

In order to check if the security measures are effective, periodically diagnose if the security measures have been put in place appropriately based on the security threats.

Additionally, evaluate if the current security measures and policies are effective for the threats and vulnerabilities, and if there are any issues, review the security policies and security measures.

# Chapter 4 Tasks of Administrators

Administrators perform the actions below as security measures when designing, building, and operating the system in accordance with the security policies formulated by the manager.

## Preparation

- Implement training

## Measures to protect against unauthorized behavior

- Perform initial setup
- Authenticate
- Control access
- Encrypt
- Control use of external media
- Restrict access
- Restrict resources

## Measures to detect and trace unauthorized behavior

- Manage logs
- Detect unauthorized access
- Analyze logs

## 4.1 Receiving Training

---

Administrators receive security-related training in accordance with the training schedule formulated by the manager. Additionally, administrators instruct users to receive training.

## 4.2 Initial Setup

---

To minimize database vulnerabilities and the possibility of unauthorized access, implement the security measures below in the initial stage of system building. Additionally, configure the database server so that it primarily operates the database system only.

### Making the server more robust

Configure the operating system and network to prevent intrusion into or destruction of a database server, so that the system operates on a secure server.

- Remove unnecessary features or services on the operating system
- Enable only the necessary protocols
- Implement the security features for services, protocols, and daemons considered to have a relatively low security level, such as file sharing and FTP

### Installing the latest version

Always download and apply the latest patches in order to reflect the latest security measures.

### Installing the minimum necessary features

Install only the necessary features in order to prevent unauthorized use of the system.

Additionally, delete or disable features and services that will not be used.

### Changing the port

To prevent unauthorized use of the system, change the default port that is set during installation.



Specify the port during setup of Fujitsu Enterprise Postgres. Refer to the Installation and Setup Guide for Server for details.

#### Access restrictions for communication features

To prevent unauthorized use of the system using the communication features, implement access restrictions for communication features.

#### Settings for prohibiting the access path to database configuration files

To prevent database destruction, implement the measures below:

- Restrict users who are permitted to access database configuration files, and periodically review the permissions
- Allow only administrators to access table or definition scripts

#### Restrictions on the access path to the database

To prevent unauthorized use or operating errors for the database, restrict the distribution range of applications used to access the database only to devices used by users who are permitted access.

#### Dealing with unauthorized programs

To prevent unauthorized intrusions into a system through a backdoor, such as by tampering with the program source code of an application, document the author of the program to be run and perform checking and testing so that the program will not be tampered with. Additionally, employ safe coding techniques so that issues with general coding vulnerabilities can be addressed.

#### System security settings

In cases where it is clear that the system security settings will impact security, set reliable security settings in the initial setup stage, such as setting appropriate security parameters.

## 4.3 Authentication

---

When accessing a database, authentication must always be performed in order to prevent tampering or information leakage from spoofing by a malicious user.

Password authentication is used when logging on to a database, and the account and password used for authentication are to be strictly managed by administrators.

Additionally, authentication must also be implemented reliably for connections to a database from clients, so that only permitted users can access the database.

### 4.3.1 Managing Accounts

---

For account management, perform the actions below.

#### Create the required accounts

To prevent unauthorized use of accounts, such as spoofing, implement the measures below when creating an account:

- Select the required account
- Specify the user privileges
- Create database administrator accounts and general user accounts separately according to the privileges



Accounts are created using the CREATE ROLE statement. Refer to "CREATE ROLE" in the PostgreSQL Documentation for details.

### Delete unnecessary accounts

Remove accounts not used on a daily basis, such as unused accounts and accounts not needed for operations that are created by default during product installation.



Accounts are deleted using the DROP ROLE statement. Refer to "DROP ROLE" in the PostgreSQL Documentation for details.

### Set up account lockout

The usage frequency of accounts is to be checked periodically, and if there are any accounts that have not been used for a long period, lock those accounts. Set a limit for failed login attempts, and if this limit is exceeded, lock the account. Additionally, set the period until a locked account is reenabled.



Account locking can be performed by using LDAP authentication. Refer to "LDAP Authentication" in the PostgreSQL Documentation for details.

### Manage database administrator accounts

Manage database administrator accounts in accordance with the account management policy formulated by the manager.

### Manage development environment and production environment accounts

To prevent unauthorized use of accounts used in a development environment, delete accounts used in the development environment before operation starts in the production environment. In cases where it is unavoidable to use an account used in the development environment in the production environment, use different passwords in each environment.

### Set up a temporary use account

If a temporary user will use the system, either provide a shared account with a temporary password for each use, or create a temporary account.

## 4.3.2 Managing Passwords

---

Manage passwords as below.

### Make strong passwords

The use of account passwords that can easily be guessed by others, such as a password that matches the ID, or the default password provided during installation, is prohibited. Set a complex and strong password.

### Change passwords regularly

Change passwords regularly to prevent others from accessing the account in case the password is obtained by unauthorized means. Additionally, configure the settings to force a password change when prompted after the first use.

### Set the password expiry period

To encourage regular changing of passwords, set a password expiry period.

In addition, by setting a password authentication policy called a profile, if a database user's password status deviates from the predefined policy, connection to the database server can be refused or the user can be forced to change their password. For information on setting password operation policies using profiles, refer to "Policy-based Login Security" in the Operation Guide.



Password setting and changing is specified using the CREATE ROLE statement or ALTER ROLE statement. Refer to "CREATE ROLE" and "ALTER ROLE" in the PostgreSQL Documentation for details.

Additionally, by using passwordcheck and LDAP authentication, the actions below can be performed:

- The default password set during installation can be changed
- The password expiry period can be set
- The number and types of characters used for the password can be checked

Refer to "passwordcheck" and "LDAP Authentication" in the PostgreSQL Documentation for details.

---

### 4.3.3 Configuring Connections and Authentication

---

Configure connections and authentication so that the database can only be accessed by permitted users.



Client authentication is configured in pg\_hba.conf. Refer to "Client Authentication" in the PostgreSQL Documentation for details.

---

## 4.4 Access Control

---

If appropriate access privileges are not set for administrators and users, security incidents may occur, such as information leakage resulting from access to information by an unauthorized person. To minimize such incidents, it is necessary to implement the security measures below for the access privileges and perform rule-based access control.



#### Notes when setting access privileges

- The creation of a special account that allows granting of privileges to all users is prohibited
- The creation of a general account that allows access to general information such as operations data is prohibited

#### Identifying the database access requirements

To set the appropriate access privileges for each usage purpose for the database, follow the procedure below to identify the access requirements:

1. Classify the usage purpose of the account, such as "For database management", "For object management", and "For data access".
2. Classify the required privileges for each usage purpose, such as "By feature" and "By object".
3. Categorize the accounts based on each privilege.
4. Identify the minimum necessary range of data and minimum necessary access content (view, update, create, delete) to be accessed for each categorized account, and decide on the database access requirements.

#### Setting the access privileges

Assign the minimum necessary privileges based on the database access requirements for each categorized account. Additionally, restrict accounts when assigning administrator privileges.

#### Reviewing access privileges

To reflect changes in access requirements in the system, periodically review the access privileges and check if there are any access privileges that are no longer needed. If any unnecessary access privileges have been set, promptly modify the access privileges.



Access privileges are set using the GRANT statement or REVOKE statement. Refer to "GRANT" and "REVOKE" in the PostgreSQL Documentation for details.

## 4.5 Encryption

---

To prevent unauthorized usage of data in the event information leakage occurs due to data theft, eavesdropping of communication, and other such activities, implement the encryption measures below.

### Encrypt communication

To protect data from eavesdropping over the network between a database server and clients, use the encryption feature to encrypt communications.

Refer to "Configuring Secure Communication Using Secure Sockets Layer" in the Operation Guide for details.

### Encrypt data

To protect data from theft, use the encryption feature to encrypt the data. The data below is targeted for encryption:

- Data to be stored on the database
- Backup data
- Data files

Refer to "Protecting Storage Data Using Transparent Data Encryption" in the Operation Guide for details.

### Manage encryption keys

Restrict the persons who can access the encryption key to a minimum number of database administrators.

Additionally, to ensure the encrypted information will not be easily decrypted, create a mechanism for appropriately managing the encryption key for the entire life cycle (generation, distribution, saving, and disposal), and strictly manage the encryption key.

Refer to "Configuring Secure Communication Using Secure Sockets Layer" and "Protecting Storage Data Using Transparent Data Encryption" in the Operation Guide for details.

## 4.6 Controlling Use of External Media

---

Information leakage can be prevented by controlling use of external media (such as CD/DVD, USB drive, and external hard disk) and PCs that are connected to the database, and restricting the removal of data from the database.

### Restricting connection of external media

Remove external media and printers that will not be used in operations, and restrict connection of external media to which information may be written.

### Restricting use of external media

Restrict connections for external media and printers to control the writing of information to these devices.

### Controlling use of connected PCs

Prevent leakage of information from PCs connected to the database:

- Limit connections of external media to PCs
- Implement security measures to make the PC robust
- Implement individual authentication for access from the PC
- Manage the installed software and monitor the software usage status
- Limit connections to printers

## 4.7 Security Measures for Servers/Applications

---

An even more robust security system can be achieved by strengthening security for servers and applications in addition to the security measures for databases. Implement the security measures below for servers and applications:

### Restrict access

Implement the measures below and restrict access to the database server:

- Install the database server inside the firewall to prevent direct access to the database server from many unspecified PCs.
- In the local network, implement measures such as using the router to restrict IP addresses, and restrict PCs and segments that can directly access the database server.

### Restrict resources

Restrict excessive use of CPU resources by general users to prevent the disruption of service and extraction of large amounts of data.

## 4.8 Log Management

---

Logs are a feature that addresses threats such as misuse of administrator privileges, and unauthorized access to a database by a user. Information for investigating/tracing processes and operations performed for the database is retrieved and managed as logs for identifying the cause in the event information leakage or unauthorized access occurs.

Fujitsu Enterprise Postgres provides the audit log feature for retrieving and managing logs. Refer to "[Chapter 7 Audit Log Feature](#)" for details.

This section describes the information that should be obtained as logs and how to maintain logs, as a measure for managing information leakage and unauthorized access.

### 4.8.1 Retrieving Logs

---

The audit logs below are retrieved in accordance with the log retrieval policy formulated by the manager.

#### Login information

Retrieves logs during login and logout.

#### Database access information (view/update)

Retrieves all access relating to the information below:

- General database information (such as personal information and confidential information used in the business)
- Database management information (system catalog, user ID/password, and so on)

#### Changed information of database objects

Retrieves logs related to creating, changing, and deleting database objects such as database accounts and tables.

#### Operation logs for audit logs

To prevent suppression of retrieved audit logs, operations such as initialization of audit logs, and stoppage of the audit log feature are retrieved as logs.

### 4.8.2 Maintaining Logs

---

Logs are maintained in accordance with the log retrieval policy formulated by the manager.

#### Storing logs

Perform the actions below and store logs securely so that the retrieved logs will not be updated by others:

- Save logs to external media, and store the external media in a secure location, such as lockable storage
- Restrict the viewing of logs to administrators only, and set access restrictions for logs, such as not assigning update rights

- Decide on the log retention period, with consideration to cases where investigation tracing back to the time of discovery of an issue is required

#### Preventing tampering of logs

Implement measures to prevent tampering of logs, such as retaining multiple copies of logs and using storage that cannot be rewritten.

#### Encrypting logs

Encrypt logs so that logs are not easily viewed.

## 4.9 Detecting Unauthorized Access

---

To address unauthorized access, it is necessary to establish a mechanism for detecting unauthorized access to databases and monitor access.

#### Communicating unauthorized access

Create a mechanism that notifies of detected unauthorized access, such as notifying the manager and the administrator, if an account lock occurs due to the limit for failed login attempts being exceeded.

#### Checking access times

Create a mechanism that can check for suspicious access to the information below outside of normal access hours, together with implementing measures to address such access.

##### Detecting access to database management information

- Monitor logs and detect access during timeframes that have not been applied for
- In the event a request for access permission outside of normal access hours is made, the log is checked for any discrepancies in the requested content and work result

##### Detecting access to general database information

- Decide on the timeframes during which access to the database is permitted for each general account
- Detect access outside of normal access hours from session information logs

#### Checking the connection source where access is not permitted

To detect access from connection sources that are not permitted, define the sources from where access is permitted, and detect access from connection sources that are not permitted.

Define the access patterns (connection source, operating system user and account) of database administrator accounts and general accounts, and check for access outside of these patterns.

## 4.10 Analyzing Logs

---

Create a mechanism that analyzes logs to detect unauthorized behavior in cases where information leakage, unauthorized access, or other such activity, is suspected. Analyses should include those shown below.

#### Periodic analysis of session information

Analyze session information of logs from the perspectives below to detect unauthorized logins:

- Trend of sessions with a large number of failed login attempts
- Trend of sessions with accounts that are logged in for long periods of time
- Trend of sessions in which a large amount of resources are used

#### Periodic analysis of database access information

Analyze SQL statements from the perspectives below to detect unauthorized access to the database:

- Trend of SQL being executed over a long period of time
- Trend of SQL using a large amount of resources

# Chapter 5 Tasks of Users

The user performs the actions below as security measures when using the system.

## 5.1 Receiving Training

---

The user must receive security-related training as instructed by the manager or the administrator to learn about security. By having users with a common awareness relating to security, an even more stable security system can be established.

## 5.2 Managing Accounts/Passwords

---

Users can use the database system by using the account and password provided by the administrator. At such times, the user is to implement the measures below so that the account and password are not misused by others:

- Be responsible for managing the ID and password in a manner that ensures the account does not become locked during login
- Change the password regularly
- Comply with the expiry period set for the password by promptly changing the password when it is about to expire

# Chapter 6 Privileged Access Management

In PostgreSQL, the database superuser (hereinafter, superuser) bypasses all permission checks except during login. Because superusers have powerful privileges, it is recommended that even database administrators typically access the database as a non-superuser. Superuser privileges are powerful and necessary for system administration, but they also carry significant security risks. Therefore, security administrators and auditors, who are responsible for supervision and auditing, must always be able to track and verify when and how superuser privileges were exercised.

Privileged access management is a framework that meets these requirements and consists of the following three elements.

- Control of privilege acquisition and exercise  
Control who can acquire and exercise superuser privileges, and under what methods and conditions.
- Recording of privilege exercise  
Record operations performed by superusers and protect them from tampering or loss.
- Auditing of privilege exercise records  
Audit privilege exercise records to detect problems retrospectively.

Fujitsu Enterprise Postgres provides the following features to support privileged access management:

- Audit log feature  
By using the audit log function, operations performed by superusers can be recorded. For details, refer to "[Chapter 7 Audit Log Feature](#)".
- set\_user  
set\_user is an extension of PostgreSQL. By using the features provided by the set\_user, you can control the elevation to superuser within the database and post-elevation operations. For details, refer to "[6.1 Controlling Elevation via set\\_user](#)".
- Integration with privileged access management products  
By integrating with privileged access management products, you can control logins and record operations on the database and database server using superuser accounts.

## 6.1 Controlling Elevation via set\_user

The set\_user extension provides the set\_user\_u function to elevate to superuser. When you use this to elevate to superuser, a log of the explicit elevation operation is output, and some dangerous operations after elevation are prohibited and blocked if you try to execute them. The database roles that can elevate are specified by granting execution privileges to the function and an explicit allow list. This makes it easier to understand who can be elevated and to track who elevated and when.

To take advantage of set\_user's benefits, it is necessary to restrict the acquisition of superuser privileges and elevation operations by other means. Specifically, change operations as follows:

- Disable direct login with a database role that has superuser privileges
- Administrators log in as non-superusers and elevate using the set\_user\_u function only when superuser privileges are required for an operation
- Set role attributes to prevent elevation to superuser except via the set\_user\_u function, specifically by blocking SET ROLE and privilege inheritance

For information on how to use set\_user, refer to the set\_user documentation. The URL for the documentation can be found in "OSS Supported by Fujitsu Enterprise Postgres" in the General Description.

### 6.1.1 Setup for set\_user

For set\_user setup, refer to "set\_user" in the Installation and Setup Guide for Server.

## 6.1.2 Constraints

---

### Processing where elevation operations cannot be inserted

Controlling superuser elevation by `set_user` requires calling the `set_user_u` function to switch users after logging in for operations that require superuser privileges. Therefore, it cannot be applied to programs that require superuser privileges, such as PostgreSQL utility programs, where elevation operations using the `set_user_u` function cannot be inserted.

If such processing is required, create a dedicated database role for each process, and superuser privileges are granted to that database role. In addition, configure client authentication so that each process can log in only through the path it uses for database connection.

### Combination with the audit log feature

When combining `set_user` with the audit log feature, it is necessary to align the output destination for logs of post-elevation operations with the audit log feature design. For details, refer to "[Chapter 7 Audit Log Feature](#)".

### Combination with the confidentiality management feature

SQL functions provided by the confidentiality management feature cannot be executed when the user is elevated or switched using the `set_user_u` function or `set_user` function. For details, refer to "[8.2.2 Determining Confidentiality Management Roles](#)".

### Combination with connection pooling feature

Connection pooling is a function that allows multiple clients to share established database connections. When using a connection pooling, it is difficult for the server to identify which end-user performed administrative access. Also, there is a risk that superuser privileges obtained via `set_user` may not be completely isolated from end users who should not have such privileges. Therefore, when performing administrative access using `set_user`, do not use the connection pooling function and connect directly to the database.

# Chapter 7 Audit Log Feature

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

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

The scalable audit log feature enables multiple output mechanisms to capture audit logs without performance degradation, even for systems with high application multiplicity and high output volumes.

If you want to take advantage of the scalable audit log feature, refer to ["7.2 Setup"](#) and ["7.3 Setting Up the Scalable Audit Log Feature"](#).



## Information

The audit log feature cannot be used if PostgreSQL is running in single-user mode.

## 7.1 Audit Log Output Modes

In pgaudit, the two types of audit log below can be output.

### Session Audit Logging

Session Audit Logging outputs information related to SQL executed in backend processes (processes generated when connection requests are received from clients), information related to starting and connecting databases, and information related to errors, as a log. In Session Audit Logging, by specifying the log output conditions and filtering the logs to be output, performance degradation due to outputting large volumes of logs can be prevented.

Refer to ["7.5 Session Audit Logging"](#) for details.

### Object Audit Logging

When SELECT, INSERT, UPDATE, and DELETE are executed for specific objects (tables, columns), Object Audit Logging outputs these as a log. TRUNCATE is not supported. Object Audit Logging outputs object operations for which privileges have been assigned to specified roles, as a log. Object Audit Logging can control log output at an even finer level of granularity than Session Audit Logging.

Refer to ["7.6 Object Audit Logging"](#) for details.



## Information

Depending on the application or command, Fujitsu Enterprise Postgres may execute SQL internally and the audit logs may be retrieved.

Also, the audit logs of multiple SQLs with the same statement ID may be retrieved. This is because before the user executes the SQL, another SQL is executed internally by Fujitsu Enterprise Postgres.

## 7.2 Setup

This section describes the setup method of pgaudit.

1. Copy the pgaudit files



As superuser, run the following command. Note that "<x>" in paths indicates the product version.

```
$ su -
Password:*****
# cp -r /opt/fsepv<x>server64/OSS/pgaudit/* /opt/fsepv<x>server64
```

W

Open a command prompt as administrator privileges, run the following command. Note that "<x>" in paths indicates the product version.

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

## 2. Create the pgaudit configuration file

Create the pgaudit configuration file, which describes the information required for pgaudit actions. Create the file using the same encoding as used for the database.

In addition, set write permissions for the database administrator only in the pgaudit configuration file so that policies related to the audit log are not viewed by unintended users.

Refer to "[7.4 pgaudit Configuration File](#)" for details.



### Information

Do not define the rule section in the pgaudit configuration file at this point.

#### Example of a pgaudit configuration file

```
[output]
logger = 'auditlog'
```

## 3. Configure postgresql.conf

Configure the parameters below in postgresql.conf to use audit logs:

### shared\_preload\_libraries

Specify "pgaudit".

When using this feature in combination with set\_user, refer to "[7.9 When Used with set\\_user](#)".

### pgaudit.config\_file

Specify the deployment destination path of the pgaudit configuration file.

If a relative path is specified, the path will be relative to the data storage directory.

### log\_replication\_commands

Specify "on".

### log\_min\_messages

Check if "ERROR" or higher has been specified.

If outputting an audit log to a server log ("serverlog" is specified in the logger parameter of the pgaudit configuration file), check the parameters below relating to server logs.

### logging\_collector

Check if "on" has been specified.

### log\_destination

Check if "stderr" has been specified.

L

### log\_file\_mode

Check if the server log permissions are appropriate, so that only the permitted persons can access it.

## Point

The default for the `log_file_mode` parameter is 0600, which only allows the database administrator to have access.

For example, to permit other members of the group to which the database administrator belongs to view the audit logs, specify 0640 for `log_file_mode`.

### Example

```
log_file_mode = 0640
```

The database administrator can also be prevented from viewing audit logs by specifying 0000. However, write privileges are assigned for outputting logs.

If outputting an audit log to a dedicated log file ("`auditlog`" is specified in the `logger` parameter of the `pgaudit` configuration file), check the parameter below.

### `max_worker_processes`

If the `max_worker_processes` parameter has been set, add 1 to the specified value.

If you want to take advantage of the scalable audit log feature, refer to "[7.3 Setting Up the Scalable Audit Log Feature](#)".

## See

Refer to "Error Reporting and Logging" in the PostgreSQL Documentation for details on server logs.

If using database multiplexing, refer to "[7.7 Database Multiplexing](#)" for details.

### Example of `postgresql.conf`

In the example below, only the parameters that need to be configured when using the audit log feature are described.

```
shared_preload_libraries = 'pgaudit'  
pgaudit.config_file = 'pgaudit.conf'  
log_replication_commands = on  
log_min_messages = WARNING
```

#### 4. Start the instance

Start the instance and check if the message below is output.

```
LOG: pgaudit extension initialized
```

#### 5. Create the `pgaudit` extension

Execute `CREATE EXTENSION` to create the `pgaudit` extension.

```
$ psql  
=# CREATE EXTENSION pgaudit;  
=# \dx  
  
List of installed extensions  
Name      | Version | Schema  | Description  
-----+-----+-----+-----  
pgaudit   | 1.0     | public  | provides auditing functionality  
plpgsql   | 1.0     | pg_catalog | PL/pgSQL procedural language  
(2 rows)
```

6. Configure the parameters in the pgaudit configuration file

Add or change the parameters in the pgaudit configuration file as required.

Refer to "[7.4 pgaudit Configuration File](#)" for details.

7. Restart the instance

Restart the instance to apply the changes to the pgaudit configuration file. After restarting, check if the changes have been reflected correctly.

**L**

### Linux

```
LOG: log_catalog = 1
LOG: log_level_string =
LOG: log_level = 15
LOG: log_parameter = 0
LOG: log_statement_once = 0
LOG: role =
LOG: logger = auditlog
LOG: log_directory = pgaudit_log
LOG: log_filename = pgaudit-%Y-%m-%d_%H%M%S.log
LOG: log_file_mode = 0600
LOG: log_rotation_age = 1440
LOG: log_rotation_size = 10240
LOG: log_truncate_on_rotation = 0
LOG: fifo_directory = /tmp
LOG: Rule 0
LOG: pgaudit extension initialized
```

**W**

### Windows


```
LOG: log_catalog = 1
LOG: log_level_string =
LOG: log_level = 15
LOG: log_parameter = 0
LOG: log_statement_once = 0
LOG: role =
LOG: logger = auditlog
LOG: log_directory = pgaudit_log
LOG: log_filename = pgaudit-%Y-%m-%d_%H%M%S.log
LOG: log_file_mode = 0600
LOG: log_rotation_age = 1440
LOG: log_rotation_size = 10240
LOG: log_truncate_on_rotation = 0
LOG: Rule 0
LOG: pgaudit extension initialized
```

## 7.3 Setting Up the Scalable Audit Log Feature

With normal audit log feature, there is only one logger process and only one dedicated log file.

Therefore, the load of audit logs, including logs output by all backend processes, in the instance concentrates to the logger process and dedicated log file.

The scalable audit log feature distributes multiple dedicated log files, as shown below, and executes multiple output processes (logger processes) in a one-to-one relationship with those files. The number of distributions can be changed by parameter settings.



### Information

Because there is a one-to-one relationship between the dedicated log file and the logger process, the logger process also operates as many dedicated log files as there are. A number is appended to the end of its command title as reported by `ps` or Process Explorer, as follows. Use this information as a reference when monitoring the process status.

#### Example

```
pgaudit logger 0
pgaudit logger 1
pgaudit logger 2
```

### Parameters in the pgaudit configuration file

Set the following parameters to the "output section":

Refer to the ["7.4 pgaudit Configuration File"](#) for details on the parameters.

#### enable\_parallel\_logger

Set "on". The default is "off".

If the logger parameter in the pgaudit configuration file is not "auditlog" (the default), it is ignored even if set to "on"

#### parallel\_loggers

Sets the number of logger processes. Specify a value greater than or equal to "2".

Refer to ["7.4 pgaudit Configuration File"](#) for an estimate of the values to set.

#### log\_rotation\_age

We strongly recommend that you set this parameter. Although the scalable audit log feature works without configuration, it is easier to avoid incorrect analysis when analyzing the output audit logs. Refer to ["Considerations when Using the Scalable Audit Log Feature"](#).

#### Example

```
[output]
enable_parallel_logger = on
parallel_loggers = 5
log_rotation_age = 1h
```

### postgresql.conf Parameters

#### max\_worker\_processes

Adds the value set for `parallel_loggers` to the value set for the `max_worker_processes` parameter. This is because the logger process acts as a background worker.

#### Example

If `max_worker_processes` is set to 8 and you want to set `parallel_loggers` to 3, do the following:

```
max_worker_processes = 11
```

## Storage location of the dedicated log file

A dedicated log file is output with the following path and name:

For more information about the `log_directory` and `log_filename` parameters, refer to "[7.4 pgaudit Configuration File](#)".

```
log_directory parameter setting/number/number-log_filename parameter setting
```

The number is a number between 0 and the number specified by `parallel_loggers`. If `parallel_loggers` is 3, it can be 0 or greater and 2 or less. For more information about how to analyze these multiple files, refer to "[7.10 Analyzing Audit Logs in SQL](#)".

### Example

If you set `parallel_loggers=3`, a dedicated log file is generated as follows:

```
pgaudit_log/0/0-pgaudit-2024-02-02_153000.log  
pgaudit_log/1/1-pgaudit-2024-02-02_153000.log  
pgaudit_log/2/2-pgaudit-2024-02-02_153000.log
```

If you can use multiple disks, you can also distribute the disk I/O load by setting the dedicated log file storage directory as a symbolic link to another disk, as shown below.

For Windows, use junction instead of symbolic links.

### Example

Place `pgaudit_log/2` on a different disk:

```
ln -s /other_disk/2 pgaudit_log/2
```

```
mklink /J pgaudit_log\2 E:\other_disk\2
```

Depending on the situation, create a symbolic link as follows. If the directory (or symbolic link) specified as the log output destination for the scalable audit log feature does not exist, the directory is automatically created. If the directory already exists, it is used as the log output destination.

- If you are using the scalable audit log feature for the first time, there is no dedicated log file directory. In this case, create a symbolic link after stopping the instance.
- If you are already using the scalable audit log feature, you already have a dedicated log file directory. After stopping the instance, move the stored dedicated log file to another location, and then delete the storage directory. Then, create a symbolic link with the same name as the deleted directory.

## 7.4 pgaudit Configuration File

In the `pgaudit` configuration file, specify the information required for `pgaudit` actions. The `pgaudit` configuration file comprises three sections: "output section", "option section", and "rule section".

### output section

The output section is specified using the format below:

- `paramName = 'value'`

The valid parameters in the output section are shown in the table below.

```
Linux
```

Parameter name	Description	Remarks
logger	Dedicated log file ( <code>auditlog</code> )/ <code>serverLog</code> ( <code>serverlog</code> ) that will be the output destination of the audit log  The default is "auditlog" (dedicated log file).	The dedicated log file is output using the same encoding as used for the database.
log_directory	Directory where the audit log is to be created  Specify the full path or the relative path from the data storage directory.  The default is "pgaudit_log".  However, make sure that audit log files are not output under the data storage directory.  This is because if you are a backup target, such as <code>pg_basebackup</code> , and recover using that backup data, the audit log from the current time to the backup time will disappear.  This is fine if the audit log file is located at the end of the symbolic link. This is because backups such as <code>pg_basebackup</code> do not track such symbolic links.  If you specify "on" for the <code>enable_parallel_logger</code> parameter, an audit log is created for "log_directory parameter setting/number". Refer to <a href="#">"7.3 Setting Up the Scalable Audit Log Feature"</a> for more information.	Enabled only if "auditlog" is specified for the logger parameter
log_filename	File name of the audit log  Specify a file name that varies according to the time, in the same manner as for <code>log_filename</code> in the <code>postgresql.conf</code> file.  The default is "pgaudit-%Y-%m-%d_%H%M%S.log".  If you specify "on" for the <code>enable_parallel_logger</code> parameter, the audit log filename is "number-log_filename parameter setting". Refer to <a href="#">"7.3 Setting Up the Scalable Audit Log Feature"</a> for more information.	Enabled only if "auditlog" is specified for the logger parameter
log_file_mode	Specify the permissions of the audit log so that only permitted persons can access it.  The parameter value is the numeric mode specified in the format permitted in <code>chmod</code> and <code>umask</code> system calls. The default is "0600".  Refer to " <a href="#">log_file_mode</a> " in <a href="#">"7.2 Setup"</a> for information on audit log file permissions.	Enabled only if "auditlog" is specified for the logger parameter
log_rotation_age	Maximum age of the audit log file  A new audit log file is generated when the time (minute units) specified here elapses. To disable generation of new log files based on time, specify "0".  The valid units are "min" (minutes), "h" (hours), and "d" (days). If the unit is omitted, "min" will be used.  The default is "1d" (1 day).  We strongly recommend that you set this parameter if you want to take advantage of the scalable audit log feature. This is because all logger processes rotate their dedicated log files at the same time, making it easier to	Enabled only if "auditlog" is specified for the logger parameter

Parameter name	Description	Remarks
	avoid incorrect analysis. For more information, refer to <a href="#">"Considerations when Using the Scalable Audit Log Feature"</a> .	
log_rotation_size	<p>Maximum size of the audit log file</p> <p>A new log file will be generated after logs of the size specified here are output to a log file. To disable generation of new log files based on size, specify "0".</p> <p>The valid units are "kB" (kilobytes), "MB" (megabytes), and "GB" (gigabytes). If the unit is omitted, "kB" will be used.</p> <p>The default is "10MB".</p>	Enabled only if "auditlog" is specified for the logger parameter
log_truncate_on_rotation	<p>If rotating audit log files based on time, this parameter is used to specify whether to overwrite (on)/not overwrite (off) existing audit log files of the same name. For example, if "on" is specified, and "pgaudit-%H.log" is specified for log_filename, 24 separate log files will be generated based on time, and those files will be cyclically overwritten.</p> <p>The default is "off". If "off" is specified, the logs will be written to the existing audit log files.</p>	Enabled only if "auditlog" is specified for the logger parameter
fifo_directory	<p>FIFO (named pipe) directory to be used between the daemon process that outputs audit log files and the backend process</p> <p>FIFO named p.PGAUDIT.nnnn (nnnn is the postmaster PID) are created in the fifo_directories directory. The files cannot be deleted manually.</p> <p>The default is "/tmp".</p> <p>If you specify "on" for the enable_parallel_logger parameter, the name of the FIFO is "p.PGAUDIT.nnnn.number". Refer to <a href="#">"7.3 Setting Up the Scalable Audit Log Feature"</a> for more information about the numbers.</p>	Enabled only if "auditlog" is specified for the logger parameter
enable_parallel_logger	<p>Specify whether to distribute the audit log output load (on)/not (off). The default is "off".</p> <p>The setting of this parameter is ignored when you specify "serverlog" for the logger parameter, because the serverlog of the database server cannot be distributed.</p>	Enabled only if "auditlog" is specified for the logger parameter
parallel_loggers	<p>Specifies the multiplicity at which the output load of the audit log is to be distributed. You can specify a number greater than or equal to 1, but to distribute the audit log output load, specify a number greater than or equal to 2.</p> <p>The default is "2".</p> <p>Estimate parallel_loggers as follows:</p> $\text{parallel\_loggers} = \text{number of cores} / 2$ <p>This estimate assumes the highest load to output audit logs, so setting a value less than the estimated value will not significantly degrade performance in most cases. The disadvantage of increasing this value is that it consumes a little more memory. For more information about estimating memory, refer to "FUJITSU Enterprise</p>	Enabled only if "auditlog" is specified for the logger parameter

Parameter name	Description	Remarks
	<p>Postgres Memory Requirements" in the Installation and Setup Guide for Server.</p> <p>Because the logger process acts as a background worker, add the value set for this parameter to the parameter <code>max_worker_processes</code> in <code>postgresql.conf</code> that sets the maximum number of background workers.</p> <p>[Note]</p> <p>If the <code>max_worker_processes</code> setting is insufficient, the instance cannot be started.</p>	

**W**

Windows

Parameter name	Description	Remarks
logger	<p>Dedicated log file (<code>auditlog</code>)/<i>serverLog</i> (<code>serverlog</code>) that will be the output destination of the audit log</p> <p>The default is "auditlog" (dedicated log file).</p>	The dedicated log file is output using the same encoding as used for the database.
log_directory	<p>Directory where the audit log is to be created</p> <p>Specify the full path or the relative path from the data storage directory.</p> <p>The default is "pgaudit_log".</p> <p>However, make sure that audit log files are not output under the data storage directory.</p> <p>This is because if you are a backup target, such as <code>pg_basebackup</code>, and recover using that backup data, the audit log from the current time to the backup time will disappear.</p> <p>This is fine if the audit log file is located at the end of the junction. This is because backups such as <code>pg_basebackup</code> do not track such junction.</p> <p>If you specify "on" for the <code>enable_parallel_logger</code> parameter, an audit log is created for "log_directory parameter setting/number". Refer to "<a href="#">7.3 Setting Up the Scalable Audit Log Feature</a>" for more information.</p>	Enabled only if "auditlog" is specified for the logger parameter
log_filename	<p>File name of the audit log</p> <p>Specify a file name that varies according to the time, in the same manner as for <code>log_filename</code> in the <code>postgresql.conf</code> file.</p> <p>The default is "pgaudit-%Y-%m-%d_%H%M%S.log".</p> <p>If you specify "on" for the <code>enable_parallel_logger</code> parameter, the audit log filename is "number-log_filename parameter setting". Refer to "<a href="#">7.3 Setting Up the Scalable Audit Log Feature</a>" for more information.</p>	Enabled only if "auditlog" is specified for the logger parameter
log_file_mode	This parameter is ignored in Windows.	
log_rotation_age	<p>Maximum age of the audit log file</p> <p>A new audit log file is generated when the time (minute units) specified here elapses. To disable generation of new log files based on time, specify "0".</p> <p>The valid units are "min" (minutes), "h" (hours), and "d" (days). If the unit is omitted, "min" will be used.</p>	Enabled only if "auditlog" is specified for the logger parameter

Parameter name	Description	Remarks
	<p>The default is "1d" (1 day).</p> <p>We strongly recommend that you set this parameter if you want to take advantage of the scalable audit log feature. This is because all logger processes rotate their dedicated log files at the same time, making it easier to avoid incorrect analysis. For more information, refer to "<a href="#">Considerations when Using the Scalable Audit Log Feature</a>".</p>	
log_rotation_size	<p>Maximum size of the audit log file</p> <p>A new log file will be generated after logs of the size specified here are output to a log file. To disable generation of new log files based on size, specify "0".</p> <p>The valid units are "kB" (kilobytes), "MB" (megabytes), and "GB" (gigabytes). If the unit is omitted, "kB" will be used.</p> <p>The default is "10MB".</p>	Enabled only if "auditlog" is specified for the logger parameter
log_truncate_on_rotation	<p>If rotating audit log files based on time, this parameter is used to specify whether to overwrite (on)/not overwrite (off) existing audit log files of the same name. For example, if "on" is specified, and "pgaudit-%H.log" is specified for log_filename, 24 separate log files will be generated based on time, and those files will be cyclically overwritten.</p> <p>The default is "off". If "off" is specified, the logs will be written to the existing audit log files.</p>	Enabled only if "auditlog" is specified for the logger parameter
enable_parallel_logger	<p>Specify whether to distribute the audit log output load (on)/not (off). The default is "off".</p> <p>The setting of this parameter is ignored when you specify "serverlog" for the logger parameter, because the serverlog of the database server cannot be distributed.</p>	Enabled only if "auditlog" is specified for the logger parameter
parallel_loggers	<p>Specifies the multiplicity at which the output load of the audit log is to be distributed. You can specify a number greater than or equal to 1, but to distribute the audit log output load, specify a number greater than or equal to 2.</p> <p>The default is "2".</p> <p>Estimate parallel_loggers as follows:</p> $\text{parallel\_loggers} = \text{number of cores} / 2$ <p>This estimate assumes the highest load to output audit logs, so setting a value less than the estimated value will not significantly degrade performance in most cases. The disadvantage of increasing this value is that it consumes a little more memory. For more information about estimating memory, refer to "FUJITSU Enterprise Postgres Memory Requirements" in the Installation and Setup Guide for Server.</p> <p>Because the logger process acts as a background worker, add the value set for this parameter to the parameter max_worker_processes in postgresql.conf that sets the maximum number of background workers.</p>	Enabled only if "auditlog" is specified for the logger parameter

Parameter name	Description	Remarks
	[Note] If the max_worker_processes setting is insufficient, the instance cannot be started.	

## Information

If the logger parameter is set to "serverlog", audit logs will be output to the server log as log messages, therefore the status information and message severity level according to the log\_line\_prefix parameter in postgresql.conf will be output to the beginning of the audit log.

If the logger parameter is omitted or set to "auditlog", audit logs will be output to a dedicated log file as dedicated logs, therefore the status information and message severity level according to the log\_line\_prefix parameter in the postgresql.conf file will not be output.

Refer to "Output format" in "7.5 Session Audit Logging" or "Output format" in "7.6 Object Audit Logging" for information on the output format of audit logs.

## L

### Point

The pgaudit log\_file\_mode configuration parameter setting is separate from, and unaffected by, the log\_file\_mode GUC parameter setting and the -g/-allow-group-access initdb option.

When using a dedicated pgaudit log file, since the pgaudit log\_directory location defaults to inside the data storage directory, it is possible for the pgaudit log\_file\_mode permissions to conflict with the intended file permissions specified by the -g/-allow-group-access initdb option. In this case, the pgaudit log\_directory should be specified to be a directory located outside of the data storage directory.

## W

### Point

If the file permissions for the server log do not satisfy the security requirements of the system, specify a dedicated log file for the output destination of the audit log, and set the access permissions for the output destination directory so that the security requirements are satisfied. At that time, assign read and write permissions to the database administrator.

## option section

The option section is specified using the format below:

- *paramName* = '*value*'

The valid parameters in the option section are shown in the table below.

Parameter name	Description	Remarks
role	Name of roles used in Object Audit Logging  If specifying a name containing uppercase characters, key words, multibyte characters and commas, enclose the name in double quotation marks.	Parameter used in Object Audit Logging only
log_catalog	Whether to enable (on)/disable (off) log output for pg_catalog	

Parameter name	Description	Remarks
	If pgAdmin, and so on, will not be retrieving audit logs that access pg_catalog, specify "off". The default is "on" (enabled).	
log_parameter	Whether to enable (on)/disable (off) output of values passed by parameters in SQL execution The default is "off" (disabled).	
log_statement_once	Whether to control (on)/not control (off) output for the second and subsequent SQL statements if the same SQL statement is the log output target The default is "off" (do not control).	
log_level	Log level of audit logs The valid values are "DEBUG5", "DEBUG4", "DEBUG3", "DEBUG2", "DEBUG1", "INFO", "NOTICE", "WARNING", and "LOG". The default is "LOG".	Enabled only if "serverlog" is specified for the logger parameter
audit_log_disconnections	When using Mirroring Controller, specify whether to enable (on) or disable (off) the output of disconnection logs for connections other than those of Mirroring Controller. The default is "off" (disabled). This parameter is valid when log_disconnections in postgresql.conf is off.	Parameter used in Session Audit Logging only
audit_set_user_switch	When combined with set_user, specify whether operations after user switching by set_user are subject to pgaudit auditing (on) or not (off). The default is "on".	Enabled only if "serverlog" is specified for the logger parameter. If "auditlog" is specified for the logger parameter, operations performed after a user switch via set_user are audited regardless of this setting.

## rule section

The rule section is used in Session Audit Logging. Refer to "[7.5 Session Audit Logging](#)" for details.

### Information

Do not specify the rule section if the role parameter has been specified in the option section. If you specify the rule section, the audit logs of Object Audit Logging and Session Audit Logging will be output intermingled. The CSV format of the role audit log and the CSV format of the rule audit log are different, you will be unable to view in CSV format.

## 7.5 Session Audit Logging

In Session Audit Logging, specify the rules for filtering logs to be output in the rule section in the pgaudit configuration file.

Rules are specified using the formats below. Multiple values can be specified, using a comma as the delimiter.

- *paramName* = 'value'

- *paramName* != 'value'

If [rule] is described on its own in the rule section with no parameters specified, all audit logs of Session Audit Logging will be output.

Example

```
[output]
logger = 'auditlog'
[rule]
```

If [rule] is not described as a section, audit logs of Session Audit Logging will not be output.

Example

```
[output]
logger = 'auditlog'
```

The valid parameters in the rule section are shown in the table below.

Parameter name	Description	Example
timestamp	Timestamp range Refer to " <a href="#">timestamp</a> " for details on how to specify timestamps.	timestamp = '09:00:00 - 10:00:00, 18:00:00 - 18:30:00'
database	Database name To specify a blank value, specify use a pair of double quotation marks ("" ) instead of the <i>value</i> . When specifying a name containing uppercase characters, key words, multibyte characters and commas, enclose the name in double quotation marks.	database = 'prodcut_db'
audit_role	Role name To specify a blank value, specify use a pair of double quotation marks ("" ) instead of the <i>value</i> . When specifying a name containing uppercase characters, key words, multibyte characters and commas, enclose the name in double quotation marks.	audit_role = 'appuser1'
class	Operation class Select from the values below. Multiple values can be specified. Refer to " <a href="#">class</a> " for details on the meaning of each class.  - BACKUP - CONNECT - DDL - ERROR - FUNCTION - MISC - READ	class = 'READ, WRITE'

Parameter name	Description	Example
	<ul style="list-style-type: none"> <li>- ROLE</li> <li>- WRITE</li> <li>- SYSTEM</li> </ul>	
object_type	<p>Object type</p> <p>This parameter is enabled when the class parameter is "READ" and "WRITE".</p> <p>Select from the values below. Multiple values can be specified.</p> <ul style="list-style-type: none"> <li>- TABLE</li> <li>- INDEX</li> <li>- SEQUENCE</li> <li>- TOAST_VALUE</li> <li>- VIEW</li> <li>- MATERIALIZED_VIEW</li> <li>- COMPOSITE_TYPE</li> <li>- FOREIGN_TABLE</li> <li>- FUNCTION</li> </ul>	object_type = 'TABLE, INDEX'
object_name	<p>Object name</p> <p>This parameter is enabled when the class parameter is "READ" and "WRITE".</p> <p>For objects that can be modified by a schema, such as a table, modify such objects by schema name.</p> <p>When specifying a name containing uppercase characters, key words, multibyte characters and commas, enclose the name in double quotation marks. When specifying the object name "<i>schemaName.tableName</i>", enclose the entire object name modified by schema name in double quotation marks.</p> <p>To specify a blank value, use a pair of double quotation marks ("" ) instead of the <i>value</i>.</p>	<p>object_name = 'myschema.tbl1'</p> <p>object_name = 'myschema.tbl1, "mySchema.TABLE"'</p>
application_name	<p>Application name</p> <p>To specify a blank value, use a pair of double quotation marks ("" ) instead of the <i>value</i>.</p>	application_name = 'myapp'
remote_host	<p>Connection source(client side) host name or IP address</p> <p>If "on" is specified for the log_hostname parameter in the postgresql.conf file, specify a host name. Otherwise, specify the IP address. If using a local host, specify "[local]".</p>	remote_host = 'ap_server'

Parameter name	Description	Example
	To specify a blank value, use a pair of double quotation marks ("" ) instead of the <i>value</i> .	

## timestamp

Specify a timestamp range from "*startTime*" to "*endTime*" for the log output target. The timestamp format is 'hh:mm:dd-hh:mm:dd' (hh is expressed in 24-hour notation, and hh, mm, and dd are expressed in two-digit notation).

The start time must be earlier than the end time. If specifying multiple ranges, specify each start and end timestamp using a comma as the delimiter.

End timestamps consider milliseconds. For example, if '11:00:00 - 11:59:59' is specified for the timestamp, "11:00:00:000" to "11:59:59:999" will be the target range.

The timestamps used by evaluation in the rule section of pgaudit are different to the timestamps issued in the log entries. That is because log entries are output after evaluation by pgaudit, with the timestamp being generated at that time.

## class

The meaning of each class specified in the class parameter is below:

- READ: SELECT, COPY FROM
- WRITE: INSERT, UPDATE, DELETE, TRUNCATE, COPY TO
- FUNCTION: Function call, DO
- ROLE: GRANT, REVOKE, CREATE ROLE, ALTER ROLE, DROP ROLE
- DDL: All DDLs (such as CREATE and ALTER) other than the DDLs of the ROLE class
- CONNECT: Events relating to connecting (request, authenticate, and disconnect)
- SYSTEM: Instance start, promotion to primary server
- BACKUP: pg\_basebackup
- ERROR: Event completed by an error (PostgreSQL error codes other than 00). This class can be used if ERROR or lower level is specified for the log\_min\_messages parameter in postgresql.conf.
- MISC: Other commands (such as DISCARD, FETCH, CHECKPOINT, and VACUUM)

## Evaluation of the rule section

- When a log event occurs, all expressions in the rule section are evaluated at once. Log entries are only output if all parameters in the rule section are evaluated as being true.

For example, if the rule below has been set, of the operations performed by 'apserver' to 'myschema.tbl1', the operations applicable to classes other than "WRITE" in the period from 10 a.m. to 11 a.m. will be output as audit logs.

```
[rule]
timestamp = '10:00:00-11:00:00'
remote_host = 'apserver'
object_name = 'myschema.tbl1'
class != 'WRITE'
```

- Multiple rule sections can be defined in the pgaudit configuration file. Log events are evaluated using each rule section, and an audit log is output for each matching rule section.

For example, if the rules below are set, duplicated audit logs will be output.

```
[rule]
object_name = 'myschema.tbl1'
```

```
[rule]
object_name = 'myschema.tbl1'
```

- If the same parameter is specified multiple times in one rule section, the last specified parameter is effective.

For example, if the rule below has been set, "object\_name = 'myschema.tbl3'" will take effect.

```
[rule]
object_name = 'myschema.tbl1'
object_name = 'myschema.tbl2'
object_name = 'myschema.tbl3'
```

## Output format

In Session Audit Logging, audit logs are output in the format below:

```
AUDIT: SESSION,READ,2022-03-12 19:00:58 PDT,
(1)      (2)      (3)
[local],19944,psql,appuser,postgres,2/8, 2, 1,SELECT,,TABLE,myschema.account, ,
(4)     (5)     (6)     (7)      (8)     (9)(10)(11)(12)(13)(14)      (15)      (16)
SELECT * FROM myschema.account;,<not logged>
(17)                                (18)
```

No	Content
(1)	Log header Fixed as "AUDIT: SESSION".
(2)	Class
(3)	SQL start time
(4)	Remote host name If using a local host, [local] is output.
(5)	Backend process ID
(6)	Application name If an application name is not specified, [unknown] is output.
(7)	User name
(8)	Database name
(9)	Virtual transaction ID
(10)	Statement ID
(11)	Substatement ID
(12)	Command tag
(13)	SQLSTATE
(14)	Object type
(15)	Object name
(16)	Error message
(17)	SQL If the SQL contains a password, such as for CREATE ROLE, and so on, it will be replaced with "<REDACTED>". Additionally, if "on" is specified for the log_statement_once parameter of the option section in the pgaudit configuration file, "<previously logged>" is output for the second and subsequent statements.

No	Content
(18)	<p>Depending on the log_parameter parameter value of the option section in the pgaudit configuration file, the output content will be as below.</p> <ul style="list-style-type: none"> <li>- log_parameter=on is specified <p>If parameters are specified in the SQL, the parameters are concatenated and output, using a space as the delimiter. If parameters are not specified in the SQL, "&lt;none&gt;" is output.</p> </li> <li>- log_parameter=off (default) is specified <p>"&lt;not logged&gt;" is output.</p> </li> </ul> <p>Additionally, if "on" is specified for the log_statement_once parameter of the option section in the pgaudit configuration file, "&lt;previously logged&gt;" is output for the second and subsequent statements.</p>

### Information

If accessing resources that use the features below, the command tag (12) may be output as "???".

- INSTEAD OF trigger
- RULE
- VIEW
- Security policy per row
- Table inheritance

### Example

Below is an example of retrieving audit logs in Session Audit Logging.

#### 1. Settings

In the pgaudit configuration file, specify the rule section below.

```
[rule]
class = 'READ, WRITE'
object_name = 'myschema.account'
```

#### 2. Retrieving logs

Execute the SQL below from the client.

```
CREATE TABLE myschema.account
(
  id int,
  name text,
  password text,
  description text
);
INSERT INTO myschema.account (id, name, password, description) VALUES (1, 'user1',
'HASH1', 'blah, blah');
SELECT * FROM myschema.account;
```

The audit log below can be retrieved.

'DDL' is not defined in the class parameter, so CREATE TABLE is not output as an audit log.

```
AUDIT: SESSION,WRITE,2022-03-12 19:00:49 PDT,[local],19944,psql,appuser,postgres,
2/7,1,1,INSERT,,TABLE,myschema.account,, "INSERT INTO myschema.account (id, name,
password, description) VALUES (1, 'user1', 'HASH1', 'blah, blah');" <not logged>
AUDIT: SESSION,READ,2022-03-12 19:00:58 PDT,[local],19944,psql,appuser,postgres,
2/8,2,1,SELECT,,TABLE,myschema.account,,SELECT * FROM myschema.account;<not logged>
```

## 7.6 Object Audit Logging

In Object Audit Logging, retrieval of audit logs is achieved by using roles.

Roles are specified in the role parameter of the option section to retrieve audit logs. If there are privileges for commands executed by a role, or if privileges have been inherited from another role, those command operations are output as audit logs.

For example, after "auditor" is set for the role parameter of the option section, the SELECT and DELETE privileges for the account table are assigned to "auditor". In this case, when SELECT or DELETE is executed for the account table, audit logs are output.

### Output format

In Object Audit Logging, audit logs are output in the format below:

```
AUDIT: OBJECT,1,1,READ,SELECT,TABLE,public.account,SELECT password FROM account;<not
logged>
      (1)      (2) (3) (4)  (5)      (6)      (7)      (8)      (9)
```

No	Content
(1)	Log header Fixed as "AUDIT: OBJECT".
(2)	Statement ID
(3)	Substatement ID
(4)	Class name
(5)	Command tag
(6)	Object type
(7)	Object name
(8)	SQL If "on" is specified for the log_statement_once parameter of the option section in the pgaudit configuration file, "<previously logged>" is output for the second and subsequent statements.
(9)	Depending on the log_parameter parameter value of the option section in the pgaudit configuration file, the output content will be as below. <ul style="list-style-type: none"> <li>- When log_parameter=on  If parameters are specified in the SQL, the parameters are concatenated and output, using a comma as the delimiter.  If parameters are not specified in the SQL, "&lt;none&gt;" is output.</li> <li>- When log_parameter=off (default)  "&lt;not logged&gt;" is output.</li> </ul> Additionally, if "on" is specified for the log_statement_once parameter of the option section in the pgaudit configuration file, "<previously logged>" is output for the second and subsequent statements.

## Information

If accessing resources that use the features below, the command tag (5) may be output as "???".

- INSTEAD OF trigger
- RULE
- VIEW
- Security policy per row
- Table inheritance

## Example

Below is an example of retrieving logs in Object Audit Logging.

By setting the target for assigning privileges to roles in detail, log output can be controlled.

In the example below, log retrieval of the account table is controlled by the privileges assigned to the columns, however, log retrieval of the account\_role\_map table is controlled by the privileges assigned to the table.

### 1. Settings

The role parameter below is specified for the option section in the pgaudit configuration file.

```
[option]
role = 'auditor'
```

### 2. Defining a role

A role is defined for Object Audit Logging.

```
CREATE USER auditor NOSUPERUSER LOGIN;
```

### 3. Retrieving logs

Execute the SQL below from the client.

```
CREATE TABLE account
(
    id int,
    name text,
    password text,
    description text
);
GRANT SELECT (password) ON public.account TO auditor;
SELECT id, name FROM account;
SELECT password FROM account;
GRANT UPDATE (name, password) ON public.account TO auditor;
UPDATE account SET description = 'yada, yada';
UPDATE account SET password = 'HASH2';
CREATE TABLE account_role_map
(
    account_id int,
    role_id int
);
GRANT SELECT ON public.account_role_map TO auditor;
SELECT account.password, account_role_map.role_id
```

```
FROM account
INNER JOIN account_role_map ON account.id = account_role_map.account_id;
```

The audit log below can be retrieved.

In the account table, only the operations for columns that privileges have been assigned to are output as logs.

In the account\_role\_map table, privileges are assigned to the table, so operations performed for the table are output as logs.

```
AUDIT: OBJECT,4,1,READ,SELECT,TABLE,public.account,SELECT password FROM account;,<not
logged>
AUDIT: OBJECT,7,1,WRITE,UPDATE,TABLE,public.account,UPDATE account SET password =
'HASH2';,<not logged>
AUDIT: OBJECT,10,1,READ,SELECT,TABLE,public.account,"SELECT account.password,
account_role_map.role_id
FROM account
INNER JOIN account_role_map ON account.id = account_role_map.account_id;",<not
logged>
AUDIT: OBJECT,10,1,READ,SELECT,TABLE,public.account_role_map,"SELECT account.password,
account_role_map.role_id
FROM account
INNER JOIN account_role_map ON account.id = account_role_map.account_id;",<not
logged>
```

## 7.7 Database Multiplexing

This section describes audit log retrieval while in database multiplexing mode.

### 7.7.1 Setup

If setting up the audit log feature in a database multiplexing environment that has already been built, follow the procedure below.

1. Copy the pgaudit files

Copy the pgaudit files on the primary server and standby server.

Refer to step 1 in ["7.2 Setup"](#) for details on copying the pgaudit files.

2. Create the pgaudit configuration file

Create the pgaudit configuration file on the primary server. Copy the pgaudit configuration file you created to the standby server.

Refer to step 2 in ["7.2 Setup"](#) for details on creating the pgaudit configuration file.

3. Configure postgresql.conf

In the postgresql.conf file on the primary server and standby server, configure the parameters for using audit logs. Set the same values for the parameters.

Refer to step 3 in ["7.2 Setup"](#) and ["7.7.2 Configuring Audit Log Retrieval"](#) for details on the parameters to configure.

4. Configure the *serverIdentifier.conf* file of Mirroring Controller

In the *serverIdentifier.conf* file on the primary server and standby server, configure the parameters for using audit logs.

Refer to ["7.7.2 Configuring Audit Log Retrieval"](#) for details on the parameters to be set.

5. Start the instance

Start the instance of the primary server and standby server.

6. Create the pgaudit extension

Execute CREATE EXTENSION on the primary server to create a pgaudit extension.

Refer to step 5 in "7.2 Setup" for details on creating pgaudit extensions.

7. Configure the parameters in the pgaudit configuration file

Add/change the parameters of the pgaudit configuration file on the primary server. Copy the pgaudit configuration file with the added/changed parameters to the standby server.

Refer to "7.4 pgaudit Configuration File" and "7.7.2 Configuring Audit Log Retrieval" for details on the parameters to set.

8. Restart the instance

Restart the instance of the primary server and standby server.

## 7.7.2 Configuring Audit Log Retrieval

---

In database multiplexing mode, Mirroring Controller periodically accesses the database to check the multiplexing status and detect failure. Due to this, audit logs are also periodically retrieved, so log files become used up. Therefore, set the parameters below so that audit logs are not retrieved by Mirroring Controller.

postgresql.conf

log\_connections

Omit, or specify "off".

log\_disconnections

Omit, or specify "off".

serverIdentifier.conf file of Mirroring Controller

target\_db

Specify "template1".



### Information

.....  
If creating a new database, create it after stopping Mirroring Controller, or specify a name other than "template1" for the template database.  
.....

pgaudit configuration file

option section audit\_log\_disconnections

If you want to output disconnection logs for connections other than Mirroring Controller, specify audit\_log\_disconnections = "on".

rule section database

Specify database != 'template1'.

## 7.8 Failover

---

This section explains audit log retrieval performed by the failover feature integrated with the cluster software.

### 7.8.1 Configuring Audit Log Retrieval

---

The failover feature accesses the database periodically to check the instance status. Due to this, audit logs are also periodically retrieved, so log files become depleted. Therefore, set the parameters below so that audit logs are not retrieved.

postgresql.conf

log\_connections

Omit, or specify "off".

log\_disconnections

Omit, or specify "off".

pgaudit configuration file

rule section database

Specify database != 'template1'.

## 7.9 When Used with set\_user

---

The audit log feature of Fujitsu Enterprise Postgres is configured to output audit logs even for operations performed after switching to a superuser with the `set_user_u` function of `set_user`. In addition, when a user is switched to a superuser using the `set_user_u` function, `set_user` automatically sets `log_statement = all`. As a result, detailed logs for operations performed after privilege elevation are also written to the server log. For this reason, when `set_user` and the audit log feature are used together, logs related to operations after elevation will be output to both the audit log and the server log. This leads to the following issues:

- Increased operational burden due to increased log volume

Since logs with the same content are output to multiple log files, the log volume increases, which can affect operations such as log storage, analysis, and monitoring.

- Expansion of the output range for sensitive information

Due to the effect of `log_statement = all`, sensitive information such as SQL statements and parameters is output not only to the audit log but also to the server log. Server logs may have different access controls compared to audit logs, which poses a risk of sensitive information being output to an unintended scope.

To avoid these problems and appropriately control the log output destination and content, we recommend one of the following configurations:

### Consolidate audit logs into a dedicated log file

If the log output destination is set to a dedicated log file, logs after elevation by the `set_user_u` function are also configured to be output only to the dedicated log file. Set the following GUC parameter for `set_user` to suppress log output by `set_user`.

- Set the `set_user.block_log_statement` parameter in the `postgresql.conf` file to off.

By suppressing log output to the server log, you can prevent sensitive information from being output to the server log.

### The output destination is the server log, but the output content and format should match the audit log feature

If you want to set the log output destination to the server log and match the output content and format to the audit log feature, set the following parameters and configure it to suppress log output by `set_user`.

- Set the `set_user.block_log_statement` parameter in the `postgresql.conf` file to off.

By suppressing unnecessary log output, you can curb the increase in log volume and reduce operational load.

### Prioritize logs automatically output during elevation by set\_user

If you prioritize the behavior of log output during elevation by `set_user`, configure it to suppress log output of post-elevation operations by the audit log feature. Set the following parameters to suppress audit log output after user switching by `set_user` by the audit log feature.

- Set the `audit_set_user_switch` parameter in the option section of the `pgaudit` configuration file to off.

By suppressing unnecessary log output, you can curb the increase in log volume and reduce operational load.

## Information

- When using the audit log feature in combination with `set_user`, ensure that "set\_user" is specified before "pgaudit" in the `shared_preload_libraries` parameter of `postgresql.conf`.

```
shared_preload_libraries = 'set_user,pgaudit'
```

- If you set the `set_user.block_log_statement` parameter to "off", users who have switched to the superuser via `set_user_u` will be able to modify the server log settings (the `log_statement` parameter). Ensure that the settings are configured to collect the necessary audit information in the audit log.

## 7.10 Analyzing Audit Logs in SQL

In general, you can quickly and easily analyze the audit log by creating a table for analysis and loading the audit log into that table.

Maybe the analysis table is not updated, so we recommend creating that table as an UNLOGGED table. This is because when loading, the write-ahead log (WAL) is not output and it is not replicated, so it can be loaded faster without interfering with other tasks.

### How to Identify the Audit Log Files to Load

To analyze audit logs for a specific time period, load all files that contain audit logs for that time period.

However, do not access dedicated log files with the latest update date. This is because it contains incomplete audit log data that is being written.

### Loading Method

The audit log is output in csv format, and can be loaded using the COPY FROM statement or the `pgx_loader` command.

The following example shows how to load Session Audit Logging output to a dedicated log file.

#### 1. Creating Tables

Define a table with the columns required to reference the audit log.

```
$ psql
=# CREATE UNLOGGED TABLE auditlog (
header text,
class text,
sql_start_time timestamp with time zone,
remote_host_name text,
backend_process_id integer,
application_name text,
session_user_name text,
database_name text,
virtual_transaction_id text,
statement_id integer,
substatement_id integer,
command_tag text,
sqlstate text,
object_type text,
object_name text,
error_message text,
sql text,
parameter text
);
```

#### 2. Creating an Index

For example, creating an index on a timestamp is useful for time-series analysis.

```
$ psql
=# CREATE INDEX auditlog_time_index ON auditlog (sql_start_time);
```

### 3. Loading the Audit Log

Specify a dedicated log file in the FROM clause of COPY FROM.

```
$ psql
=# COPY auditlog FROM
'pgaudit-2024-02-02_150000.log' WITH CSV DELIMITER ',';
```

You can also use the PROGRAM clause to load multiple dedicated log files together.

The following example loads one day of audit logs by specifying wildcard for day part of file name:

L

```
$ psql
=# COPY auditlog FROM
PROGRAM 'cat pgaudit-2024-02-02_*' WITH CSV DELIMITER ',';
```

W

```
$ psql
=# COPY auditlog FROM
PROGRAM 'type pgaudit-2024-02-02_*' WITH CSV DELIMITER ',';
```

### 4. Audit Log Analysis

Access the table that loaded the audit log.

Loaded records are not in chronological order, but can be sorted by timestamp

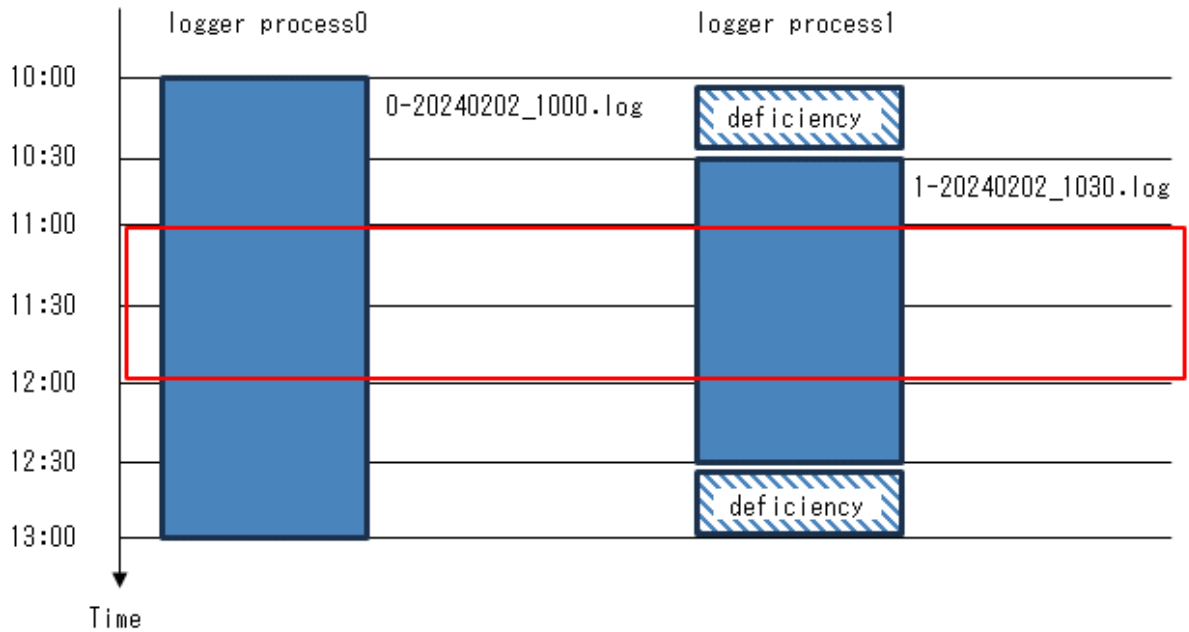
```
$ psql
=# SELECT * FROM auditlog ORDER BY sql_start_time;
header      | class | sql_start_time | remote_host_name | ...
-----+-----+-----+-----+-----
AUDIT: SESSION | DDL   | 2024-02-02 15:00:49+09 | ::1              | ...
AUDIT: SESSION | SYSTEM | 2024-02-02 15:00:58+09 |                  | ...
```

## Considerations when Using the Scalable Audit Log Feature

- When you enable the scalable audit log feature, there are many dedicated log files for the target period as you specify for `parallel_loggers`, so make sure to load all the dedicated log files.
- When you disable `log_rotation_age`, do not analyze the time except target period because not all audit logs may have been loaded.

For example, suppose that you are analyzing the audit log at 11:00-12:00, and you load the file output by logger process 0 at 10:00-13:00 and the file output by logger process 1 at 10:30-12:30, as shown in the following figure. Audit log records up to 10:00-13:00 appears in the table loaded with log files, but some audit logs (Audit log output by logger process 1 at 10:00-10:30 and 12:30-13:00) have not been loaded.

When `log_rotation_age` is enabled, the timing of the dedicated log file changeover is consistent across all logger processes to prevent such erroneous analysis.



## 7.11 Removing Setup

This section describes how to remove the setup of pgaudit.

1. Start the instance
2. Remove the pgaudit extension

Execute `DROP EXTENSION` to remove the pgaudit extension from the database.

```
$ psql -d <database name>
=# DROP EXTENSION pgaudit;
=# \q
```

3. Change postgresql.conf

Remove the parameters below relating to pgaudit.

- `shared_preload_libraries`
- `pgaudit.config_file`

4. Restart the instance

5. Remove the pgaudit files



As superuser, run the following command. Note that "<x>" in paths indicates the product version.

```
$ su -
Password:*****
# rm -rf /opt/fsepv<x>server64/filesCopiedDuringSetup
```



### Information

The files copied during setup can be checked below.

```
# find /opt/fsepv<x>server64/OSS/pgaudit
```

W

Open a command prompt as administrator privileges, run the following command. Note that "<x>" in paths indicates the product version.

```
> del "c:\Program Files\Fujitsu\fsepv<x>server64\filesCopiedDuringSetup"
```



### Information

The files copied during setup can be checked below.

```
> dir /b /s "c:\Program Files\Fujitsu\fsepv<x>server64\OSS\pgaudit"
```

# Chapter 8 Confidentiality Management

In order to prevent unauthorized use of the various data stored in the database, it is necessary to set appropriate privileges for each database resource for database users. However, there are multiple users to whom privileges are granted, and there are also many target database resources. So it takes a lot of effort to set it up. The confidentiality management feature reduces the time and effort required to do so and supports the setting and maintenance of appropriate privileges.

This chapter describes how to install, design and use confidentiality management features. For usage examples of the feature, refer to "[8.8 Usage Example of Confidentiality Management](#)".

## 8.1 Setup

This feature is provided as an EXTENSION of PostgreSQL. The name is `pgx_confidential_management_support`. Register the extension with the database cluster using the CREATE EXTENSION statement as follows:

This must be run by superuser. Because this extension registers PostgreSQL event triggers. By registering an event trigger, when a database object such as a table is deleted, related information is deleted from the information managed by the confidentiality management feature.

This extension can be created for any schema.

```
CREATE EXTENSION pgx_confidential_management_support
```



### Information

- The various definitions described below are registered in the tables included in this extension. Note that dropping this extension with the DROP EXTENSION statement will therefore drop all these definitions as well.
- If the superuser also serves as the confidentiality management role, or if you have only one confidentiality management role, setup is complete. For confidentiality management roles, refer to "[8.2.2 Determining Confidentiality Management Roles](#)". However, if multiple confidentiality management roles manage different confidentiality matrices, they must be prevented from manipulating each other's confidentiality matrices. To do so, a superuser should run the script provided by the product as follows. This script defines policies for the row level security feature on the tables provided by the confidentiality management. `${install_dir}` refers to the directory where you installed the product.

```
psql -f ${install_dir}/share/extension/pgx_confidential_management_support_policy.sql
```

- 'public' is granted SELECT on tables included by this extension when CREATE EXETENSION statement is executed. Don't revoke the privilege from 'public'. For example, event trigger of the extension confirms to update its tables when a general user drop some table. The privilege is required at the time. Also, there is no problem that users refer content of the tables like `pg_catalog`.

## 8.2 Designing Confidentiality Management

Describes the design of confidentiality management.

### 8.2.1 Designing a Confidentiality Matrix

A confidentiality matrix is a matrix of confidentiality levels and confidentiality groups. Elements of the confidentiality matrix are confidentiality privileges. Here we define them.

A confidentiality level is a group of data with the same degree of confidentiality, and a confidentiality group is a group of roles with the same access to confidentiality data. The details are described below.

The end result is to classify database objects into confidentiality levels and roles into confidentiality groups. But for the first step, define your confidentiality matrix abstractly without thinking about concrete tables, roles, etc. By doing so, the relationship between confidentiality levels and confidentiality groups can also be applied to databases with different sets of database objects.

You can define multiple confidentiality matrices. These confidentiality matrices are identified by a name, but the name is unique within the database (not the database cluster).

A role that manages the confidentiality matrix is called a confidentiality management role. The details are described below.

## How to create multiple confidentiality matrices

There are various design methods, but an example is shown below.

- Create one confidentiality matrix for a dataset
  - For example, create one confidentiality matrix for one section of your organization.
- Copy the confidentiality matrix if the two sections have the same confidentiality design
- When multiple organizations share datasets, create a single confidentiality matrix for the shared datasets

An example of this case is shown below.

- Create a confidentiality matrix *A* for the dataset that only section 1 has access to.
- Create a confidentiality matrix *B* for the dataset that only section 2 has access to.
- Create a confidentiality matrix *S* for the dataset shared by section 1 and section 2.
- Section 1 administrator *M1* is in charge of the confidentiality management role for confidentiality matrix *A*.
- Section 2 administrator *M2* is in charge of the confidentiality management role for confidentiality matrix *B*.
- Role group *G* to which *M1* and *M2* belong is in charge of confidentiality management role of confidentiality matrix *S*.

### 8.2.1.1 Defining Confidentiality Levels

Confidentiality level is a concept that indicates degree of confidentiality. For example, determine a confidentiality level that classifies tables that contain personal information and those that do not.

A collection of data that you classify as confidentiality level is a database object, such as a table. Such objects are called confidentiality objects. Refer to "[What is a confidentiality objects?](#)" for confidentiality objects.

Attributes can be set for the confidentiality level. For example, if a confidentiality level has an attribute that requires encryption, tables belonging to that confidentiality level are required to be encrypted. The confidentiality management feature may or may not automatically change the attribute of the confidentiality object to which it belongs. Also, each attribute targets a different confidentiality object.

The table below shows the attribute description, the confidentiality object targeted, and whether to change the attribute of the confidentiality object automatically.

Confidentiality level attributes	Specifiable values	Confidentiality object type	Auto change	Description
encryption_algorithm	AES128 AES256 none	table column rowset	No	Requires confidentiality objects to be encrypted with the specified algorithm. For details, refer to "Protecting Storage Data Using Transparent Data Encryption" in the Operation Guide. If <i>none</i> is specified, no encryption is requested. Default is <i>none</i> . The attribute is not changed automatically because encryption and decryption involve large data updates.

Confidentiality level attributes	Specifiable values	Confidentiality object type	Auto change	Description
				Instead, you cannot add a confidentiality object with an encryption strength lower than the specified encryption strength. Also, when changing to increase the encryption strength, confidentiality objects that do not meet the post-change conditions must not be included.

### What is a confidentiality objects?

Confidentiality objects have the types shown in the following table. For example, some data contained in table *T* can be added to confidentiality level *L1* and other data to confidentiality level *L2*.

Confidentiality object type	Description
schema	It does not mean all the tables contained in the schema. It simply means an object that exists in the system catalogs. Means the schema itself that is the target of access control by the GRANT statement.
table	Includes views, materialized views and partitions as well as tables. Currently do not support foreign tables.
column	Means column.
rowset	A set of rows that satisfy a specified condition.  Although there is no rowset object in the PostgreSQL system catalog, it is treated as a confidentiality object in the confidentiality management feature for ease of use.  Access control for rowset-type confidentiality objects uses PostgreSQL's row-level security functionality internally. Therefore, how access is controlled follows the row-level security specification.

The following objects may also contain confidential information. However, the current confidentiality management feature does not support them as confidentiality objects. When managing these, it is recommended to manage them in a role different from the confidentiality management role. The details are described in "[8.5 Suggestions for Monitoring Methods](#)".

- Function
- Procedure
- Foreign server
- Foreign data wrapper
- Foreign table

### 8.2.1.2 Defining Confidentiality Groups

A confidentiality group is an object that indicates which confidentiality level roles in confidentiality group can access to. For example, decide which role groups have access to personal information and which groups do not. Confidentiality groups are actually automatically defined as PostgreSQL roles (role groups). This role is called a confidentiality group role.

### What is a confidentiality group role?

Confidentiality group roles are targets of GRANT and REVOKE statements executed internally by the confidentiality management feature. And when you add a role to the confidentiality group, the added role becomes a member of the

confidentiality group role. The privileges of the added role are given by inheriting the privileges given to the confidentiality group role by the confidentiality management feature.

The following attributes can be set for the confidentiality group role via the function that creates the confidentiality group. The meaning and default value of attribute are the same as in the CREATE ROLE statement specification. These are limited to attributes that may be necessary when managing confidentiality. Therefore, use the confidentiality management feature to change it. The confidentiality management feature automatically sets the following attributes for the role added to the confidentiality group.

- SUPERUSER
- CREATEDB
- CREATEROLE
- REPLICATION
- BYPASSRLS

Other attributes follow the PostgreSQL CREATE ROLE statement defaults. Even if the attribute is changed using the ALTER ROLE statement, the operation of confidentiality management feature will not be disturbed.

The naming convention for confidentiality group roles is as follows.

```
pgx_cgroup_role_${serial_number}
```

When deleting a confidentiality group, you can choose not to delete the confidentiality group role. In such a case, it is a means to know that it was a confidentiality group role later.

### Information

Do not define roles using strings that follow this naming convention. This is because the name of the new confidentiality group role will be the number next to the maximum *serial\_number* of the role that follows this naming convention. For example, if a role named 'pgx\_cgroup\_role\_999999999999999999' exists, you will not be able to create new confidentiality group roles.

## 8.2.1.3 Defining Confidentiality Privilege

Confidentiality privileges indicate how a confidentiality group can access confidentiality objects registered at a certain confidentiality level.

For example, a confidentiality group *R* can be defined to have SELECT, INSERT, UPDATE, and DELETE privileges on table-type confidentiality objects belonging to confidentiality level *L*.

Once a confidentiality object or role has been added to a confidentiality level or confidentiality group, the feature automatically uses the GRANT or CREATE POLICY statement to grant privileges to the actual database object according to the confidentiality privilege definition to the confidentiality group role.

For rowset types, confidentiality privileges are those that can be specified in the CREATE POLICY statement. In cases other than rowset, it is a privilege that can be specified in the GRANT statement. For detailed privileges, refer to the respective SQL statement reference in the PostgreSQL Documentation.

### Information

Do not change the definition of POLICY that this feature creates internally. Even if you change it, the function of this feature may return to the state before the change without warning.

## Point

---

For roles registered in the confidentiality matrix, if the privilege to a certain confidentiality object is more than the confidentiality privilege, revoke those privileges according to the confidentiality privilege.

If a role not registered in the confidentiality matrix or PUBLIC was granted privileges to a confidentiality object registered in the confidentiality matrix, functions of this feature for adding confidentiality objects or roles will fail. Because this function is not allowed to manage.

Even in a slightly more complicated situation, failing functions that add confidentiality objects and roles prevent roles added to the confidentiality matrix from deviating from confidentiality privileges:

- It is set so that a role registered in the confidentiality matrix can inherit a role not registered in the confidentiality matrix, and as a result, it has more privileges than the confidentiality privileges.
  - Roles registered in the confidentiality matrix are granted privileges by a POLICY not created by this feature. In this case, the function will fail if such a POLICY exists without rigorously checking for violations of confidentiality privilege.
- 

## 8.2.2 Determining Confidentiality Management Roles

---

The confidentiality management role performs all operations on the confidentiality matrix. As such, the confidentiality management role requires very strong privileges to either:

1. Has all of the following rights
  - Superuser privileges
  - Ownership of database objects belonging to the confidentiality level
2. Has all of the following privileges
  - CREATEROLE privilege
  - SELECT, INSERT, UPDATE, and DELETE privileges on all tables included in the extension 'public' is granted SELECT when CREATE EXTENSION statement is executed.
  - Ownership of database objects belonging to the confidentiality level

## Information

---

Note that the previous owner may not be able to execute GRANT statements, etc. once the confidentiality management role becomes the owner. This is because ownership of database objects cannot be shared by multiple roles belonging to different role groups. This is the PostgreSQL specification.

---

The confidentiality management role can create and manage multiple confidentiality matrices, but it is safer to distribute the authority. For that reason, we recommend that you decide on confidentiality management roles in the following order of priority.

- Create one-to-one confidentiality matrices and confidentiality management roles.
- Create a role group with multiple confidentiality management roles as members, or manage multiple confidentiality matrices with one confidentiality management role.
- A superuser is also a confidentiality management role.

## Information

---

- To assign attributes that only a superuser can grant to a role managed using the confidentiality management feature, the superuser must also serve as the confidentiality management role. For example, the REPLICATION attribute is such an attribute. Refer to the reference of the PostgreSQL Documentation for details.

- If you migrate from Fujitsu Enterprise Postgres 15 SP2 or earlier, the CREATEROLE privilege has been changed. For this reason, if you want to use a non-superuser role as the confidentiality management role, you must first add the privilege (CREATEDB, BYPASSRLS, etc.) that you set for creating and updating confidentiality groups to the confidentiality management role. If these privileges are not set, the confidentiality group operation may fail.
- You cannot use the confidentiality management feature (SQL function) when using the user switching feature provided by set\_user. Therefore, you must log in to the database directly with the confidentiality management role.

## 8.2.3 Classify Confidentiality Objects According to the Definition of Confidentiality Level

### 8.2.3.1 Defining Confidentiality Objects

First, refer to list of definitions of schemata, tables in your database, etc. that you want to manage, and define confidentiality objects. As mentioned earlier, there are various types of confidentiality objects such as column type and rowset type, so you can define confidentiality objects based on their data content.

Access control for rowset-type confidentiality objects uses PostgreSQL's row-level security functionality internally. Therefore, the method of specifying a set of rows follows the specifications of the AS clause, USING clause, and WITH CHECK clause of the CREATE POLICY statement. Also, the definition of the rowset-type confidentiality object is specified in the argument of the function that registers the rowset-type confidentiality object to the confidentiality level. This feature executes the ALTER TABLE statement with the ENABLE ROW LEVEL SECURITY clause to enable POLICY when a rowset-type confidentiality object is added.

Currently it is not possible to register a foreign table as a table-type confidentiality object.

#### Point

Recommend that you do not grant privileges on confidentiality objects to PUBLIC. Granting privileges to PUBLIC is the same as granting privileges to all roles registered in the confidentiality matrix. This makes no sense if all roles that access confidentiality objects are managed using this feature. If PUBLIC is granted privileges to confidentiality objects, the functions included in this feature will check that the privileges granted to each role are not exceeded, and will fail if they are.

#### Information

- Be careful when confidentiality objects are of column type. This is because if the table-type confidentiality object is set at the same time, the table type privilege takes precedence.

For example, if you want to revoke the SELECT privilege only from a special column *C* in some table *T*, list the columns other than column *C* and grant the SELECT privilege to them without granting the SELECT privilege to table *T*. This follows the PostgreSQL's GRANT statement specification for column.

If you have to enumerate a large number of columns, it might be a good idea to move the special columns to a new table and present your existing application with a VIEW that JOINS both tables.

- If the confidentiality object is of rowset type, set the same privileges for the table type as those specified for the rowset type. This is because when a SQL statement accesses data, it first checks that you have privilege to access the table. After that, the rowset privileges are checked for matching rows. This follows the PostgreSQL row level security specification.

### 8.2.3.2 Classify Confidentiality Objects

Classify confidentiality objects according to the confidentiality level definition.

A single confidentiality object cannot be classified in multiple confidentiality levels, even if they are in different confidentiality matrices.

## 8.2.4 Classify Roles According to Confidentiality Group Definitions

A role can belong to multiple confidentiality groups in different confidentiality matrices at the same time. However, a single role cannot be classified into multiple confidentiality groups within a single confidentiality matrix.

### Information

- We strongly recommend that you do not group confidentiality group roles into other confidentiality groups. This is because the confidentiality management feature does not prohibit such usage, but it probably only complicates security management. Also note that PostgreSQL does not allow cyclic groupings of roles.
- After adding roles to a confidentiality group, it is highly recommended that such roles not be made members of role groups not managed by the confidentiality matrix. This is because although the confidentiality management feature does not prohibit such situations, for example, unreasonably increasing the privileges of such role groups may become a way out for roles managed by the confidentiality management feature.  
Note that this feature does not allow roles with such loopholes to be added to confidentiality groups.
- If you migrate from Fujitsu Enterprise Postgres 15 SP2 or earlier, the CREATEROLE privilege has been changed. If you want to use a non-superuser as a confidentiality management roles, you need ADMIN OPTION privileges for the role to be added to the confidentiality group. Therefore, the roles that can be added to the confidentiality group must be:

- Roles created with confidentiality management role privileges
- A role that previously granted the ADMIN OPTION privilege for a role to a confidentiality management role

Example) To grant the confidentiality management role manager\_role only the ADMIN OPTION privilege for role user\_role1:

```
GRANT user_role1 TO manager_role WITH ADMIN TRUE, INHERIT FALSE, SET FALSE;
```

## 8.3 How to Use Confidentiality Management Feature (Definition)

Confidentiality management feature provides functions and tables.

Some functions define, change, or drop confidentiality matrices, etc. These functions will output an audit log indicating that they were executed, so you can later confirm that an illegal operation was performed.

In addition, you can directly refer to the table provided by this feature and check the defined contents using functions that help referencing.

Also, some functions output the attributes of confidentiality objects defined in the PostgreSQL system catalog and the attributes that should be set for that confidentiality object. You can use these to compare attributes.

For details, refer to "[Appendix B System Management Functions Used by Confidentiality Management Feature](#)".

### Tables included by the pgx\_confidential\_management\_support extension

Table name	Description
pgx_confidential_matrix	A list of confidentiality matrices. You can refer to the attributes of the registered confidentiality matrix, the update time, or the time when the confidentiality level or confidentiality group was registered or deleted.
pgx_confidential_level	A list of confidentiality levels. You can refer to the registered confidentiality level attributes, update time, or the time when a confidentiality object was registered to the confidentiality level or removed from the confidentiality level.

Table name	Description
pgx_confidential_group	A list of confidentiality groups. You can refer to the registered confidentiality group attributes, update time, or the time when a role was registered to the confidentiality group or removed from the confidentiality group.
pgx_confidential_privilege	A list of confidentiality privileges. You can refer to confidentiality privilege set for each confidentiality object, update time, and so on.
pgx_confidential_object	A list of confidentiality objects. You can refer to object attributes or update time, and so on.
pgx_confidential_role	A list of roles registered in the confidentiality group. You can refer to role attributes or update time, and so on.
pgx_confidential_policy	This is a list of policies created to set privileges for rowset-type confidentiality objects. You can refer to the name of the policy you created and the privileges it has set.  Rows in this table are inserted when you add a rowset-type confidentiality object.

Rows in each table are added, deleted, or updated when you execute functions to add, delete, or update definitions.

Also, if the target confidentiality object is deleted by a DROP TABLE statement, etc., it will also be deleted from the following tables.

- pgx\_confidential\_object
- pgx\_confidential\_policy

### Functions for adding, removing or updating definitions

These functions will print an audit log indicating that they were executed.

Function name	Description
pgx_create_confidential_matrix	Create a confidentiality matrix.
pgx_alter_confidential_matrix	Change the attributes of the confidentiality matrix.
pgx_drop_confidential_matrix	Drop the confidentiality matrix.
pgx_copy_confidential_matrix	Copy the confidentiality matrix.
pgx_create_confidential_level	Create confidentiality levels and register them in the confidentiality matrix.
pgx_alter_confidential_level	Change the confidentiality level attribute.
pgx_drop_confidential_level	Remove a confidentiality level from the confidentiality matrix and drop a confidentiality level.
pgx_create_confidential_group	Create a confidentiality group and register it in the confidentiality matrix.
pgx_alter_confidential_group	Change the attributes of a confidentiality group.
pgx_drop_confidential_group	Remove confidentiality groups from the confidentiality matrix and drop confidentiality group.
pgx_grant_confidential_privilege	Grant confidentiality privileges.

Function name	Description
pgx_revoke_confidential_privilege	Revoke confidentiality privileges.
pgx_add_object_to_confidential_level	Add confidentiality objects to the confidentiality level.
pgx_remove_object_from_confidential_level	Removes confidentiality objects from the confidentiality level.
pgx_add_role_to_confidential_group	Add a role to a confidentiality group.
pgx_remove_role_from_confidential_group	Remove a role from a confidentiality group.

### Functions that Support Definition Referencing and Comparison with System Catalogs

Function name	Description
pgx_get_attribute_of_objects	For all confidentiality objects registered in the specified confidentiality matrix, displays the attributes defined in the confidentiality matrix and the attributes actually set in the database.
pgx_get_attribute_of_roles	For all roles registered in the specified confidentiality matrix, displays the attributes defined in the confidentiality matrix and the attributes actually set in the system catalog.
pgx_get_privileges_on_level_and_group	Displays a list of combinations of confidentiality objects registered with the specified confidentiality level and roles registered with the specified confidentiality group, such that the following can be compared. <ul style="list-style-type: none"> <li>- Privileges specified by the confidentiality privilege settings that should be granted to the specified role</li> <li>- Privileges granted in the actual system catalog</li> </ul>
pgx_get_privileges_on_object	For the specified confidentiality object, display a list that allows you to compare the following: <ul style="list-style-type: none"> <li>- Privileges dictated by confidentiality privilege settings that should be granted to all roles</li> <li>- Privileges granted in the actual system catalog</li> </ul>
pgx_get_privileges_on_role	For all confidentiality objects, display a list that allows you to compare: <ul style="list-style-type: none"> <li>- Privileges specified by the confidentiality privilege settings that should be granted to the specified role</li> <li>- Privileges granted in the actual system catalog</li> </ul>
pgx_get_privileges_on_matrix	For all objects registered in the specified confidentiality matrix, display a list that allows you to compare: <ul style="list-style-type: none"> <li>- Privileges defined by confidentiality privilege settings that should be granted to all roles registered in the confidentiality matrix.</li> <li>- Privileges granted in the actual system catalog</li> </ul>

Describe the definition procedure.

### 8.3.1 Creating a Confidentiality Management Role

---

Create a confidentiality management role using the CREATE ROLE statement, or use an existing role as a confidentiality management role. Set the privileges and attributes shown in "8.2.2 Determining Confidentiality Management Roles" for the confidentiality management role.

#### Information

---

- Use caution when renaming a confidentiality management role. If you want to do so, delete all confidentiality matrices managed by that confidentiality management role, rename the confidentiality management role, and then define the same confidentiality matrix again. Otherwise, you will not be able to operate the confidentiality matrix. For example, you cannot change confidentiality privileges or remove confidentiality objects from the confidentiality level. If you accidentally renamed it first, change it back to the original name and then proceed as described above. In the future, we will simply allow the changed name of the confidentiality management role to be set in the confidentiality matrix.
  - If you want to delete the confidentiality management role, delete the confidentiality management role after deleting the confidentiality matrix. Otherwise, you will not be able to create a confidentiality matrix with the same name because you will not be able to delete the confidentiality matrix left behind. If you accidentally delete the confidentiality management role before deleting the confidentiality matrix, create a confidentiality management role with the same name again and delete the confidentiality matrix, or delete the confidentiality matrix with a role that has SUPERUSER privileges.
- 

### 8.3.2 Creating a Confidentiality Matrix

---

Create a confidentiality matrix with comments describing the confidentiality matrix as follows. This comment is stored in the pgx\_confidential\_matrix table.

```
select pgx_create_confidential_matrix('matrix_foo', 'This matrix is defined for foo.')
```

#### Point

---

The role executing this function is considered the confidentiality management role. If this function is executed after SET ROLE, the role specified in SET ROLE will be regarded as the confidentiality management role, not the role that executed SET ROLE.

---

You can also create a new confidentiality matrix 'matrix\_dest' by duplicating the already created confidentiality matrix 'matrix\_src' as follows. If you copy the confidentiality matrix and create it, the comments will also be copied. Comments can be changed using the pgx\_alter\_confidential\_matrix function.

```
select pgx_copy_confidential_matrix('matrix_dest', 'matrix_src')
```

You can also check the created matrix by referring to the pgx\_confidential\_matrix table.

The source must be within the same database. If you want to copy from different databases or different database instances, choose one of the following methods.

- Define using a function in the same way as the original.
- Follow the steps below. However, before proceeding, refer to the <Information> below.

#### Information

---

- If you set drop\_role to true in step 7, you will no longer be able to continue managing with the confidential matrix of the replication source. The reason is as follows.

After completing step 6, you should have:

- a) Confidentiality matrix of source database
- b) the confidentiality matrix loaded in step 4
- c) Confidentiality matrix duplicated in step 6

The b) and c) confidentiality matrices use different confidentiality group roles. This is the effect of the *pgx\_copy\_confidential\_matrix* function. However, the confidentiality matrices of a) and b) share the same confidentiality group role. Although this is temporary, it is a bad situation. To do so, remove the confidentiality matrix in b) in step 7. If *true* is specified for *drop\_role* at this time, the confidentiality group role used in a) will be deleted. Therefore, management by the confidentiality matrix of a) cannot be continued.

- The destination must have just installed this extension with CREATE EXTENSION. This is because the table data contained in this extension contains an ID representing the confidentiality matrix and confidentiality level, and this ID is generated by the SEQUENCE included in the extension. For example, if you have created at least one confidentiality matrix, the confidentiality matrix is numbered with an ID of 1, but that matrix is not the object numbered with an ID of 1 in the copy source.



Execute steps other than step 6 with a role that has SUPERUSER privileges.

1. Set up extensions in the target database.
2. From the table below, specify the ID of the source confidentiality matrix in the WHERE clause and extract the data using the COPY TO statement.
  - pgx\_confidential\_matrix
  - pgx\_confidential\_level
  - pgx\_confidential\_group
  - pgx\_confidential\_privilege
3. Refer to the *pg\_sequences* system view and save the value of the *last\_value* column of the following SEQUENCE in a file.
  - pgx\_confidential\_matrix\_cmatid\_seq
  - pgx\_confidential\_level\_clevid\_seq
  - pgx\_confidential\_group\_cgroid\_seq
  - pgx\_confidential\_privilege\_cpriid\_seq
4. Load the data extracted in step 2 into the above table in the target database using the COPY FROM statement. This operation can be performed by any confidentiality management role in any confidentiality matrix.
5. Use the *pg\_catalog.setval* function to set the *last\_value* for each SEQUENCE on the copy destination, specifying the value saved in step 3. An example is shown below.
 

```
SELECT pg_catalog.setval('pgx_confidential_matrix_cmatid_seq', 5)
```
6. Create a new confidentiality matrix using the *pgx\_copy\_confidential\_matrix* function and specifying the confidentiality matrix loaded in step 4 as the copy source. Perform this operation in the confidentiality management role of the newly created confidentiality matrix.
7. Use the *pgx\_drop\_confidential\_matrix* function specifying *false* for *drop\_role* to drop the confidentiality matrix loaded in step 4.

### 8.3.3 Adding Confidentiality Levels to the Confidentiality Matrix

Create a confidentiality level and add it to the confidentiality matrix as follows: In the example below, the JSON format third argument requests that confidentiality objects belonging to 'level1' are encrypted. Specify *NULL* if you don't want anything. Refer to "[B.2 Confidentiality Level Manipulation Functions](#)" for the content of the request.

```
select pgx_create_confidential_level('matrix_foo', 'level1',
'{"encryption_algorithm":"AES256"}', 'The strongest encryption is required for this
level.')
```

You can also check the added confidentiality level by referring to the `pgx_confidential_level` table.

### 8.3.4 Adding Confidentiality Groups to the Confidentiality Matrix

Create a confidentiality group and add it to the confidentiality matrix as follows. Attributes are given to the confidentiality group by the third argument in JSON format. Refer to "[B.3 Confidentiality Group Manipulation Functions](#)" for the attributes that can be assigned.

In the example below, a role belonging to this confidentiality group grants the `CREATEDB` privilege.

```
select pgx_create_confidential_group('matrix_foo', 'group1', '{"CREATEDB':true}', 'Roles
belonging to this confidentiality role are permitted to create db.')
```

You can also check the added confidentiality group by referring to the `pgx_confidential_group` table.

### 8.3.5 Granting Confidentiality Privileges to Confidentiality Groups

In the example below, roles belonging to 'group1' are granted `SELECT`, `INSERT` and `DELETE` privileges for tables belonging to 'level1'. Specify privileges for each type of confidentiality object by the fourth argument in JSON format. Refer to "[B.4 Confidentiality Privilege Manipulation Functions](#)" for the privileges that can be granted.

```
select pgx_grant_confidential_privilege('matrix_foo', 'level1', 'group1', '{"schema":
["ALL"], "table":["SELECT","INSERT","DELETE"]}')
```

If a confidentiality object of the target type is registered in the confidentiality level, when this function is executed, the confidentiality management feature internally uses the `GRANT` statement to grant a confidentiality group role access to confidentiality objects that are registered at a confidentiality level. The effect of the `GRANT` statement can be checked as follows.

```
select pgx_get_privileges_on_level_and_group('matrix_foo', 'level1', '["role1","role2"]')
```

For the format of the table returned by a function that outputs authority information, such as this function, refer to "[B.7 Functions that Support Definition Referencing and Comparison with System Catalogs](#)".

### 8.3.6 Adding Confidentiality Objects to Confidentiality Level

Classify the intended database object into the appropriate confidentiality level according to the confidentiality level design. In the example below, 'table1' of 'schema1' and 'table2' of 'schema1' are added to 'level1' at the same time. Specify the confidentiality object type and confidentiality object name in the third argument in JSON format. If the type is same, you can enumerate multiple confidentiality objects. You can enumerate multiple objects you want to register for each sensitive object type. In this example, only the *table* type is registered, but *schema* type can also be registered at the same time.

```
select pgx_add_object_to_confidential_level ('matrix_foo', 'level1',
'[{
  "type": "table",
  "object": [
    {
      "schema": "schema1",
      "table": ["table1", "table2"]
    }
  ]
}]')
```

When specifying a rowset, you must declare what kind of set of rows it is in the value of the *rowset\_expression* key. For details on how to specify it, refer to "[B.5 Confidentiality Object Manipulation Functions](#)".

When this function is executed, the confidentiality management feature internally uses the GRANT statement or CREATE POLICY statement to grant privileges to the specified confidentiality objects to all confidentiality groups set in the confidentiality matrix. The effect of the GRANT statement can be checked as follows. This example checks the privileges granted on 'table1' in 'schema1' and 'table2' in 'schema1'.

```
select pgx_get_privileges_on_object('matrix_foo',
  ' [{
    "type": "table",
    "object": [
      "schema": "schema1",
      "table": ["table1", "table2"]
    ]
  } ]')
```

### 8.3.7 Adding Roles to Confidentiality Groups

---

Classify the intended role into the appropriate confidentiality group according to the confidentiality group design. You can add multiple roles at once. In the example below, role role1 and role role2 are added to confidentiality group 'group1'.

```
select pgx_add_role_to_confidential_group('matrix_foo', 'group1', '["role1","role2"]')
```

The confidentiality management feature adds the specified role to the members of the confidentiality group role.

The feature does not execute GRANT statements for confidentiality objects at this time. This is because we have already executed the GRANT statement between the confidentiality group role and the confidentiality object, and we will not execute the GRANT statement between the individual roles and the confidentiality object.

By using the function below, you can check the privileges that can be exercised by the specified role after inheriting the privilege of the confidentiality group role with the INHERIT attribute on the role side or changing to the confidentiality group role with the SET ROLE statement. In other words, you can check the access privileges that the actually registered roles can exercise. For details on the output format, refer to "[B.7 Functions that Support Definition Referencing and Comparison with System Catalogs](#)". The example below checks the privileges of role 'role1' and role 'role2'.

```
select pgx_get_privileges_on_role('matrix_foo', '["role1","role2"]')
```

## 8.4 How to Use Confidentiality Management Feature (Change and Deletion)

---

### 8.4.1 Renaming Confidentiality Objects

---

The current confidentiality management feature cannot follow when confidentiality objects are renamed. Therefore, if the name is changed, remove the old name confidentiality object with the *pgx\_remove\_object\_from\_confidential\_level* function specifying the old name, and register the confidentiality object with the new name using the *pgx\_add\_object\_to\_confidential\_level* function specifying the new name.

### 8.4.2 Renaming Roles

---

The current confidentiality management feature cannot follow when roles are renamed. Therefore, when renaming, exclude the role with the old name with the *pgx\_remove\_role\_from\_confidential\_group* function specifying the old name, and after renaming the role, register the role with the new name with the *pgx\_add\_role\_to\_confidential\_group* function specifying the new name.

### 8.4.3 Deleting Roles

---

The current confidentiality management feature cannot follow when roles are dropped. Therefore, when dropping, exclude the role with *pgx\_remove\_role\_from\_confidential\_group* function.

## 8.4.4 Changing Confidentiality Matrix

---

Change the name or attributes of the confidentiality matrix as follows. Specify the name and attribute value to be changed in JSON format. You can change the name and multiple attributes at once. Attributes not specified here are not changed.

```
select pgx_alter_confidential_matrix('matrix_foo', '{"name": "matrix_bar",
"comment": "This matrix is defined for bar."}')
```

## 8.4.5 Deleting Confidentiality Matrix

---

Delete the confidentiality matrix as follows.

```
select pgx_drop_confidential_matrix('matrix_foo', false, false)
```

By specifying a second or third argument, you can choose to delete all confidentiality levels and confidentiality groups added to the confidentiality matrix, or remove confidentiality group roles. Confidentiality group roles that were not deleted can be identified using the previously mentioned naming convention. No matter what you choose, the roles registered in the confidentiality group will not be deleted. The above example simply deletes the confidentiality matrix only.

## 8.4.6 Changing Confidentiality Level

---

Change the name and attributes of the confidentiality levels that were added to the confidentiality matrix 'matrix\_foo' as follows. Specify the value after changing the name and attributes in JSON format. You can change the name or multiple attributes at once. Attributes not specified here are not changed.

```
select pgx_alter_confidential_level('matrix_foo', 'level1', '{"name": "levelX",
"comment": "This level required the highest confidential clearance."}')
```

If a confidentiality object was already registered with the confidentiality level, the attributes of the confidentiality object are also changed automatically. However, as mentioned in ["8.2.1.1 Defining Confidentiality Levels"](#), it may not be possible to automatically change the attributes of confidentiality objects. At that time, the function checks that the attributes of the registered confidentiality object are stricter than the attribute of the modified confidentiality level. If the check fails, this function will also fail. For example, if you try to change the encryption\_algorithm from AES128 to AES256 and an AES128 encrypted table is registered as a confidentiality object, this function will fail. In such a case, please execute this function again after encrypting the table with AES256.

## 8.4.7 Deleting Confidentiality Level

---

Delete the confidentiality level as follows.

```
select pgx_drop_confidential_level('matrix_foo', 'level1', false)
```

With a third argument, you can choose to delete those that depend on the specified confidentiality level. Of course, confidentiality objects registered with the confidentiality level are not deleted.

## 8.4.8 Changing Confidentiality Group

---

Change the name and attributes of the confidentiality groups that were added to the confidentiality matrix 'matrix\_foo' as follows. Specify the value after changing the name and attributes in JSON format. You can change the name or multiple attributes at once. Attributes not specified here are not changed.

```
select pgx_alter_confidential_group('matrix_foo', 'group1', '{"name": "groupX",
"comment": "Members of this group have the highest confidential clearance."}')
```

## 8.4.9 Deleting Confidentiality Group

---

Delete the confidentiality group as follows.

```
select pgx_drop_confidential_group('matrix_foo', 'group1', true, true)
```

By specifying a second or third argument, you can choose to delete dependencies on the specified confidentiality group or delete the confidentiality group role. Any choice will not delete the roles belonging to the confidentiality group.

### Information

Even if you leave the confidentiality group role, this function will revoke privileges from the confidentiality group role. The privilege to revoke is the privilege defined in confidentiality privileges. Therefore, note that a role registered in a confidentiality group will be deprived of various privileges granted by inheriting the confidentiality group role.

## 8.4.10 Revoking Confidentiality Privileges

Revoke confidentiality privileges from confidentiality groups as follows. Specify the confidentiality object and confidentiality group and the privileges to revoke in JSON format. The following example revokes INSERT and DELETE privileges for table-type confidentiality objects belonging to 'level1' from 'group1'.

```
select pgx_revoke_confidential_privilege('matrix_foo','level1', 'group1', '{"table":
["INSERT", "DELETE"]}')
```

## 8.4.11 Removing Confidentiality Objects from Confidentiality Level

Remove confidentiality objects from the confidentiality level as follows. Specify the type and name of confidentiality objects to be excluded in JSON format. The following example removes table-type confidentiality objects schema1.table1 and schema1.table2 from the confidentiality level 'level1'.

```
select pgx_remove_object_from_confidential_level('matrix_foo', 'level1'
'[{
  "type": "table",
  "object": [
    {
      "schema": "schema1",
      "table": ["table1", "table2"]
    }
  ]
}]')
```

## 8.4.12 Removing Roles from Confidentiality Groups

Remove the role from the confidentiality group as follows. Specify the roles to remove in JSON format. The following example removes roles 'role1' and 'role2' from confidentiality group 'group1'.

```
select pgx_remove_role_from_confidential_group('matrix_foo', 'group1',
'["role1", "role2"]')
```

## 8.5 Suggestions for Monitoring Methods

Confidentiality management role must ensure that the database is operating securely as intended, even after confidentiality managements are defined. If you are only using the confidentiality management feature, you do not have to worry about such things. However, the confidentiality management feature does not prohibit changing the definitions of tables and roles without using this feature.

So you have to detect when such an action has taken place.

However, even if they detect it, they may forget to deal with it. Therefore, it is necessary to periodically check the difference between the confidentiality level and confidentiality group and the actual definition of confidentiality objects and roles. Of course, even if there was a mismatch, it would not be a problem if the confidentiality object or role had stricter attributes and privilege.

The procedure presented here aims at matching.

## 8.5.1 How to Detect Privilege Changes without Using Confidentiality Management feature

---

Use the audit log to detect unauthorized modification of the attributes of confidentiality objects or roles, or modification of privileges without going through the confidentiality management feature.

The basic method for detection is to detect actions by roles other than the confidentiality management role. However, there are exceptions such as:

When performing an operation without using the confidentiality management feature for a legitimate reason

For example, when changing the authority of a function that the confidentiality management feature does not treat as a confidentiality object. In order to identify this, it is recommended to determine roles that perform changes that do not involve the confidentiality management feature. This is because when various roles do this, it becomes difficult to detect audit logs that deviate from operational rules.

When monitoring the activity of the confidentiality management role

For example, it would be a good idea to create rules that allow access only at specified times and from specified terminals, and to detect activities that violate those rules from audit logs. It is important to set rules so that violations cannot be covered up. For example `application_name` is not suitable as it can be easily spoofed.

## 8.5.2 How to Check Confidentiality Objects and Roles

---

This section describes how to confirm that the contents registered in the confidentiality matrix match the confidentiality objects and roles managed in the PostgreSQL system catalog. The method shown here is just an example.

### Inspector

The inspector must be granted SELECT privilege on tables with the `pg_confidential_management_support` extension and SELECT privilege on the confidentiality object and role information added to the subject confidentiality matrix in the system tables. At a minimum, the confidentiality management role for the confidentiality matrix being reviewed and superuser certainly have such privileges.

### Procedure

The confirmation procedure is divided into the following three phases. If you detect a discrepancy, use the audit log to investigate what caused the discrepancy. And fix it to match.



### Information

---

- Various objects may be changed (modified) in a chain reaction when returning to the correct state. Doing so may erase traces of unauthorized manipulation. Therefore, it is recommended that you investigate any design inconsistencies and ensure that you have the necessary audit logs for the investigation before fixing them.
  - The `pgx_get_privileges_on_matrix` function presented here can output a very large table if the number of confidentiality objects or roles is large. If the size of this table exceeds the value of PostgreSQL's `work_mem` parameter, I/O will occur according to PostgreSQL's specifications, resulting in a slowdown. Therefore, it is recommended that `work_mem` be set as high as possible in the session in which this function is executed.
- 

#### (1) Make sure the state of the confidentiality matrix is the same as the design

If you detect a mismatch, please restore the correct state by executing functions such as `pgx_alter_confidential_level` function provided by this extension.

- Attributes of the confidentiality matrix

Check using the `pgx_confidential_matrix` table.

- Number of registered confidentiality levels and attributes of each confidentiality level

Check the row that matches the `clvmatid` to be checked from the `pgx_confidential_level` table. To know the confidentiality matrix identifier, see `cmatid` in the `pgx_confidential_matrix` table.

- Number of registered confidentiality groups and attributes of each confidentiality group

Check the row in the *pgx\_confidential\_role* table whose *ccolmatid* matches the identifier of the confidentiality matrix to be checked.

- Confidentiality group privileges set to confidentiality level

Refer to the *pgx\_confidential\_privileges* table. However, in this table confidentiality levels and confidentiality groups are represented as identifiers. Join with *pgx\_confidential\_level* table and *pgx\_confidential\_role* table if you want to check by confidentiality level name or confidentiality group name.

## (2) Verify correct confidentiality objects and roles registered

Confirm that the definition of the confidentiality management feature matches the definition of the confidentiality objects and roles. If you detect a mismatch, please restore the correct state by executing functions such as *pgx\_alter\_confidential\_level* function provided by this extension.

- Is the number of confidentiality objects registered in the confidentiality level the same as designed?
- Do the attributes of each confidentiality object match the attributes of the confidentiality level to which it belongs?

It is recommended to check these using *pgx\_get\_attribute\_of\_objects* function. Because this function gets the state of the confidentiality object from the PostgreSQL system catalog and outputs a table that can be compared with the definition. For the format of the table, refer to "[B.7 Functions that Support Definition Referencing and Comparison with System Catalogs](#)".

- Is the number of roles registered in the confidentiality group the same as designed?
- Do the attributes of each role match the attributes of the confidentiality group to which it belongs?

It is recommended to check these using *pgx\_get\_attribute\_of\_roles* function. For the format of the table returned by a function that outputs authority information, such as this function, refer to "[B.7 Functions that Support Definition Referencing and Comparison with System Catalogs](#)".

## (3) Check privilege of confidentiality objects

Verify that the roles granted access to confidentiality objects and what the privileges are are consistent with the definitions in the confidentiality matrix. A good help is to use the *pgx\_get\_privileges\_on\_matrix* function. For details on the output format, see "[8.3.5 Granting Confidentiality Privileges to Confidentiality Groups](#)".

If you want to focus on any confidentiality level or role, use the function below.

- *pgx\_get\_privileges\_on\_level\_and\_group()*
- *pgx\_get\_privileges\_on\_object()*
- *pgx\_get\_privileges\_on\_role()*

### Point

It is inefficient to output a large table many times in order to analyze it with various SQL statements. You can perform efficient analysis by specifying a query that executes this function in the INSERT statement as shown below.

```
INSERT INTO temp_table_for_analysis SELECT pgx_get_privileges_on_matrix('matrix_foo')
```

## 8.6 Backup/Restore

There are 4 methods below.

- Use *pg\_basebackup*
- Use *pg\_dumpall* to back up the entire database cluster as a script file
- Use *pg\_dump* and *pg\_restore* to back up and restore only some databases
- Use *pg\_dump*'s *pg\_restore* to back up and restore only some tables and schemas

Of these, if you adopt the method of using `pg_basebackup`, you do not need to be particularly careful. The following special precautions should be taken when adopting other methods.

### **Use `pg_dumpall` to back up the entire database cluster as a script file**

Note the following.

- Do not use the `-O` option of the `pg_dumpall` command. This feature does not work properly if the confidentiality object changes ownership.
- Do not use the `-x` option of the `pg_dumpall` command. Of course, this would change tightly controlled privileges.
- Execute the backup script file as a superuser. A backup script file contains the `CREATE EXTENSION` statement for this extension. This is because the `CREATE EXTENSION` statement must be executed as a superuser.

### **Backup and restore only some databases using `pg_dump` and `pg_restore`**

Note the following.

- Use the `-r` option of the `pg_dumpall` command to back up and restore role information as well. Of course, because roles are essential to confidentiality management.
- Do not use the `-O` option of the `pg_dump` command. This feature will not work properly if the confidentiality object changes ownership.
- Do not use the `-x` option of the `pg_dump` command. Of course, this would change tightly controlled privileges.
- Restore by a superuser for `CREATE EXTENSION`.

### **Use `pg_dump`'s `pg_restore` to backup and restore only some tables and schema**

Note the following.

- Use the `-e` option of the `pg_dump` command to back up and restore this extension as well. This is because the tables contained in this extension store the information required for confidentiality management.
- Use the `-r` option of the `pg_dumpall` command to back up and restore role information as well. Of course, because roles are essential to confidentiality management.
- Do not use the `-O` option of the `pg_dump` command. This feature will not work properly if the confidentiality object changes ownership.
- Do not use the `-x` option of the `pg_dump` command. Of course, this would change tightly controlled privileges.
- Restore by a superuser for `CREATE EXTENSION`.

## **8.7 Removing Setup**

---

If you are no longer using this feature and want to continue using your database as before, we recommend simply `DROP EXTENSION`. This is because if you delete definitions such as confidentiality matrices and confidentiality levels created using this extension using the deletion functions provided by this extension, the roles automatically created by this extension will be deleted, the privileges of roles granted by this extension are revoked. Simply dropping this extension will only drop the tables it contains, it will not do any of these things.

## **8.8 Usage Example of Confidentiality Management**

---

Here is an example of the concept of confidentiality levels and confidentiality groups.

This section assumes a simple business that handles customer purchase information.

First, create a confidentiality matrix for confidentiality management of this information.

```
SELECT pgx_create_confidential_matrix('matrix_purchase_managenet' , 'Confidentiality
management of customer purchase information');
```

The data we process may also contain personally identifiable information. Access to such data should be restricted to those who have access to it to minimize the risk of information disclosure.

Customer purchase information, including personally identifiable information, is managed in the following table.

```
CREATE TABLE purchase.customer_info(      -- Customer information
  customer_id    integer,                  -- Customer ID
  name           text,
  address        text,
  phone_number   char(12),
  rank           integer                    -- Customer's service rank
);

CREATE TABLE purchase.history(           -- History of a customer's purchase of goods
  customer_id    integer,                  -- ID of the customer who purchased
  purchase_date  date,                    -- Date the goods were purchased
  item_code      char(12),                 -- Code of the goods purchased
  purchase_number integer,                 -- Quantity purchased
  purchase_amount integer                 -- Amount of purchased goods
);
```

Among customer information, name, address, and telephone number are personal information because they can identify an individual when combined. In order to properly handle such information, we have made it so that it can only be handled by employees belonging to a specific group who have received appropriate training.

Therefore, we have prepared two confidentiality levels: "level\_personal\_info", which means highly confidential personal information, and "level\_customer\_info", which means other information.

```
SELECT pgx_create_confidential_level('matrix_purchase_management', 'level_personal_info',
                                     NULL, 'Personally identifiable information');
SELECT pgx_create_confidential_level('matrix_purchase_management', 'level_customer_info',
                                     NULL, 'Non-personally identifiable information');
```

In addition, we will prepare two confidentiality groups: "group\_qualified" who have been educated about handling personal information and can handle personal information appropriately, and "group\_non\_qualified" who are not qualified.

```
SELECT pgx_create_confidential_group ('matrix_purchase_management', 'group_qualified',
                                     NULL, 'Qualified staff handling personal information');
SELECT pgx_create_confidential_group ('matrix_purchase_management', 'group_non_qualified',
                                     NULL, 'General employee');
```

Let's take a closer look at the data we're dealing with.

Since the customer information table contains personal information, it corresponds to personal information. However, the customer\_id and rank contained in the customer information table are not personal information because they are not personally identifiable information. In addition, since this customer\_id and rank are also information necessary for business analysis, it is inconvenient that only those who are qualified to handle personal information can access such information.

Therefore, the customer information table uses columns for confidentiality management. The entire customer information table is protected as personal information, and the range of access is expanded by making the columns that are not personal information general customer information.

Follow this policy to set confidentiality level privilege for confidentiality group.

```
SELECT pgx_grant_confidential_privilege('matrix_purchase_management',
                                         'level_personal_info',
                                         'group_qualified', '{"table":["ALL"]}');
SELECT pgx_grant_confidential_privilege('matrix_purchase_management',
                                         'level_customer_info',
```

```

                                'group_qualified', '{"table":["ALL"]}');
SELECT pgx_grant_confidential_privilege('matrix_purchase_management',
                                'level_customer_info',
                                'group_not_qualified',
                                '{"table":["ALL"], "column":["SELECT"]}');

```

Only "qualified personnel" can handle "personal information". "Customer information" can be handled by both "qualified personnel" and "general employees". Some columns of tables that handle "personal information" are allowed to be referred to as "customer information".

This completes the authorization settings in the confidentiality matrix.

Next, we will register the database objects that handle purchase information in the confidentiality matrix.

```

SELECT pgx_add_object_to_confidential_level('matrix_purchase_management',
'level_personal_info',
                                ' [{
                                "type":"table",
                                "object":{
                                        "schema": "purchase",
                                        "table": ["customer_info"]
                                }
                                } ]');
SELECT pgx_add_object_to_confidential_level('matrix_purchase_management',
'level_customer_info',
                                ' [{
                                "type":"column",
                                "object":{
                                        "schema": "purchase",
                                        "table": "customer_info",
                                        "column": ["customer_id", "rank"]
                                }
                                } ]');
SELECT pgx_add_object_to_confidential_level('matrix_purchase_management',
'level_customer_info',
                                ' [{
                                "type":"table",
                                "object":{
                                        "schema": "purchase",
                                        "table": ["history"]
                                }
                                } ]');

```

The entire customer information table is "personal information", the customer\_id column and rank column of the customer information table are "customer information", and the entire purchase history table is also "customer information".

Finally, enroll the employee in the confidentiality group. "Alex" and "Bola" are "qualified persons" who have received training in personal information management. Also, "Charlie" and "Dana" are "general employees" because they have not yet received training on personal information management.

```

SELECT pgx_add_role_to_confidential_group('matrix_purchase_management',
                                'group_qualified',
                                ['Alex', 'Bola']);
SELECT pgx_add_role_to_confidential_group('matrix_purchase_management',
                                'group_non_qualified',
                                ['Charlie', 'Dana']);

```

# Appendix A Tables Used by Confidentiality Management Feature

This section describes the tables used by the confidentiality management feature.

## A.1 pgx\_confidential\_matrix

A list of confidentiality matrices.

You can refer to the attributes of the registered confidentiality matrix, the update time, or the time when the confidentiality level or confidentiality group was registered or deleted.

Column name	Type	Constraint	Description
cmatid	bigint	primary key generated always as identity	Identifier of the confidentiality matrix.
cmatname	varchar(63)	unique not null	Name of the confidentiality matrix.
cmatowner	name	not null	Owner of the confidentiality matrix. The role that created the confidentiality matrix becomes the owner.
cmatcomment	text		Comment.
cmatupdatetime	timestamp with time zone	not null	Update time of the confidentiality matrix itself.
cmatoperationtime	timestamp with time zone		The time when the confidentiality level and confidentiality group were added/ deleted.

## A.2 pgx\_confidential\_level

A list of confidentiality levels.

You can refer to the registered confidentiality level attributes, update time, or the time when a confidentiality object was registered to the confidentiality level or removed from the confidentiality level.

Column name	Type	Constraint	Description
clevid	bigint	primary key generated always as identity	Identifier of the confidentiality level.
clevname	varchar(63)	not null	Name of the confidentiality level.
clevmatid	bigint	not null references pgx_confidential_matr ix(cmatid)	Identifier of the confidentiality matrix to which the confidentiality level belongs.
clevcomment	text		Comment.
clevupdatetime	timestamp with time zone	not null	Update time of the confidentiality level itself.

Column name	Type	Constraint	Description
clevoperationtime	timestamp with time zone		The time when the confidentiality object was added/deleted.
clevencalgorithm	text	not null	Encryption method. "none" for no encryption. AES128 and AES256 can be set.

### A.3 pgx\_confidential\_group

A list of confidentiality groups.

You can refer to the registered confidentiality group attributes, update time, or the time when a role was registered to the confidentiality group or removed from the confidentiality group.

Column name	Type	Constraint	Description
cgroid	bigint	primary key generated always as identity	Identifier of the confidentiality group.
cgroname	varchar(63)	not null	Name of the confidentiality group.
cgromatid	bigint	not null references pgx_confidential_matrix(c matid)	Identifier of the confidentiality matrix to which the confidentiality group belongs.
cgrocomment	text		Comment.
cgroupdatetime	timestamp with time zone	not null	Update time of the confidentiality group itself.
cgrooperationtime	timestamp with time zone		The time when the role was added/deleted.
cgrorolename	name	not null	Name of the confidentiality group role.
cgrosuperuser	bool	not null	true if the confidentiality group role has SUPERUSER privileges.
cgrocreatedb	bool	not null	true if the confidentiality group role has CREATEDB privileges.
cgrocreatorole	bool	not null	true if the confidentiality group role has CREATEROLE privileges.
cgroreplication	bool	not null	true if the confidentiality group role has REPLICATION privileges.
cgrobypassrls	bool	not null	true if the confidentiality group role has BYPASSRLS privileges.

### A.4 pgx\_confidential\_privilege

A list of confidentiality privileges.

You can refer to confidentiality privilege set for each confidentiality object, update time, and so on.

Column name	Type	Constraint	Description
cpriid	bigint	primary key generated always as identity	Identifier of the privilege.

Column name	Type	Constraint	Description
cprimatid	bigint	not null references pgx_confidential_matrix(c matid)	Identifier of the confidentiality matrix to which the privilege belongs.
cprilevid	bigint	not null references pgx_confidential_level(cle vid)	Identifier of the confidentiality level for which privilege is set.
cprigroid	bigint	not null references pgx_confidential_group(cg roid)	Identifier of the confidentiality group for which privilege is set.
cpritype	text	not null	Type of the confidentiality object which privilege is set.
cpriupdateime	timestamp with time zone	not null	Update time when privilege was set/ changed.
cpriacl	text[]	not null	Access privileges that have been set. (*1)

\*1: The character string indicating authority appears in the following order in the text type array of cpriacl.

ALL, SELECT, INSERT, UPDATE, DELETE, TRUNCATE, REFERENCES, TRIGGER, CREATE, CONNECT, TEMPORARY, EXECUTE, USAGE

If the appropriate privilege is not set, the string simply does not appear. For example, {'INSERT','TRUNCATE'} if you only have INSERT and TRUNCATE privileges.

## A.5 pgx\_confidential\_object

A list of confidentiality objects.

You can refer to object attributes or update time, and so on.

Column name	Type	Constraint	Description
cobjid	bigint	primary key generated always as identity	Identifier of the confidentiality object.
cobjmatid	bigint	not null references pgx_confidential_matrix(cma tid)	Identifier of the confidentiality matrix to which the confidentiality object belongs.
cobjlevid	bigint	not null references pgx_confidential_level(clevi d)	Identifier of the confidentiality level to which the confidentiality object belongs.
cobjtype	text	not null	The type of confidentiality object.
cobjschema	name	not null	Schema name of the confidentiality object.
cobjtable	name	not null	Table name of the confidentiality object.
cobjname	text	not null	Name of the confidentiality object.

Column name	Type	Constraint	Description
cobjupdate	timestamp with time zone	not null	Registration time of the confidentiality object.
cobjpolicy	jsonb		Conditions that determine the range of rowsets when the type is rowset.  It is expressed by the setting contents in POLICY.

## A.6 pgx\_confidential\_role

---

A list of roles registered in the confidentiality group.

You can refer to role attributes or update time, and so on.

Column name	Type	Constraint	Description
croolid	bigint	primary key generated always as identity	Identifier of the role.
croimatid	bigint	not null references pgx_confidential_matrix( cmatid)	Identifier of the confidentiality matrix to which the role belongs.
crolgroid	bigint	not null references pgx_confidential_group(c groid)	Identifier of the confidentiality group to which the role belongs.
crolname	name	not null	Name of the role.
croupdate	timestamp with time zone	not null	Registration time of the role.

## A.7 pgx\_confidential\_policy

---

This is a list of policies created to set privileges for confidentiality objects of rowset type. You can refer to the name of the policy you created and the privileges it has set.

Rows in this table are inserted when you add a rowset type confidentiality object.

Column name	Type	Constraint	Description
cpolid	bigint	primary key generated always as identity	Identifier of the policy.
cpolmatid	bigint	not null references pgx_confidential_matrix(c matid)	Identifier of the confidentiality matrix to which the policy belongs.
cpollevid	bigint	not null references pgx_confidential_level(cle vid)	Identifier of the confidentiality level to which the policy belongs.
cpolgroid	bigint	not null references	Identifier of the confidentiality group to which the policy belongs.

Column name	Type	Constraint	Description
		pgx_confidential_group(cgroid)	
cpolobjid	bigint	not null references pgx_confidential_object(cobjid)	Identifier of the rowset object using this policy.
cpolprivilege	text	not null	Privilege this policy has (SELECT, INSERT, UPDATE, DELETE, ALL).
cpolname	name	not null	Name of the policy.
cpolexpression	jsonb	not null	Expression of the policy.

# Appendix B System Management Functions Used by Confidentiality Management Feature

This section describes the system management functions used by the confidentiality management feature. All functions abort the transaction on failure.

## Information

- Be careful when performing operations that involve deleting confidentiality groups.

If you remove a confidentiality group along with a confidentiality group role, you simply no longer have the role that can access the confidentiality object. However, when you leave the confidentiality group role, the function revokes privileges from the confidentiality group role. The privilege to revoke is the privilege defined in confidentiality privileges.

- The `pgx_get_privileges_on_matrix` function may output a very large table if the number of confidentiality objects or roles is large. If the size of this table exceeds the value of PostgreSQL's `work_mem` parameter, I/O will occur according to PostgreSQL's specifications and will be slow. To prevent this, it is recommended that `work_mem` be set as high as possible in the session in which this function is executed.

## B.1 Confidentiality Matrix Manipulation Functions

Function name	Return value	Description
<code>pgx_create_confidential_matrix(confidential_matrix_name varchar, comment text)</code>	void	<p>Create the confidentiality matrix with the specified name.</p> <p>The created confidentiality matrix is registered in the <code>pgx_confidential_matrix</code> table along with the comment.</p> <p>Only roles with the required entitlements for the confidentiality management role can execute this function. The role that executes this function is the confidentiality management role for the confidentiality matrix. Functions such as executing by specifying a confidentiality matrix name require that the executed role is a confidentiality management role for the specified confidentiality matrix. For details, please refer to "<a href="#">8.2.2 Determining Confidentiality Management Roles</a>".</p> <p>The length of <code>confidential_matrix_name</code> must be less than 64 characters. Note that the units are not bytes.</p> <p>There are no restrictions on the characters that can be used in the <code>confidential_matrix_name</code>.</p> <p>When you specify the name of the confidentiality matrix to other functions, you must specify the same string that you specified to this function.</p> <p>Note that unlike most CREATE statements, the name of the confidentiality matrix is case sensitive.</p>
<code>pgx_copy_confidential_matrix(confidential_matrix_name varchar, source_confidential_matrix_name varchar)</code>	void	<p>Copy the source confidentiality matrix specified by <code>source_confidential_matrix_name</code> to the confidentiality matrix named <code>confidential_matrix_name</code>. Confidentiality matrix, confidentiality levels, confidentiality groups, and confidentiality privileges details are replicated. However, the information of confidentiality objects registered in the confidentiality level and roles registered in the confidentiality group are not duplicated.</p>

Function name	Return value	Description
		<p>Any confidentiality management role in any confidentiality matrix can execute this function.</p> <p>The owner of the cloned confidentiality matrix is the role that executed this function.</p> <p>Restrictions on <i>confidential_matrix_name</i> as strings are the same as for the <i>pgx_create_confidential_matrix</i> function.</p> <p>Comments are also duplicated.</p>
<p><i>pgx_alter_confidential_matrix</i>(<i>confidential_matrix_name</i> varchar, <i>alter_object</i> json)</p>	void	<p>Change the attributes of the confidentiality matrix named by <i>confidential_matrix_name</i>.</p> <p>Only confidentiality management role for the specified confidentiality matrix can execute this function.</p> <p>For <i>alter_object</i>, specify the attribute you want to change and the value after change in key-value format as follows. Attributes not specified remain unchanged.</p> <pre>{   "name": "matrix_foo",   "comment": "This matrix is defined for foo." }</pre> <p><i>name</i>: Specify the name of the modified confidentiality matrix. Cannot be <i>null</i>. The function will fail if you specify the name of a confidentiality matrix that already exists.</p> <p><i>comment</i>: Specify a comment after the change. Can be <i>null</i>.</p>
<p><i>pgx_drop_confidential_matrix</i>(<i>confidential_matrix_name</i> varchar, <i>cascade</i> bool, <i>drop_role</i> bool)</p>	void	<p>Drop the confidentiality matrix with the specified name.</p> <p>Only confidentiality management role for the specified confidentiality matrix can execute this function.</p> <p>Specify <i>true</i> for <i>cascade</i> to recursively check and remove objects that depend on this confidentiality matrix. For example, delete the confidentiality groups and confidentiality levels registered in this confidentiality matrix. Then remove the confidentiality privileges associated with that confidentiality level. To drop a confidentiality level, execute internally <i>pgx_drop_confidential_level</i> function with a <i>cascade</i> value. When dropping a confidentiality group, execute internally <i>pgx_drop_confidential_group</i> function with <i>cascade</i> and <i>drop_role</i> values. See also the descriptions of these functions. In particular, how the privilege of confidentiality objects are changed is important.</p> <p>Specify <i>false</i> for <i>cascade</i> simply removes the confidentiality matrix. The function will fail if there are objects that depend on this confidentiality matrix.</p> <p>If <i>true</i> is specified for <i>drop_role</i>, the confidentiality group role registered in this confidentiality matrix will be deleted. Naturally, it only makes sense when <i>cascade</i> is <i>true</i>.</p>

## B.2 Confidentiality Level Manipulation Functions

Function name	Return value	Description
<p><code>pgx_create_confidential_level</code>(confidential_matrix_name varchar, confidential_level_name varchar, options json, comment text)</p>	<p>void</p>	<p>Creates a confidentiality level, registers it with the specified confidentiality matrix, and adds it to the <i>pgx_confidential_level</i> table with the specified comment and attributes specified in <i>options</i>.</p> <p>Only confidentiality management role for the specified confidentiality matrix can execute this function.</p> <p>The length of <i>confidential_level_name</i> must be less than 64 characters.</p> <p>Note that the units are not bytes.</p> <p>There are no restrictions on the characters that can be used in the <i>confidential_level_name</i>.</p> <p>When specifying a confidentiality level name for any other function, you must specify the same string as specified for this function</p> <p>Note that unlike most CREATE statements, confidentiality level names are case-sensitive.</p> <p>Specify a comment for <i>comment</i>.</p> <p>For <i>options</i>, specify the attribute of the confidentiality level as follows. If you specify NULL, the default value for each attribute will be set.</p> <pre>{   "encryption_algorithm": "AES256" }</pre> <p><i>encryption_algorithm</i>: Specify the encryption algorithm. The algorithms and default values that can be specified are the same as the <i>tablespace_encryption_algorithm</i> parameter of Transparent Data Encryption of Fujitsu Enterprise Postgres. Must not be <i>null</i>.</p>
<p><code>pgx_alter_confidential_level</code>(confidential_matrix_name varchar, confidential_level_name varchar, alter_object json)</p>	<p>void</p>	<p>Change the confidentiality level attribute.</p> <p>Only confidentiality management role for the specified confidentiality matrix can execute this function.</p> <p>For <i>alter_object</i>, specify the attribute you want to change and the value after change in key-value format as follows.</p> <pre>{   "name": "level_new",   "comment": "This level is the highest confidentiality level.",   "encryption_algorithm": "AES256" }</pre> <p><i>name</i>: Specify the name of the modified confidentiality level. Cannot be null.</p> <p><i>comment</i>: Specify a comment after the change. Can be null.</p> <p>Other attributes are the same as options of <i>pgx_create_confidential_level</i> function. Attributes not specified are not changed.</p> <p>Be careful when increasing the degree of confidentiality. For example, if you increase the encryption strength and</p>

Function name	Return value	Description
		there are confidentiality objects with a lower strength than the new strength, this function will fail.
pgx_drop_confidential_level(confidential_matrix_name varchar, confidential_level_name varchar, cascade bool)	void	<p>Remove a confidentiality level from the confidentiality matrix and delete a confidentiality level.</p> <p>Only confidentiality management role for the specified confidentiality matrix can execute this function.</p> <p>If <i>true</i> is specified for <i>cascade</i>, the confidentiality level can be deleted even if confidentiality objects are registered in this confidentiality level.</p> <p>If <i>false</i> is specified for <i>cascade</i>, it is not possible to delete a confidentiality level that has confidentiality objects registered.</p>

## B.3 Confidentiality Group Manipulation Functions

Function name	Return value	Description
pgx_create_confidential_group(confidential_matrix_name varchar, confidential_group_name varchar, options json, comment text)	void	<p>Create the confidentiality group, registers it with the specified confidentiality matrix, and adds it to the <i>pgx_confidential_group</i> table with the specified comment and attributes specified in <i>options</i>.</p> <p>Only confidentiality management role for the specified confidentiality matrix can execute this function. However, setting some attributes requires superuser privileges, as described below. Therefore, by necessity, the confidentiality management role must be superuser if those attributes are to be set.</p> <p>This function internally uses the CREATE ROLE statement to create a confidentiality group role.</p> <p>The length of <i>confidential_group_name</i> must be less than 64 characters. Note that the units are not bytes.</p> <p>There are no restrictions on the characters that can be used in the <i>confidential_group_name</i>.</p> <p>When you specify the name of the confidentiality group to other functions, you must specify the same string that you specified to this function.</p> <p>Note that unlike most CREATE statements, confidentiality group names are case-sensitive.</p> <p>For <i>options</i>, specify the attributes of the confidentiality group as follows. If you specify NULL, the default value for each attribute will be set.</p> <pre>' {   "SUPERUSER": false,   "CREATEDB": true,   "CREATEROLE": false,   "REPLICATION": false,   "BYPASSRLS": false } '</pre>

Function name	Return value	Description
		<p>The only attributes that can be specified are the SUPERUSER, CREATEDB, CREATEROLE, REPLICATION, and BYPASSRLS attributes that relate to access privileges to data.</p> <p>These attributes are some of the role attributes that can be specified in the CREATE ROLE statement. The attribute semantics and default values are the same as in the CREATE ROLE statement specification.</p> <p>As noted in the CREATE ROLE statement description, this function fails if a non-superuser specifies <i>true</i> for the SUPERUSER, REPLICATION, and BYPASSRLS attributes.</p>
<p><code>pgx_alter_confidential_group(confidential_matrix_name varchar, confidential_group_name varchar, alter_object json);</code></p>	<p>void</p>	<p>Change the attributes of a confidentiality group.</p> <p>Only confidentiality management role for the specified confidentiality matrix can execute this function. As described in <i>pgx_create_confidential_group</i> function, you must be superuser to set some attributes.</p> <p>This function internally uses the ALTER ROLE statement to change the attributes of the confidentiality group role.</p> <p>For <i>alter_object</i>, specify the attribute you want to change and the value after change in key-value format as follows.</p> <pre>{   "name": "group_new" ,   "comment": "Members of this group have the highest confidential clearance.",   "CREATEDB": false }</pre> <p><i>name</i>: Specify the name of the confidentiality group after modification. Cannot be null.</p> <p><i>comment</i>: Specify a comment after the change. Can be null.</p> <p>Other attributes are the same as options of <i>pgx_create_confidential_group</i> function. Attributes not specified are not changed.</p> <p>For example, if you change the attribute to a weaker one, such as changing CREATEDB to <i>false</i>, the attributes of roles registered in the confidentiality group will be similarly weakened.</p>
<p><code>pgx_drop_confidential_group(confidential_matrix_name varchar, confidential_group_name varchar, cascade bool, drop_role bool)</code></p>	<p>void</p>	<p>Drop the confidentiality group from the confidentiality matrix and delete a confidentiality group.</p> <p>Only confidentiality management role for the specified confidentiality matrix can execute this function.</p> <p>If <i>true</i> is specified for <i>cascade</i>, the confidentiality group can be deleted even if roles are registered in this confidentiality group.</p> <p>If <i>false</i> is specified for <i>cascade</i>, confidentiality groups that have roles registered cannot be deleted.</p> <p>Specify <i>true</i> for <i>drop_role</i> to drop the confidentiality group role. Roles registered in confidentiality groups remain. If</p>

Function name	Return value	Description
		<p>false is specified for <i>drop_role</i>, the confidentiality group role will not be deleted.</p> <p>When you leave the confidentiality group role, the function revokes privileges from the confidentiality group role. The privilege to revoke is the privilege defined in confidentiality privileges.</p>

## B.4 Confidentiality Privilege Manipulation Functions

Function name	Return value	Description
<p><code>pgx_grant_confidential_privilege(confidential_matrix_name varchar, confidential_level_name varchar, confidential_group_name varchar, privilege json)</code></p>	void	<p>Grant confidentiality privileges.</p> <p>Only confidentiality management role for the specified confidentiality matrix can execute this function.</p> <p>Grants access to the confidentiality level specified by <i>confidential_level_name</i> to the confidentiality group specified by <i>confidential_group_name</i>.</p> <p>When you run this function repeatedly, it simply adds more privileges to grant. Granting the same privilege more than once does not result in an error.</p> <p>The privilege to be granted is specified in <i>privilege</i>. <i>privilege</i> specifies the type of the confidentiality object as the key and an array of privileges as the value, like this:</p> <pre>{   "table": ["SELECT", "INSERT", "UPDATE", "DELETE"],   "schema": ["CREATE", "USAGE"],   "rowset": ["ALL"] }</pre> <p>The privileges that can be specified depend on the type of confidentiality object.</p> <p>Privilege for rowset type confidentiality object is privilege that can be specified in the FOR clause of the CREATE POLICY statement.</p> <p>Except for the rowset type, it is a privilege that can be granted with the GRANT statement according to the confidentiality object type.</p> <p>If ALL is specified, it is assumed that all privileges that can be specified for that type are listed.</p> <p>That is, ALL does not appear in the <i>cpriacl</i> column of the <i>pgx_confidential_privilege</i> table.</p> <p>The same specification as the WITH GRANT OPTION clause of the GRANT statement cannot be specified. This is because only confidentiality management roles should use this feature to change privilege to confidentiality objects.</p> <p>Be carefull when PUBLIC is granted to target confidentiality object. This is because granting privileges to PUBLIC is the same as granting privileges to all roles</p>

Function name	Return value	Description
		registered in the confidentiality matrix. This function will fail if a privilege granted indirectly to each role using PUBLIC is defined in the confidentiality privileges that should not be granted to that role. Similarly, this function also checks privileges granted indirectly through group roles that are not registered in the confidentiality matrix. In doing so, it recursively checks the chain of inheritance.
pgx_revoke_confidential_privilege( confidential_matrix_name varchar, confidential_level_name varchar, confidential_group_name varchar, privilege json)	void	<p>Revoke confidentiality privileges.</p> <p>Only confidentiality management role for the specified confidentiality matrix can execute this function.</p> <p>Revokes privilege to the confidentiality level specified by <i>confidential_level_name</i> from the confidentiality group specified by <i>confidential_group_name</i>.</p> <p>Revoking ungranted privileges does not fail.</p> <p>Privilege to be revoked is specified in <i>privilege</i>. The specification method is the same as <i>pgx_grant_confidential_privilege</i> function.</p> <p>Be carefull when PUBLIC is granted to target confidentiality object. This is because granting privileges to PUBLIC is the same as granting privileges to all roles registered in the confidentiality matrix. This function will fail if a privilege granted indirectly to each role using PUBLIC is defined in the confidentiality privileges that should not be granted to that role. Similarly, this function also checks privileges granted indirectly through group roles that are not registered in the confidentiality matrix. This time, it checks the chain of inheritance up to its ancestors.</p>

## B.5 Confidentiality Object Manipulation Functions

Function name	Return value	Description
pgx_add_object_to_confidential_level( confidential_matrix_name varchar, confidential_level_name varchar, object_name json)	void	<p>Only confidentiality management role for the specified confidentiality matrix can execute this function.</p> <p>Adds the confidentiality object specified by <i>object_name</i> to the confidentiality level specified by <i>confidential_level_name</i>.</p> <p>This function internally uses the GRANT statement to grant privileges to the confidentiality group role according to the confidentiality privileges associated with this confidentiality level.</p> <p>However, if the confidentiality object is of type rowset, it internally uses the CREATE POLICY statement to grant privileges to the confidentiality group role according to the confidentiality privileges associated with this confidentiality level. Also, to enable POLICY, execute the ALTER TABLE statement on the target table with the ENABLE ROW LEVEL SECURITY clause.</p>

Function name	Return value	Description
		<p>Currently it is not possible to register a foreign table as a table type confidentiality object.</p> <p>Be carefull when PUBLIC is granted to target confidentiality object. This is because granting privileges to PUBLIC is the same as granting privileges to all roles registered in the confidentiality matrix. This function will fail if a privilege granted indirectly to each role using PUBLIC is defined in the confidentiality privileges that should not be granted to that role. Similarly, this function also checks privileges granted indirectly through group roles that are not registered in the confidentiality matrix. In doing so, it recursively checks the chain of inheritance.</p> <p>For rowset type confidentiality objects, this function will fail if the target table has a POLICY defined that was not created using this feature, regardless of what privileges are granted.</p> <p>These checks only apply to the confidentiality group role or to roles that are registered with the confidentiality group. If a POLICY exists that targets a role that is not, the function will not fail.</p> <p>Specify object_name as follows: Only the rowset is slightly different. The example below attempts to register multiple types of objects in one go.</p> <pre> ' [{   "type": "schema",   "object": [     { "schema": "schema1" },     { "schema": "schema2" }   ] }, {   "type": "table",   "object": [     {       "schema": "schema1",       "table": ["table1", "table2"]     },     {       "schema": "schema2",       "table": ["table8", "table9"]     }   ] }, {   "type": "column",   "object": [     {       "schema": "schema1",       "table": "table1",       "column": ["column1", "column2"]     },     {       "schema": "schema1",       "table": "table2", </pre>

Function name	Return value	Description
		<pre> "column":["column8", "column9"] } ] }]'</pre> <p>For the rowset type, you define rowset and give it a name, as in the example below. This name is used by the <i>pgx_remove_object_from_confidential_level</i> function to identify the rowset type confidentiality object when removing it.</p> <p>This example shows:</p> <ul style="list-style-type: none"> <li>- In <code>schema1.table1</code>, represents a set of rows (rowset) for which the conditional expression (<code>user = current_user</code> OR <code>manger = current_user</code>) is true.</li> <li>- If <code>manager=current_user</code> is true, the <code>coll</code> value of the record to be INSERTed or the record after UPDATE must be greater than zero.</li> </ul> <pre> '[{   "type": "rowset",   "object": [     {       "schema": "schema1",       "table": "table1",       "rowset_name": "rowset1",       "rowset_expression": [         {           "as": "permissive",           "using": "user = current_user"         },         {           "as": "permissive",           "using": "manager = current_user",           "with check": "coll &gt; 0"         }       ]     }   ] }]'</pre> <p>Each key (<i>as</i>, <i>using</i>, <i>with check</i>) has the same meaning as the clause of the same name in the CREATE POLICY statement.</p> <p>As you can see from this, one element of the array specified in <i>rowset_expression</i> corresponds to one POLICY object created by the CREATE POLICY statement.</p> <p>In fact, this function internally executes as many CREATE POLICY statements as there are elements in the array. The name of POLICY at this time is <code>'pgx_cms_policy_\${cpolid}'</code>.</p> <p><code>#{cpolid}</code> is automatically numbered by this extension.</p> <p>&lt;Note&gt;</p>

Function name	Return value	Description
		Do not create policies with names that begin with <i>pgx_cms_policy_</i> . Because this function may fail.
<code>pgx_remove_object_from_confidential_level(confidential_matrix_name varchar, confidential_level_name varchar, object_name json)</code>	void	<p>Removes confidentiality objects from the confidentiality level.</p> <p>Only confidentiality management role for the specified confidentiality matrix can execute this function.</p> <p>Removes the confidentiality object specified by <i>object_name</i> from the confidentiality level specified by <i>confidential_level_name</i>.</p> <p>The format of <i>object_name</i> is the same as <i>object_name</i> in the <i>pgx_add_object_to_confidential_level</i> function. But the <i>rowset_expression</i> key is optional. If specified, it is simply ignored. Therefore, <i>object_name</i> in <i>pgx_add_object_to_confidential_level</i> can be specified to this function without modifying it.</p> <p>At this time, any privileges granted to the confidentiality group based on confidentiality privileges are revoked.</p> <p>If the confidentiality object is of rowset type, delete the internally created policy.</p> <p>At this time, if there are zero policies associated with the table, the ALTER TABLE statement with the DISABLE ROW SECURITY clause is internally executed to disable row level security.</p>

## B.6 Role Manipulation Functions

Function name	Return value	Description
<code>pgx_add_role_to_confidential_group(confidential_matrix_name varchar, confidential_group_name varchar, role_name json)</code>	void	<p>Add a role to a confidentiality group.</p> <p>Only confidentiality management role for the specified confidentiality matrix can execute this function.</p> <p>Adds the role specified by <i>role_name</i> to the confidentiality group specified by <i>confidential_group_name</i>.</p> <p>If the role to be added has been granted broader privileges than the confidentiality privileges, revoke the privileges according to the confidentiality privileges.</p> <p>Be carefull when PUBLIC is granted to target confidentiality object. This is because granting privileges to PUBLIC is the same as granting privileges to all roles registered in the confidentiality matrix. This function will fail if a privilege granted indirectly to each role using PUBLIC is defined in the confidentiality privileges that should not be granted to that role. Similarly, this function also checks privileges granted indirectly through group roles that are not registered in the confidentiality matrix. In doing so, it recursively checks the chain of inheritance.</p>

Function name	Return value	Description
		<p>Also, if the added role has stronger attributes than the confidentiality group, change the attributes to match the confidentiality group.</p> <p>This function will fail if a strong attribute is indirectly assigned using a group role that is not registered in the confidentiality matrix.</p> <p><i>role_name</i> is specified as follows.</p> <pre>'{"role1","role"}'</pre>
<p><code>pgx_remove_role_from_confidential_group(confidential_matrix_name varchar, confidential_group_name varchar, role_name json)</code></p>	void	<p>Remove a role from a confidentiality group.</p> <p>Only confidentiality management role for the specified confidentiality matrix can execute this function.</p> <p>Removes the role specified by <i>role_name</i> from the confidentiality group specified by <i>confidential_group_name</i>.</p> <p>This function internally executes a REVOKE statement to remove the role from the confidentiality group role.</p> <p>It does not change the attributes of the role being removed, nor the privileges granted to that role.</p> <p>Simply banish it from the group so that it cannot inherit privileges.</p> <p>The method of specifying <i>role_name</i> is the same as <code>pgx_add_role_to_confidential_group</code> function.</p>

## B.7 Functions that Support Definition Referencing and Comparison with System Catalogs

Functions that Support Definition Referencing and Comparison with System Catalogs

Function name	Return value	Description
<p><code>pgx_get_attribute_of_objects(confidential_matrix_name varchar)</code></p>	setof record	<p>Returns a table of attributes defined in the confidentiality matrix and attributes actually set in the database for all confidentiality objects registered in the specified confidentiality matrix. Refer to "<a href="#">Tables returned by pgx_get_attribute_of_objects</a>" for the table format.</p> <p>Only confidentiality management role for the specified confidentiality matrix can execute this function.</p>
<p><code>pgx_get_attribute_of_roles(confidential_matrix_name varchar)</code></p>	setof record	<p>Returns a table of attributes defined in the confidentiality matrix and attributes actually set in the system catalog for all roles registered in the specified confidentiality matrix. Refer to "<a href="#">Table returned by pgx_get_attribute_of_roles</a>" for the format of the table.</p> <p>Only confidentiality management role for the specified confidentiality matrix can execute this function.</p>
<p><code>pgx_get_privileges_on_level_and_group(confidential_matrix_name varchar, confidential_level_name)</code></p>	setof record	<p>Returns a table that allows you to compare the following for combinations of confidentiality objects registered in the specified confidentiality level and roles registered in the</p>

Function name	Return value	Description
varchar, confidential_group_name varchar)		<p>specified confidentiality group. Refer to "<a href="#">Table returned by pgx_get_privileges_on_level_and_group</a>" for the format of the table.</p> <ul style="list-style-type: none"> <li>- Privileges specified by the confidentiality privilege settings that should be granted to the specified role</li> <li>- Privileges granted in the actual system catalog</li> </ul> <p>Only confidentiality management role for the specified confidentiality matrix can execute this function.</p>
pgx_get_privileges_on_object(confidential_matrix_name varchar, object_name json)	setof record	<p>Returns a table that allows you to compare the following for the specified confidentiality objects. Refer to "<a href="#">Table returned by pgx_get_privileges_on_object</a>" for the format of the table.</p> <ul style="list-style-type: none"> <li>- Privileges dictated by confidentiality privilege settings that should be granted to all roles</li> <li>- Privileges granted in the actual system catalog</li> </ul> <p>Only confidentiality management role for the specified confidentiality matrix can execute this function.</p>
pgx_get_privileges_on_role(confidential_matrix_name varchar, role_name json)	setof record	<p>Returns a table that allows you to compare the following for the all confidentiality objects. Refer to "<a href="#">Table returned by pgx_get_privileges_on_role</a>" for the format of the table.</p> <ul style="list-style-type: none"> <li>- Privileges specified by the confidentiality privilege settings that should be granted to the specified role</li> <li>- Privileges granted in the actual system catalog</li> </ul> <p>Only confidentiality management role for the specified confidentiality matrix can execute this function.</p>
pgx_get_privileges_on_matrix(confidential_matrix_name varchar)	setof record	<p>Returns a table that allows you to compare the following for the all confidentiality objects in the specified confidentiality matrix. Refer to "<a href="#">Table returned by pgx_get_privileges_on_matrix</a>" for the format of the table.</p> <ul style="list-style-type: none"> <li>- Privileges defined by confidentiality privilege settings that should be granted to all roles registered in the confidentiality matrix.</li> <li>- Privileges granted in the actual system catalog</li> </ul> <p>Only confidentiality management role for the specified confidentiality matrix can execute this function.</p>

#### Tables returned by pgx\_get\_attribute\_of\_objects

Column name	Type	Description
matrix_name	varchar(63)	Confidentiality matrix name
confidential_level_name	varchar(63)	Confidentiality level name
object_type	text	Confidentiality object type
object_schema	name	Confidentiality object schema name

Column name	Type	Description
object_table	name	Confidentiality object table name
object_name	text	Confidentiality object name
rowset_expression	json	Conditional expression when the confidentiality object is a row
encrypt_on_matrix	text	Encryption method and strength specified in the confidentiality matrix
encrypt_on_object	text	Actual encryption method and strength of confidentiality objects

#### Tables returned by pgx\_get\_attribute\_of\_roles

Column name	Type	Description
matrix_name	varchar(63)	Confidentiality matrix name
confidential_group_name	varchar(63)	Confidentiality group name
role_name	name	Role name
confidential_group_role	bool	Indicates whether it is a confidentiality group role or not. <i>true</i> if it is a confidentiality group role
superuser_on_matrix	bool	SUPERUSER attribute specified in the confidentiality matrix
superuser_on_role	bool	Actual SUPERUSER attribute of the role
createdb_on_matrix	bool	CREATEDB attribute specified in the confidentiality matrix
createdb_on_role	bool	Actual CREATEDB attribute of the role
createrole_on_matrix	bool	CREATEROLE attribute specified in the confidentiality matrix
createrole_on_role	bool	Actual CREATEROLE attribute of the role
replication_on_matrix	bool	REPLICATION attribute specified in the confidentiality matrix
replication_on_role	bool	Actual REPLICATION attribute of the role
bypassrls_on_matrix	bool	BYPASSRLS attribute specified in the confidentiality matrix
bypassrls_on_role	bool	Actual BYPASSRLS attribute of the role

#### Tables returned by pgx\_get\_privileges\_on\_level\_and\_group, pgx\_get\_privileges\_on\_object, pgx\_get\_privileges\_on\_role and pgx\_get\_privileges\_on\_matrix

Column name	Type	Description
matrix_name	varchar(63)	Confidentiality matrix name
confidential_level_name	varchar(63)	Confidentiality level name

Column name	Type	Description
confidential_group_name	varchar(63)	Confidentiality group name
object_type	text	Confidentiality object type
object_scheme	name	Confidentiality object schema name
object_table	name	Confidentiality object table name
object_name	text	Confidentiality object name
role_name	name	Role name
privilege_list_on_matrix	text[]	Privileges specified by the confidentiality matrix settings. Output is separated by commas
privilege_list_on_object	text[]	Privileges specified by the confidentiality matrix settings. Output is separated by commas
policy_name	name	NULL if the confidentiality object is not of type rowset  In the case of rowset type, rowset name (*1)
policy_setting_on_matrix	jsonb	NULL if the confidentiality object is not of type rowset  In the case of rowset type, the rowset policy information set in the confidentiality matrix (*1)
policy_setting_on_policy	jsonb	NULL if the confidentiality object is not of type rowset  In case of rowset type, row policy information set in the actual policy (*1)

\*1: When adding a rowset type confidentiality object, multiple privileges can be set at once, which is not represented by a single row in this table. For example, if you set SELECT and DELETE privileges, you will see a row for SELECT privileges and a row for DELETE privilege. This is because rowset type access control uses PostgreSQL's row-level security POLICY. In this specification, POLICY for SELECT privilege is different from POLICY for DELETE privilege.