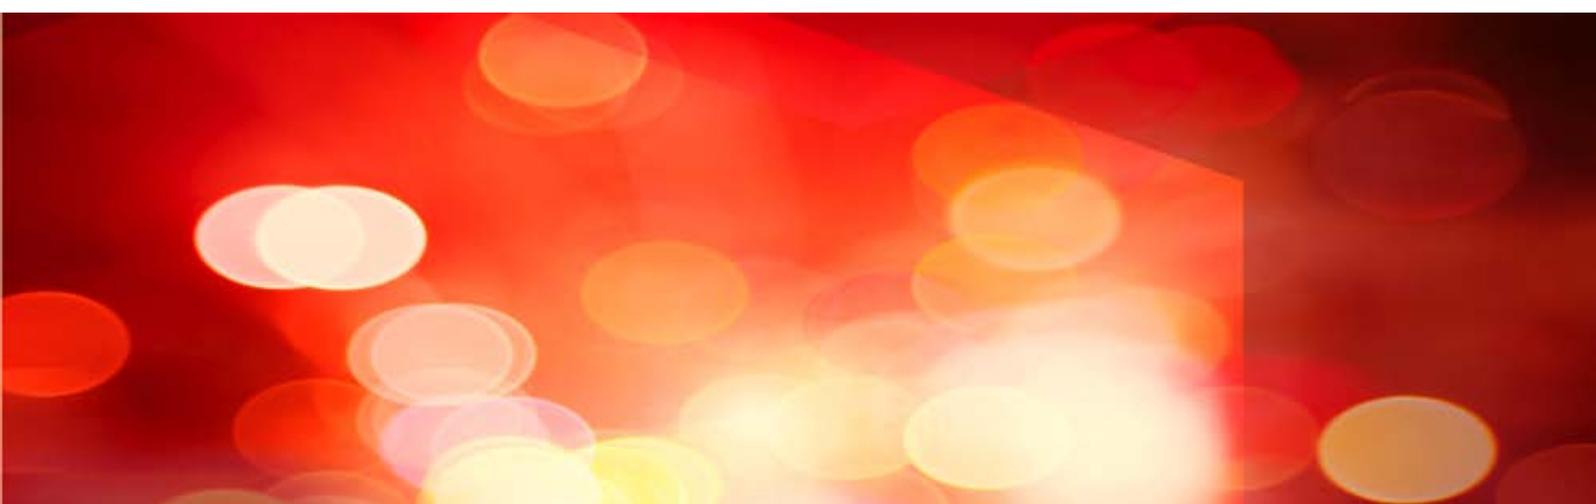


FUJITSU Enterprise Postgres 9.4



Installation and Setup Guide for Server

Linux

J2UL-2035-01ENZ0(00)
July 2015

Preface

Purpose of This Document

The Enterprise Postgres database system extends the PostgreSQL features and runs on the Linux platform.

This document describes how to install and setup "Enterprise Postgres".

Intended Readers

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

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

- PostgreSQL
- SQL
- Linux

Structure of This Document

This document is structured as follows:

[Chapter 1 Overview of Installation](#)

Describes the installation types and procedures

[Chapter 2 Operating Environment](#)

Describes the operating environment required to use Enterprise Postgres

[Chapter 3 Installation](#)

Describes how to perform a new installation of Enterprise Postgres

[Chapter 4 Setup](#)

Describes the setup to be performed after installation

[Chapter 5 Uninstallation](#)

Describes how to uninstall Enterprise Postgres

[Appendix A Installation in Silent Mode](#)

Provides specifications for installation in silent mode

[Appendix B Enterprise Postgres Migration Procedures](#)

Describes how to migrate Enterprise Postgres from an old version/level.

[Appendix C Using WebAdmin in a Multi-version Installation Environment](#)

Describes the required operations for using WebAdmin in multiple Enterprise Postgres versions.

[Appendix D Recommended WebAdmin Environments](#)

Describes the recommended WebAdmin environment.

[Appendix E Setting Up and Removing WebAdmin](#)

Describes how to set up and remove WebAdmin

[Appendix F Configuring Parameters](#)

Describes Enterprise Postgres parameters.

[Appendix H Uninstall \(middleware\)](#)

Describes the Uninstall (middleware)

[Appendix G Uninstall \(middleware\) Messages](#)

Explains the messages output by the Uninstall (middleware) tool.

[Appendix I Estimating Database Disk Space Requirements](#)

Describes how to estimate database disk space requirements

[Appendix J Estimating Memory Requirements](#)

Describes the formulas for estimating memory requirements

[Appendix K Quantitative Limits](#)

Describes the quantity range

[Appendix L Configuring Kernel Parameters](#)

Describes the settings for kernel parameters

Export Restrictions

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

Issue Date and Version

First Edition : July, 2015

Copyright Notice

Copyright 2015 FUJITSU LIMITED

Contents

Chapter 1 Overview of Installation.....	1
1.1 Features that can be Installed.....	1
1.2 Installation Types.....	1
1.2.1 New Installation.....	1
1.2.2 Reinstallation.....	1
1.2.3 Upgrade Installation.....	1
1.2.4 Multi-Version Installation.....	1
1.3 Installation Procedure.....	1
1.3.1 Installation in Interactive Mode.....	2
1.3.2 Installation in Silent Mode.....	2
1.4 Uninstallation.....	2
Chapter 2 Operating Environment.....	3
2.1 Required Operating System.....	3
2.2 Related Software.....	4
2.3 Excluded Software.....	4
2.4 Required Patches.....	4
2.5 Hardware Environment.....	5
2.6 Disk Space Required for Installation.....	5
2.7 Supported System Environment.....	5
2.7.1 TCP/IP Protocol.....	5
2.7.2 File System.....	5
2.8 Enterprise Postgres-based PostgreSQL Versions.....	6
2.9 Notes on Using Streaming Replication.....	6
Chapter 3 Installation.....	7
3.1 Pre-installation Tasks.....	7
3.2 Installation in Interactive Mode.....	8
3.3 Installation in Silent Mode.....	12
3.4 Post-installation Tasks.....	13
Chapter 4 Setup.....	14
4.1 Operating Method Types and Selection.....	14
4.2 Preparations for Setup.....	15
4.2.1 Creating an Instance Administrator.....	15
4.2.2 Preparing Directories for Resource Deployment.....	15
4.2.3 Estimating Resources.....	18
4.2.4 Configuring Corefile Names.....	18
4.3 Creating an Instance.....	18
4.3.1 Using WebAdmin.....	19
4.3.1.1 Logging into WebAdmin.....	19
4.3.1.2 Creating an Instance.....	20
4.3.1.3 Changing the Settings.....	25
4.3.2 Using the initdb Command.....	32
4.3.2.1 Editing Kernel Parameters.....	32
4.3.2.2 Creating an Instance.....	32
4.4 Configuring Remote Connections.....	35
4.4.1 When an Instance was Created with WebAdmin.....	35
4.4.2 When an Instance was Created with the initdb Command.....	35
4.5 Other Settings.....	36
4.5.1 Error Log Settings.....	36
4.5.2 Configuring Automatic Start and Stop of an Instance.....	36
4.5.3 Settings when Using the features compatible with Oracle databases.....	38
4.6 Integration with Message-Monitoring Software.....	38
4.7 Deleting Instances.....	39

4.7.1 Using WebAdmin.....	39
4.7.2 Using Server Commands.....	40
Chapter 5 Uninstallation.....	41
5.1 Uninstallation in Interactive Mode.....	41
5.2 Uninstallation in Silent Mode.....	43
Appendix A Installation in Silent Mode.....	46
A.1 Specification Format.....	46
A.2 List of Parameters.....	46
A.3 Messages and Return Values.....	49
A.4 CSV File Format.....	50
Appendix B Enterprise Postgres Migration Procedures.....	52
B.1 Migration Tasks for Basic Operations.....	52
B.2 Migration Procedures when Performing Failover Operations using the Standby Feature.....	52
Appendix C Using WebAdmin in a Multi-version Installation Environment.....	55
Appendix D Recommended WebAdmin Environments.....	57
D.1 Recommended Browser Settings	57
D.2 How to Set Up the Pop-up Blocker.....	57
Appendix E Setting Up and Removing WebAdmin.....	58
E.1 Setting Up WebAdmin.....	58
E.1.1 Setting Up WebAdmin.....	58
E.1.2 Activating the Web Server Feature of WebAdmin.....	59
E.1.3 Stopping the Web Server Feature of WebAdmin.....	59
E.2 Removing WebAdmin.....	60
Appendix F Configuring Parameters.....	61
Appendix G Uninstall (middleware) Messages.....	64
G.1 Messages output by FJSVcir.....	64
G.2 Messages output by Uninstall (middleware).....	67
Appendix H Uninstall (middleware).....	70
H.1 Features that are Installed.....	70
H.2 Notes.....	70
H.2.1 Notes on the Uninstall (middleware) Tool.....	70
Appendix I Estimating Database Disk Space Requirements.....	72
I.1 Estimating Table Size Requirements.....	72
I.2 Estimating Index Size Requirements.....	73
I.3 Sizes of Data Types.....	74
I.3.1 Sizes of Fixed-Length Data Types.....	74
I.3.2 Sizes of Variable-Length Data Types.....	74
I.3.3 Sizes of Array Data Types.....	75
I.3.4 Number of Bytes per Character.....	75
I.4 Estimating Transaction Log Space Requirements.....	76
I.5 Estimating Archive Log Space Requirements.....	76
I.6 Estimating Backup Disk Space Requirements.....	76
Appendix J Estimating Memory Requirements.....	77
J.1 Enterprise Postgres Memory Requirements.....	77
J.2 Database Multiplexing Memory Requirements.....	78
J.3 PL/extJava Memory Requirements.....	78
Appendix K Quantitative Limits.....	79
Appendix L Configuring Kernel Parameters.....	84

Chapter 1 Overview of Installation

This chapter provides an overview of Enterprise Postgres installation.

1.1 Features that can be Installed

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

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

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

Y: Can be installed

N: Cannot be installed

1.2 Installation Types

The following four installation types are available for Enterprise Postgres:

- New installation
- Reinstallation

1.2.1 New Installation

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

1.2.2 Reinstallation

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

1.2.3 Upgrade Installation

Perform upgrade installation to upgrade installed program files to the latest version.

1.2.4 Multi-Version Installation

Perform multi-version installation to install different versions to the installed program files separately.



When performing multi-version installation to reinstall a version already used in the past, caution regarding WebAdmin is advised. Refer to "[Appendix C Using WebAdmin in a Multi-version Installation Environment](#)" for details.

1.3 Installation Procedure

The following installation procedures are available for Enterprise Postgres:

- Installation in interactive mode
- Installation in silent mode

Select the installation procedure that corresponds to your environment.

1.3.1 Installation in Interactive Mode

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

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

- New installation
- Reinstallation

1.3.2 Installation in Silent Mode

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

New installations can be performed in silent mode.

1.4 Uninstallation

Uninstallation removes the system files of the installed Enterprise Postgres.

Chapter 2 Operating Environment

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



See

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

2.1 Required Operating System

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

- RHEL6 (Intel 64)
- RHEL7



Note

The SELinux (Security-Enhanced Linux) feature is not supported for RHEL6.



Information

Select the x86_64 architecture package when installing the 64-bit product.

- The following packages are required for operations on RHEL6 (Intel64).

Package name	Remarks
alsa-lib	-
audit-libs	-
glibc	-
libgcc	-
libstdc++	-
libtool-ltdl	Required when using ODBC drivers.
ncurses-libs	-
nss-softokn-freebl	-
pam	Required when using PAM authentication.
python-libs	Required when using PL/Python
tcl	Required when using PL/Tcl
unzip	-
xz-libs	-
zlib	-

- The following packages are required for operations on RHEL7.

Package name	Remarks
alsa-lib	-
audit-libs	-

Package name	Remarks
glibc	-
libgcc	-
libstdc++	-
libtool-ldt	Required when using ODBC drivers.
ncurses-libs	-
nss-softokn-freebl	-
pam	Required when using PAM authentication.
perl	Install v5.10.1 if using PL/Perl.
python-libs	Required when using PL/Python Install 2.6.6.
tcl	Required when using PL/Tcl
unzip	-
xz-libs	-
Zlib	-

2.2 Related Software

The following software required to use Enterprise Postgres.

- NetCOBOL Standard Edition V7.0L10 or later

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

Table 2.1 Connectable servers

OS	Product name
Windows Linux	FUJITSU Software Enterprise Postgres Standard Edition 9.4 or later

2.3 Excluded Software

This section describes excluded software.

Enterprise Postgres

If all the following conditions are met, Enterprise Postgres cannot be installed:

- The product generations are the same
- The editions are different

Other products

There are no exclusive products.

2.4 Required Patches

There are no required patches.

2.5 Hardware Environment

The following hardware is required to use Enterprise Postgres.

Memory

At least 512 MB of memory is required.

2.6 Disk Space Required for Installation

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

Table 2.2 Disk space required for installation

Directory	Required disk space (Unit: MB)
/etc	1 + 1(*1)
/var	1 + 4(*1) + 1(*2)
/opt	85(*1) + 1(*2)
Installation destination of the server	600
Installation destination of the client (32-bit)	95
Installation destination of the client (64-bit)	100

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

*2: Required for the installation of FJQSS.

2.7 Supported System Environment

This section describes the supported system environment.

2.7.1 TCP/IP Protocol

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



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

2.7.2 File System

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

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

The recommended file system is "ext4".



In terms of software reliability, ext3 is preferable to ext4 on RHEL 6.4 and earlier.

2.8 Enterprise Postgres-based PostgreSQL Versions

The Enterprise Postgres-based PostgreSQL version is PostgreSQL 9.4.3.

2.9 Notes on Using Streaming Replication

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

*1: The product version is indicated by "x" in the notation "x.y.z".



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

Chapter 3 Installation

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

3.1 Pre-installation Tasks

Check the system environment below before installing Enterprise Postgres.

Check the disk space

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

Refer to "[2.6 Disk Space Required for Installation](#)" for information on the required disk space.

Reconfigure the disk partition if disk space is insufficient.

Check the installed product and determine the installation method

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

Example

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

Currently installed products
1. Enterprise Postgres Client(64bit) 9.4
2. Enterprise Postgres Client(32bit) 9.4
3. Enterprise Postgres Standard Edition(64bit) 9.4

Type [number] to select the software you want to uninstall.
[number,q]
=>q

Exiting Uninstaller.
```

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

- Reinstallation

Remove applied updates

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



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

- Performing reinstallation

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

Perform reinstallation after removing the update.

1. Display the applied updates

Execute the following command to display the applied updates:

Example

```
# /opt/FJSVfupde/bin/uam showup
2015-05-21T18:50 T009843LP-01 S A - O - - - - (RHEL6)Enterprise Postgres
[product-name Enterprise Postgres Client 9.4]
2015-05-21T18:53 T010230LP-01 S A - O - - - - (RHEL6)Enterprise Postgres Client
[product-name Enterprise Postgres Client 9.4]
2015-05-21T18:52 T009844LP-01 S A - O - - - - (RHEL6) Enterprise Postgres Client
```

2. Remove the updates

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

Example

```
# /opt/FJSVfupde/bin/uam remove -i T009843LP-01
2015-05-21T18:50 T009843LP-01 S A - O - - - - (RHEL6)Enterprise Postgres
Are you sure you want to restore the updated product to its pre-update state? (Y/N)y
Restoration to the pre-update state is completed.
The update application management ledger has been updated.
Starting from now, the updated update application management ledger is used to research and
apply updates.
```

Prepare when performing migration

To perform migration from an old version/level, refer to "[Appendix B Enterprise Postgres Migration Procedures](#)", check the procedure, and then start the migration tasks.

3.2 Installation in Interactive Mode

Install according to the following procedure:



The following characters can be used as input values:

Alphanumeric characters, hyphens and forward slashes

1. Stop applications and the client program

If the installation method is one of the following, the applications and the program must be stopped:

- Reinstallation

Before starting the installation, stop the following:

- Applications that use the product
- Instance

Using WebAdmin

In the [Monitor] menu window, click [stop].

Using server commands

Execute the pg_ctl command with the stop mode option specified.

Example

```
$ /opt/fsepserver64/bin/pg_ctl stop -D /database/inst1
```

- Web server feature of WebAdmin

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

Example

```
# cd /opt/fsepserver64/gui/sbin
# ./WebAdminStop
```

In the example above, /opt/fsepserver64 is the installation directory.

2. Change to the superuser

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

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

3. Mount the DVD drive

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

Example

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

/dev/dvd is the DVD drive device, and /media/dvd is the mount point (which must already exist before calling the command).



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

4. Run the installation

Execute the following command:

Example

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

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

5. Select the product for installation

The list of installation target products is displayed as shown below.

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

```
The following products can be installed:
1: Enterprise Postgres Standard Edition (64bit) 9.4
2: Enterprise Postgres Client (32bit) 9.4
3: Enterprise Postgres Client (64bit) 9.4

Select the product to be installed.
To select multiple products, separate using commas (,). (Example: 1,2)
[number,all,q](The default value is all): all

Selected product
  Enterprise Postgres Standard Edition (64bit) 9.4
  Enterprise Postgres Client (32bit) 9.4
  Enterprise Postgres Client (64bit) 9.4

Do you want to install the above product?
y: Proceed to the next step
n: Select the product again
```

```
q: Quit without installing
[y,n,q](The default value is y): y
```

Information

- To develop or execute a 32-bit application in a 64-bit environment, Enterprise Postgres Client 32bit is required.
- The window below is displayed for each product that that was selected but has already been installed.

```
Select the method to install the server product.
If performing an upgrade installation (reinstallation), select the number of the product to be
updated. (*1)

1: Enterprise Postgres Standard Edition (64bit) 9.4
q: Quit without installing
[number,q]: 1
```

*1: Output if the installed product can undergo installation or reinstallation.

6. Check the settings

The window for checking the installation information is displayed as shown below.

If there is no problem with the settings, type "y" and press Enter to start the installation (refer to "[10. Display the installation status](#)").

To change the settings, type "c" and press Enter.

```
Product to be installed
  Enterprise Postgres Standard Edition (64bit) 9.4
    New installation
  Enterprise Postgres Client (32bit) 9.4
    New installation
  Enterprise Postgres Client (64bit) 9.4
    New installation

Installation directory information
  Enterprise Postgres Standard Edition (64bit) 9.4
    /opt/fsepserver64
  Enterprise Postgres Client (32bit) 9.4
    /opt/fsepclient32
  Enterprise Postgres Client (64bit) 9.4
    /opt/fsepclient64

Setup information
  WebAdmin setup: Execute
  Web server port number: 27515
  WebAdmin internal port number: 27516

Start installation using the above information?
y: Start the installation
c: Change the information          (*1)
q: Quit without installing
[y,c,q](The default value is y): c          (*1)
```

*1: This option is not displayed if there is no information that can be modified.

If the user types "c", then the window mentioned in the next step is displayed.

7. Enter the installation directory

If the installation method is one of the following, the window below is displayed for each product:

- New installation

Enter the directory and press Enter.

```
Enterprise Postgres Standard Edition (64bit) 9.4 : Specify the installation directory.
[directory name,q](The default value is /opt/fsepsv64): /opt/fsepsv64
```

8. Enter the WebAdmin setup information

If a server product was selected, and the installation method is one of the following, the window below will be displayed:

- New installation

Type the item name and press Enter.

```
Do you want to execute WebAdmin setup?
y: Execute
n: Do not execute
[y,n,q](The default value is y):

Specify the Web server port number.
Web server port number [1024-32767,q](The default value is 27515):

Specify the WebAdmin internal port number.
WebAdmin internal port number [1024-32767,q](The default value is 27516):
```

If you have not set up WebAdmin, refer to "[Appendix E Setting Up and Removing WebAdmin](#)" for details.

9. Check the changed settings

The changed settings are displayed as shown below.

If there is no problem with the settings, type "y" and press Enter.

To change the settings again, type "c" and press Enter.

```
Product to be installed
  Enterprise Postgres Standard Edition (64bit) 9.4
    New installation
  Enterprise Postgres Client (32bit) 9.4
    New installation
  Enterprise Postgres Client (64bit) 9.4
    New installation

Installation directory information
  Enterprise Postgres Standard Edition (64bit) 9.4
    /opt/fsepsv64
  Enterprise Postgres Client (32bit) 9.4
    /opt/fsepcl32
  Enterprise Postgres Client (64bit) 9.4
    /opt/fsepcl64

Setup information
  WebAdmin setup: Execute
  Web server port number: 27515
  WebAdmin internal port number: 27516

Start installation using the above information?
y: Start the installation
c: Change the information
q: Quit without installing
[y,c,q](The default value is y): y
```

10. Display the installation status

The installation status is displayed as follows:

```
Starting installation.

Enterprise Postgres Standard Edition (64bit) 9.4 Installation
```

```
Installation is complete.

Enterprise Postgres Client (32bit) 9.4 Installation

Installation is complete.

Enterprise Postgres Client (64bit) 9.4 Installation

Installation is complete.

Starting setup.          (*1)

Sets up WebAdmin.       (*1)

Setup is complete.      (*1)
```

*1: Displayed only when "Execute" is selected for WebAdmin setup.
Even if WebAdmin setup fails, the process will still continue.
If the following message is displayed, manually execute WebAdmin setup after installation is completed.
Refer to "[E.1 Setting Up WebAdmin](#)" for information on the WebAdmin setup procedure.

```
Starting setup.

Setting up WebAdmin.

ERROR: WebAdmin setup failed.
```

11. Finish the installation

Upon completion, a message is displayed showing the status.
If installation was successful, a message like the one shown below is displayed:

```
Installed successfully.
```



If an error occurs during the installation, read the error message and remove the cause of the error, and then reexecute the `install.sh` command.

3.3 Installation in Silent Mode

Installation in silent mode can be performed only when the installation method is one of the following:

- New installation

The installation procedure is described below.

1. Change to the superuser

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

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

2. Mount the DVD drive

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

Example

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

/dev/dvd is the DVD drive device, and /media/dvd is the mount point (which must already exist before calling the command).

Note

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

3. Create an installation parameters CSV file

Consider the features that will be required for system operations, and then create an installation parameters CSV file that uses the following specification format.

```
sectionName, parameterName, value  
sectionName, parameterName, value  
:
```

Refer to "[Appendix A Installation in Silent Mode](#)" for information on installation parameters CSV files.

Information

The template for the installation parameters CSV file is "mountpoint/sample/sample.csv".

4. Run the installation

Execute the following command:

Example

```
# cd /media/dvd  
# ./silent.sh /home/work/inspara.csv
```

In the example above, /media/dvd is the DVD mount point, and /home/work/inspara.csv is the installation parameter CSV.

If the installer ends in an error, a message is output to the log file and return values are returned. Refer to "[Appendix A Installation in Silent Mode](#)" for details.

3.4 Post-installation Tasks

There are no post-installation tasks.

Chapter 4 Setup

This chapter describes the setup procedures to be performed after installation completes.

4.1 Operating Method Types and Selection

This section describes how to operate Enterprise Postgres.

There are two methods of managing Enterprise Postgres operations - select one that suits your purposes:

The Operation Guide describes the operating method using WebAdmin, and the equivalent operating method using the server commands.

Simple operation management using a web-based GUI tool (WebAdmin)

Suitable when using frequently used basic settings and operations for operation management.

This method allows you to perform simple daily tasks such as starting the system before beginning business, and stopping the system when business is over, using an intuitive operation.

Usage method

Usage is started by using WebAdmin to create the instance.

By using an external scheduler and the `pgx_dmpall` command, periodic backups can be performed, which can then be used in recovery using WebAdmin.



- Do not use a server command other than `pgx_dmpall` or a server application. Operation modes that use server commands and server applications cannot be used in conjunction with WebAdmin. If used, WebAdmin will not be able to manage the instances correctly.
Refer to Reference and the PostgreSQL Documentation for information on server commands and server applications.
- An instance that was created using WebAdmin can be set up to operate database multiplexing, however once the setup is complete, it will no longer be possible to select that instance from the WebAdmin management window. Refer to the Cluster Operation Guide for information on how to perform database multiplexing operations.

Advanced operation management using server commands

When operating in a system that is automated by operation management middleware (Systemwalker Centric Manager, for example), this method allows you to use more detailed settings and operations and perform higher level operation management.

An overview of the operating method using the GUI, and its relationship with the operating method using the server commands, are shown below.

Refer to the Operation Guide for details.

Operation		Operation with the GUI	Operation with commands
Setup	Creating an instance	WebAdmin is used. The server machine capacity, and the optimum parameter for operations using WebAdmin, are set automatically.	The configuration file is edited directly using the <code>initdb</code> command.
	Changing the configuration files	WebAdmin is used.	The configuration file is edited directly.
Activating and stopping an instance		WebAdmin is used.	The <code>pg_ctl</code> command is used.
Creating a database		This is defined using pgAdmin of the GUI tool, or using the <code>psql</code> command or the application after specifying the DDL statement.	

Operation		Operation with the GUI	Operation with commands
Backing up the database		WebAdmin, or the <code>pgx_dmpall</code> command, is used.	It is recommended that the <code>pgx_dmpall</code> command be used. Recovery to the latest database can be performed.
Database recovery		WebAdmin is used.	To use the backup that was performed using the <code>pgx_dmpall</code> command, the <code>pgx_rcvall</code> command is used.
Monitoring	Database errors	The status in the WebAdmin window can be checked. (*1)	The messages that are output to the database server log are monitored (*1)
	Disk space	The status in the WebAdmin window can be checked. A warning will be displayed if the free space falls below 20%. (*1)	This is monitored using the <code>df</code> command of the operating system, for example. (*1)
	Connection status	This can be checked using pgAdmin of the GUI tool, or referencing <code>pg_stat_activity</code> of the standard statistics view from <code>psql</code> or the application.	

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



See

Refer to "Periodic Operations" and "Actions when an Error Occurs" in the Operation Guide for information on monitoring and database recovery.

4.2 Preparations for Setup

This section describes the preparation required before setting up Enterprise Postgres.

4.2.1 Creating an Instance Administrator

Decide which OS user account will be assigned the instance administrator role. You can assign it to a new user or to an existing one, but you cannot assign it to the OS superuser (root).

The following example shows an OS user account with the name "fsepuser" being assigned the instance administrator role.

Example

```
# useradd fsepuser
# passwd fsepuser
```

4.2.2 Preparing Directories for Resource Deployment

Prepare the directories required when creating instances.

Considerations when deploying resources

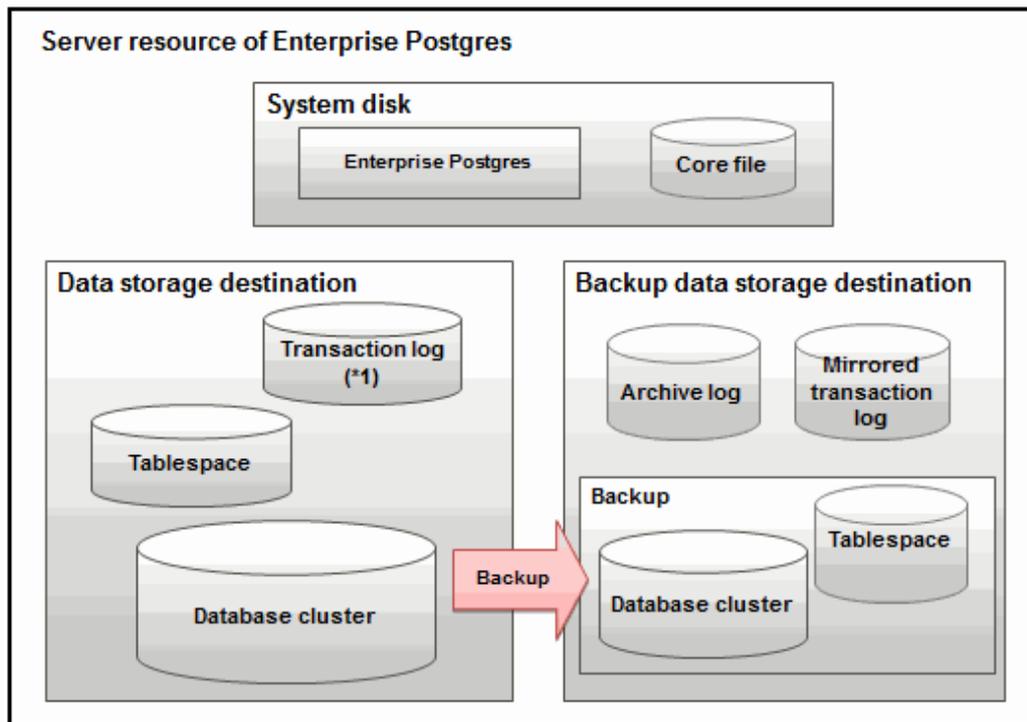
The disk configuration on the resource deployment destination is important, because it affects not only recovery following disk corruption, but normal operation as well. The points for determining the disk configuration are as follows:

1. If the backup data storage destination and the data storage destination are both lost, it will not be possible to recover the data, so deploy them to separate disks.
2. To shorten the recovery time following a single disk fault, deploy the system disk and data storage destination to separate disks.

3. The backup data storage destination requires at least double the capacity of the data storage destination, so deploy it to the disk with the most space available.
4. When large amounts of data are updated, the write-to load for the data storage destination, transaction log storage destination, and backup data storage destination (mirrored transaction log) will also be great. For this reason, deploy them to separate disks, out of consideration for performance.

Note

When using the volume manager provided by the operating system, be aware of which physical disk the file system has been created on, for example, by deploying the data storage destination and the backup data storage destination to separate disks.



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

Resource	Role
Database cluster	The area where the database is stored. It is a collection of databases managed by an instance.
Tablespace	Stores table files and index files in a separate area from the database cluster.
Transaction log	Stores log information in preparation for a crash recovery or rollback.
Archive log	Stores log information for recovery
Corefile	Enterprise Postgres process corefile output when an error occurs with a Enterprise Postgres process.

Examples of disk deployment

The following are examples of disk deployment:

Number of disks	Disk	Deployment
3	System disk	Enterprise Postgres program

Number of disks	Disk	Deployment
		Corefile
	Connected physical disk	Data storage destination, transaction log storage destination
	Connected physical disk	Backup data storage destination
2	System disk	Enterprise Postgres program
		Corefile
		Data storage destination, transaction log storage destination
	Connected physical disk	Backup data storage destination

Proposal for disk deployment using WebAdmin

To generate an instance using WebAdmin, we propose an optimum deployment that takes into account the status of all disks at the time of instance is generation, and item 1 to 3 in the "Considerations when deploying resources" subheading above, based on the criteria below (note that a different deployment can also be specified).

- The mount point does not include national characters
- The instance administrator has the proper permissions to read and write on the mount point

Information

- Even if the mount point of each file system was deployed to a different partition or slice, it will still be verified if they are on the same physical disk.
- If the operating system volume manager was used, it will not be verified if the file systems were created on the same physical disk.

Preparing directories

The directories to be prepared depend on the way that you create the instances.

The following shows the directories that need to be prepared:

Directory to be prepared	Using WebAdmin	Using the initdb command
Data storage destination	Y (*1)	Y
Backup data storage destination	Y (*1)	O
Transaction log storage destination	O (*1) (*2)	O
Corefile output destination	N (*3)	O

Y: Required

O: Optional

N: Not required

*1: WebAdmin automatically creates a directory

*2: The default is to create in a directory in the data storage destination. When it is necessary to distribute the I/O load for the database data and the transaction log, consider putting the transaction log storage destination on a different disk from the data storage destination

*3: The default is to output to `/var/tmp/fsepuser_ Version/instanceAdmin_instanceName/core` - to change it, configure the `core_directory` and the `core_contents` parameters in `postgresql.conf` (refer to "Parameters" in the Operation Guide for details)



- The directories must meet the following conditions:
 - The directory owner must be the OS user account that you want to be the instance administrator
 - The directory must have write permission
 - The directory must be empty
- It is not possible to use a directory mounted by NFS (Network File System) when using WebAdmin.

Example

The following example shows the OS superuser creating `/database/inst1` as the directory for storing the database data and changing the owner of the directory to the OS user account "fsepuser".

```
# mkdir /database/inst1
# chown fsepuser:fsep /database/inst1
# chmod 700 /database/inst1
```

4.2.3 Estimating Resources

Estimate the resources to be used on the Enterprise Postgres.

Refer to "[Appendix I Estimating Database Disk Space Requirements](#)" for information on estimating database disk space requirements.

Refer to "[Parameters automatically set by WebAdmin according to the amount of memory](#)" when creating multiple instances with WebAdmin.

Refer to "[Appendix J Estimating Memory Requirements](#)" when creating instances with the `initdb` command, to estimate memory usage.

4.2.4 Configuring Corefile Names

If a process crashes, a corefile for the process will be generated by the operating system. If a corefile is generated with the same name as an existing corefile generated for a different process, the newly-generated corefile will overwrite the previously dumped corefile. To prevent this, configure a unique corefile name for each crash by appending the process ID, program name, and datetime.

Corefile names can be configured using the "kernel.core_pattern" and "kernel.core_uses_pid" kernel parameters.

Refer to the "man page" in "core(5)" for information on how to use these parameters.

Note that with regard to the location for storing corefiles, the operating system settings take precedence over the `core_directory` parameter of `postgresql.conf`.

4.3 Creating an Instance

There are two methods that can be used to create an instance:

- [4.3.1 Using WebAdmin](#)
- [4.3.2 Using the `initdb` Command](#)

Creating multiple instances

Multiple instances can be created.

The memory allocated needs to be adjusted when multiple instances are created with WebAdmin (refer to "[Parameters automatically set by WebAdmin according to the amount of memory](#)" for details).

Features that cannot be set up using WebAdmin

The feature below cannot be set up using WebAdmin. After creating the instance using WebAdmin, perform the additional setup tasks according to the manual for each feature described in the Operation Guide.

- Storage data protection using transparent data encryption

Note

- Instances created using the `initdb` command cannot be managed using WebAdmin.
- Always use WebAdmin to delete instances that were created using it. Because WebAdmin management information cannot be deleted, WebAdmin will determine that the instance is abnormal.
- Databases with the names 'template0' and 'template1' are automatically created when an instance is created. These databases are used as the templates for databases created later. Furthermore, a default database with the name 'postgres' is automatically created, which will be used with Enterprise Postgres commands. It is important that you do not delete these databases created by default.

4.3.1 Using WebAdmin

This section describes how to create an instance using WebAdmin.

If WebAdmin is not set up when the installation is performed, refer to "[E.1 Setting Up WebAdmin](#)" and then set up WebAdmin.

Use Windows(R) Internet Explorer 8.0, 9.0, 10.0, or 11.0 as the browser, and allow cookies and pop-up requests from the server on which Enterprise Postgres is installed.

Refer to "[Appendix D Recommended WebAdmin Environments](#)" for information on how to change the pop-up request settings and other recommended settings.

Note

- WebAdmin does not run in Windows(R) safe mode.
- If the same instance is operated from multiple WebAdmin windows, it will not work correctly.
- The browser [Back] button will not work correctly. If you need to navigate to a previous page, close the browser and then log in again.
- If the browser was not operated for a fixed period (about 15 minutes), the login window will be displayed again for the next operation. However, this will not happen if the [Monitor] window was displayed.
- Port access permissions
If a port is blocked (access permissions have not been granted) by a firewall, enable use of the port by granting access. Refer to the vendor document for information on how to grant port access permissions.
Consider the security risks carefully when opening ports.

4.3.1.1 Logging into WebAdmin

This section describes how to log in to WebAdmin.

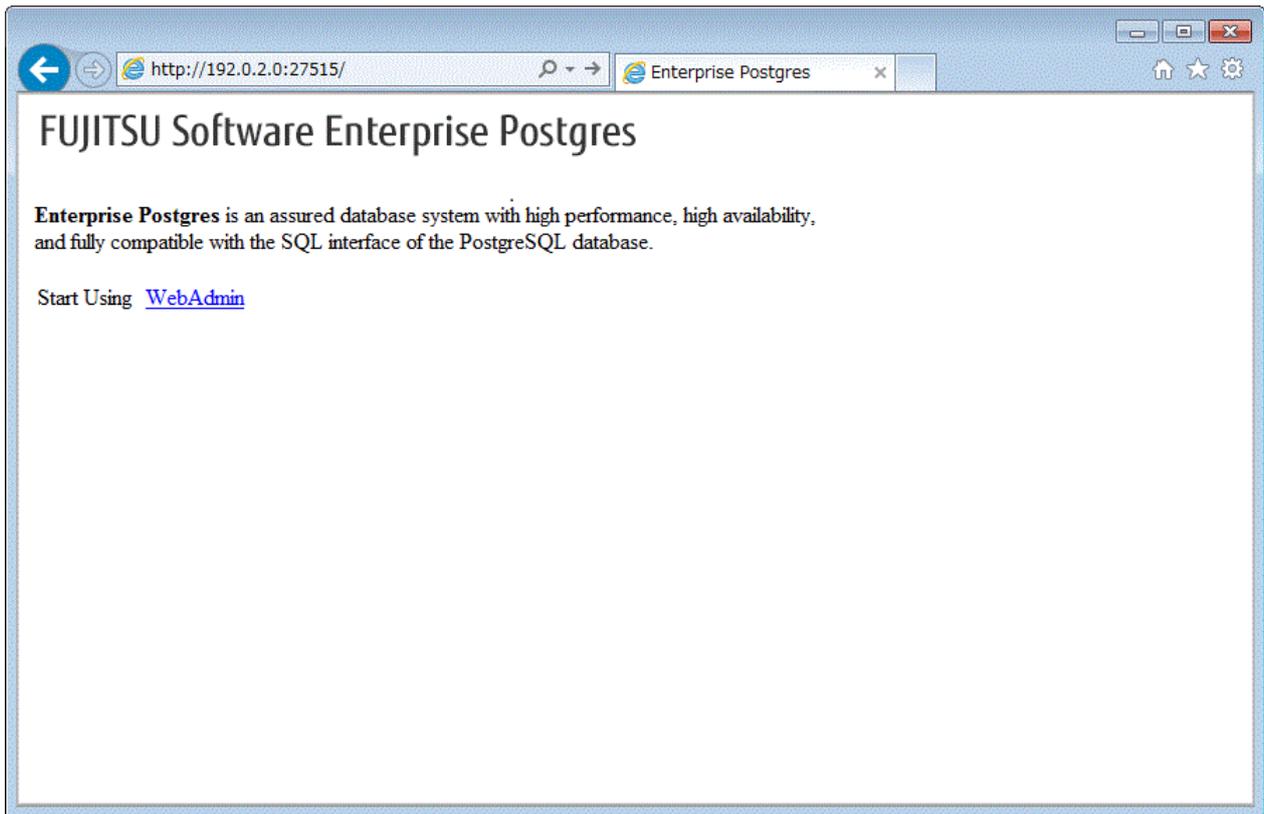
Activation URL for WebAdmin

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

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

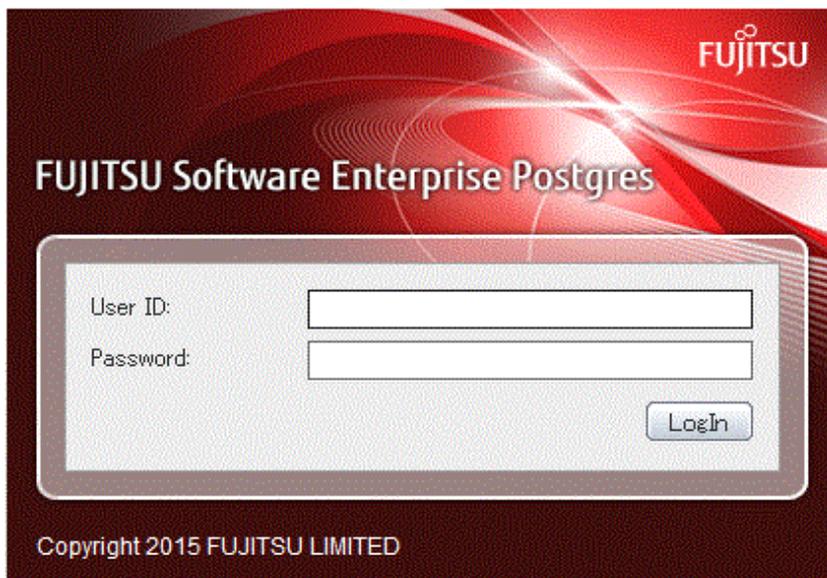
- *hostNameOrIpAddress*: Host name or IP address of the server where Enterprise Postgres is installed.
- *portNumber*: Port number of WebAdmin. The default port number is 27515.

The activation URL window shown below is displayed.



Logging in to the database server

Clicking "WebAdmin" on the startup window displays the WebAdmin [Login] window. Enter the instance administrator user ID (operating system user account name) and password, and log in to Enterprise Postgres.



4.3.1.2 Creating an Instance

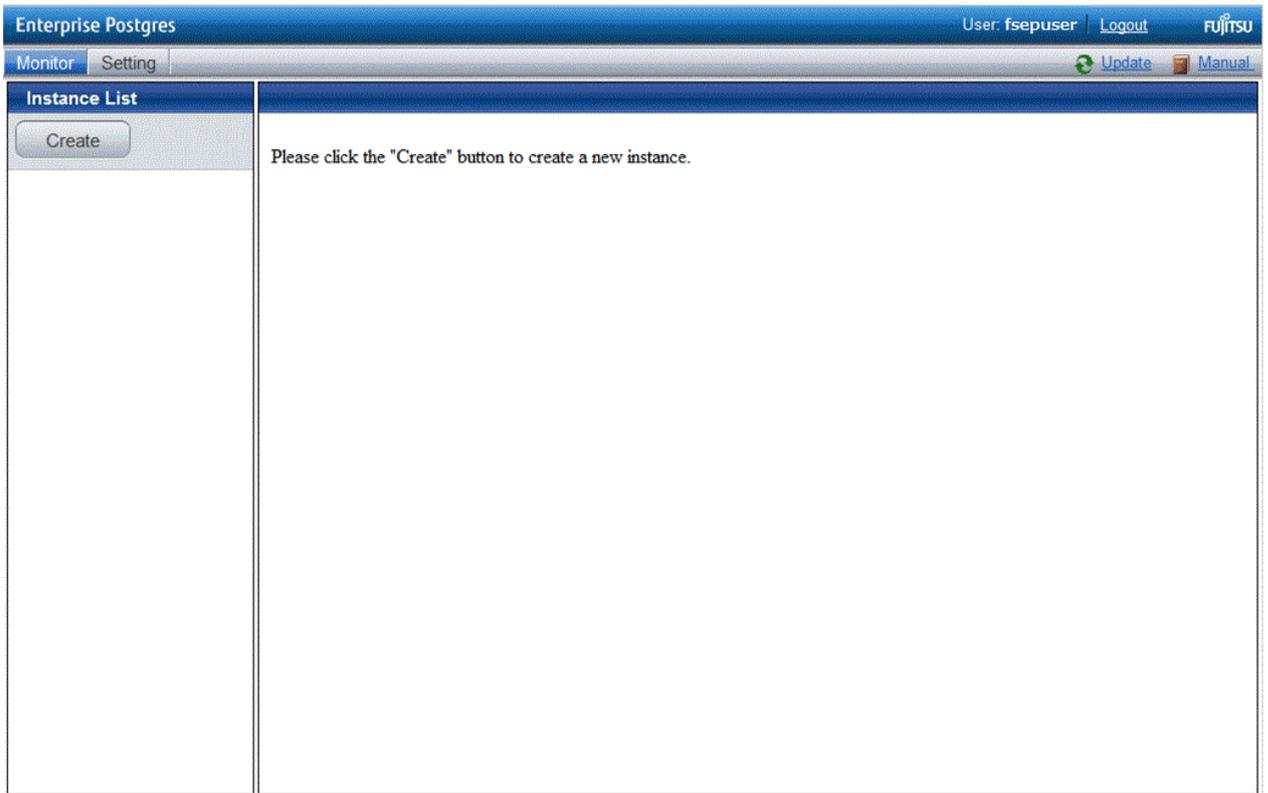
This section describes how to create an instance.

Information

WebAdmin calculates values using the formula indicated in "Managing Kernel Resources" under "Server Administration" in the PostgreSQL Documentation, and configures these in the kernel parameters. Refer to "[Appendix L Configuring Kernel Parameters](#)" for information on configuring parameter names.

Refer to "[Appendix F Configuring Parameters](#)" for information on the postgresql.conf values required to derive the set values.

1. Activate WebAdmin, and log in to the database server.
2. Click [Create] in the [Instance List] window of the [Monitor].



3. Enter the information for the instance to be created.

Create Instance

 Create a new instance.
[Please refer to Enterprise Postgres Installation and Setup Guide for Server for detail.](#)

Instance Name(*):

Data storage destination(*):

Backup data storage destination(*):

Port(*):

(*)Necessary
▶ **Advanced**

For security reasons, it is strongly recommended that the Backup data storage destination is to be placed on a disk separately from where Enterprise Postgres install directory, Data storage destination or Transaction log storage destination is located.

Enter the following items:

- [Instance Name]: Name of the database instance to manage
The name must meet the conditions below:
 - Maximum of 16 characters
 - The first character must be an ASCII alphabetic character
 - The other characters must be ASCII alphanumeric characters
- [Data storage destination]: Directory where the database data will be stored
- [Backup data storage destination]: Directory where the database backup will be stored
- [Port]: Port number of the database server

Note

- Refer to "4.2.2 Preparing Directories for Resource Deployment" - "Considerations when deploying resources" for information on points to consider when determining the data storage destination, backup data storage destination, and transaction log storage destination.
- Only the port number can be modified after the instance has been created.
- Do not specify directories that include symbolic link or multibyte characters when specifying the data storage destination or backup data storage destination.

4. Click [Advanced] to change the transaction log storage destination and database encoding.

Create Instance

 Create a new instance.
[Please refer to Enterprise Postgres Installation and Setup Guide for Server for detail.](#)

Instance Name(*):

Data storage destination(*):

Backup data storage destination(*):

Port(*):

(*)Necessary

▼ **Advanced**

Transaction log storage destination:

Encoding:

For security reasons, it is strongly recommended that the Backup data storage destination is to be placed on a disk separately from where Enterprise Postgres install directory, Data storage destination or Transaction log storage destination is located.

Enter the following items:

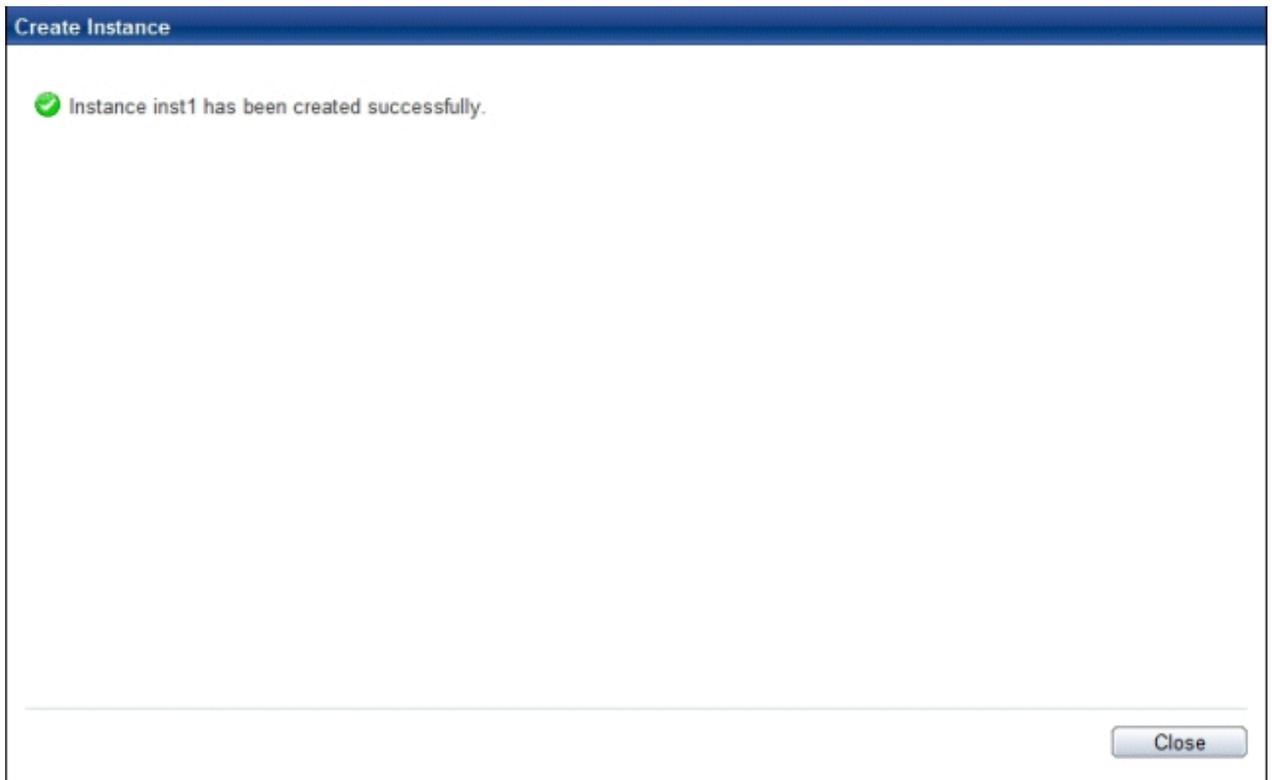
- [Transaction log storage destination]: Directory where the transaction log will be stored
- [Encoding]: Database encoding system

Note

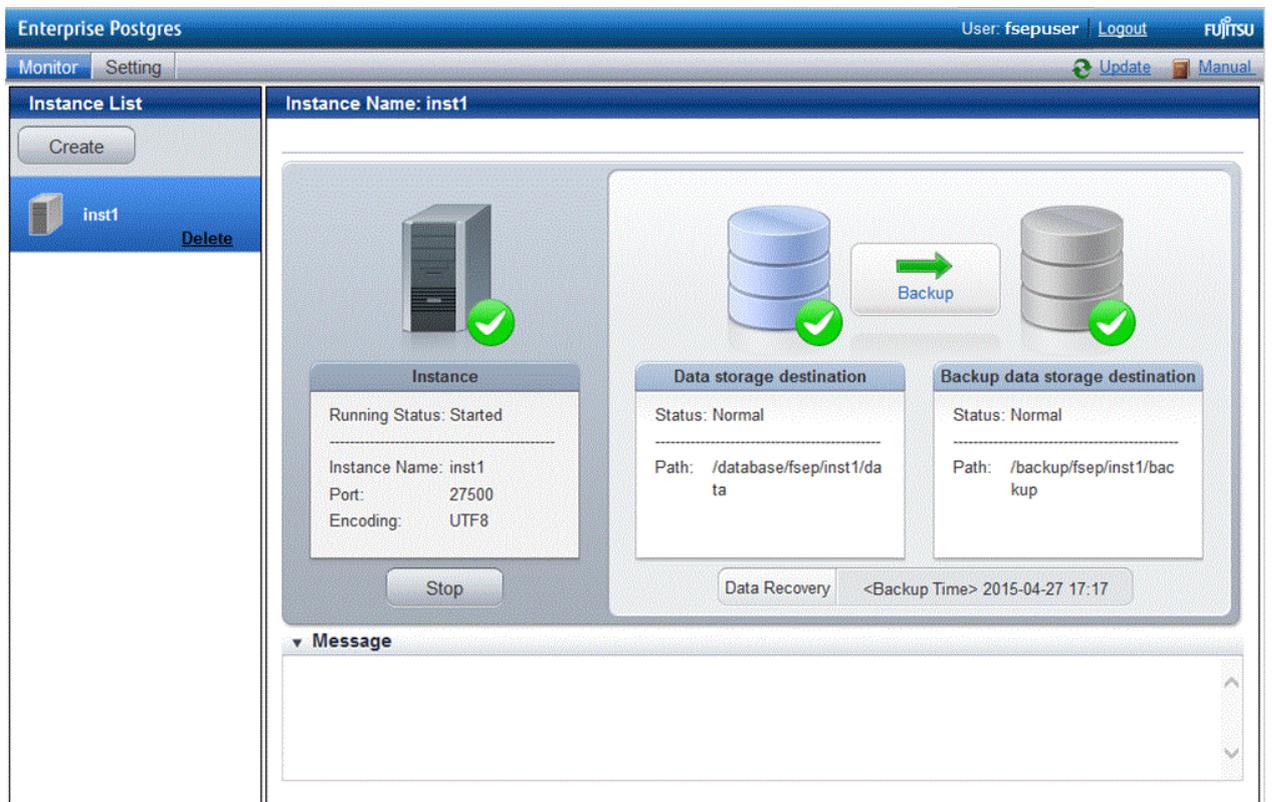
- Refer to "4.2.2 Preparing Directories for Resource Deployment" - "**Considerations when deploying resources**" for information on points to consider when determining the transaction log storage destination.
- Do not specify directories that include symbolic link or multibyte characters when specifying the transaction log storage destination.
- In the instance that was created using WebAdmin, the locale of the character set to be used in the database, and the locale of the collating sequence, are fixed using C.

5. Click [Create] to create an instance.

When the creation of the instance is complete, the following window will be displayed:



6. The instance will be started when it is created successfully.



7. Back up the basic information that was set

Perform a backup of the file shown below. The file to be backed up contains the basic information specified in this screen, and this information will be needed to recover the WebAdmin management information when a fault occurs on the system disk.

```
installDir/gui/data/gui-instances
```

4.3.1.3 Changing the Settings

You can change the following information that is set when instances are created.

You can change the character set and maximum number of connections, for example, to suit the operating and management environment for Enterprise Postgres.

- [Character set](#)
- [Client authentication](#)
- [Communication](#)
- [SQL option](#)
- [Used memory](#)

Information

These settings are the same as the parameters that can be set in the files shown below. Refer to "[Appendix F Configuring Parameters](#)" for information on the equivalence relationship between the item name and the parameter.

- postgresql.conf
- pg_hba.conf

Note

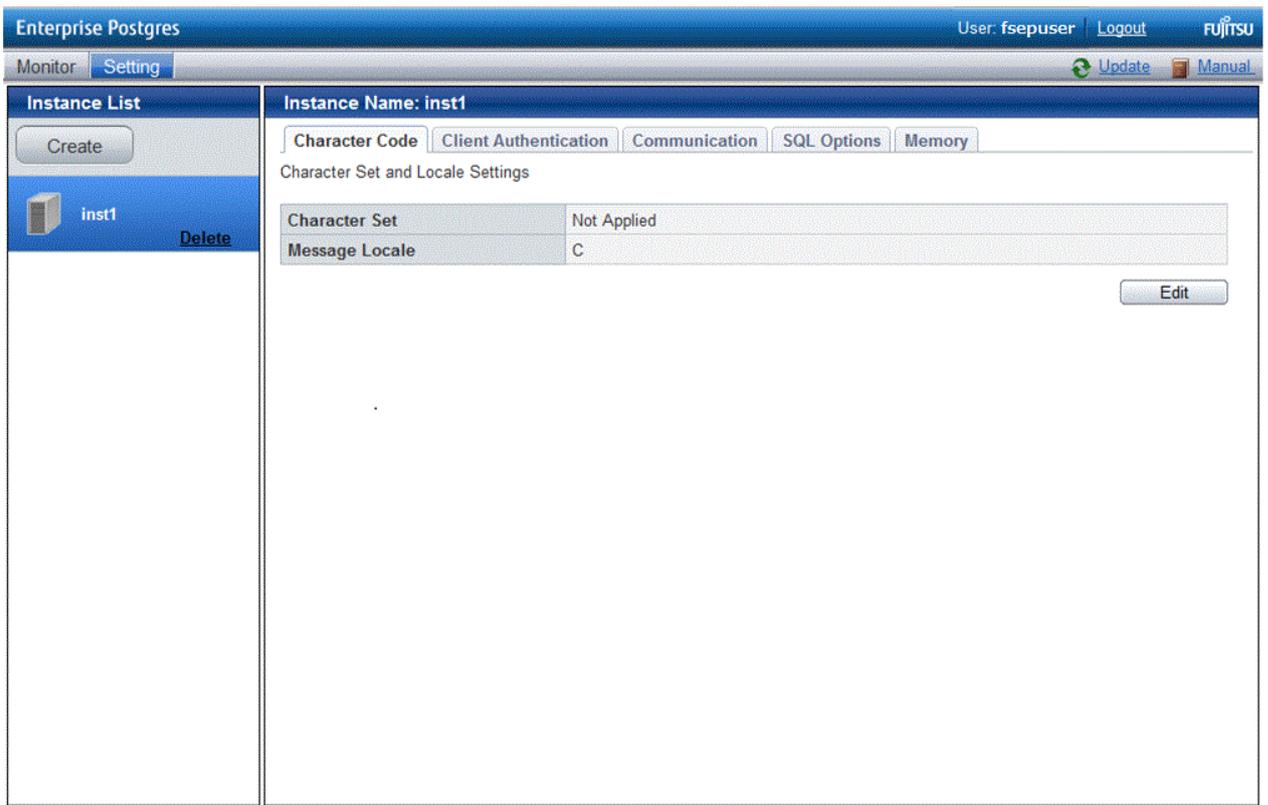
The files shown below can also be modified directly, however if a parameter not described in "[Appendix F Configuring Parameters](#)" was edited by mistake, WebAdmin may not run correctly.

- postgresql.conf
- pg_hba.conf

Changing the character set

1. Start WebAdmin and log in to the database server.
2. Select [Character Code] from the [Setting] menu.

3. Click [Edit].



4. Edit [Character Set] and [Message Locale], and then click [Modify].



See

Select a client-side encoding system that can be converted to/from the database encoding system. Refer to "PostgreSQL Documentation" - "Server Administration" - "Automatic Character Set Conversion Between Server and Client" for information on the encoding system combinations that can be converted.

Changing client authentication

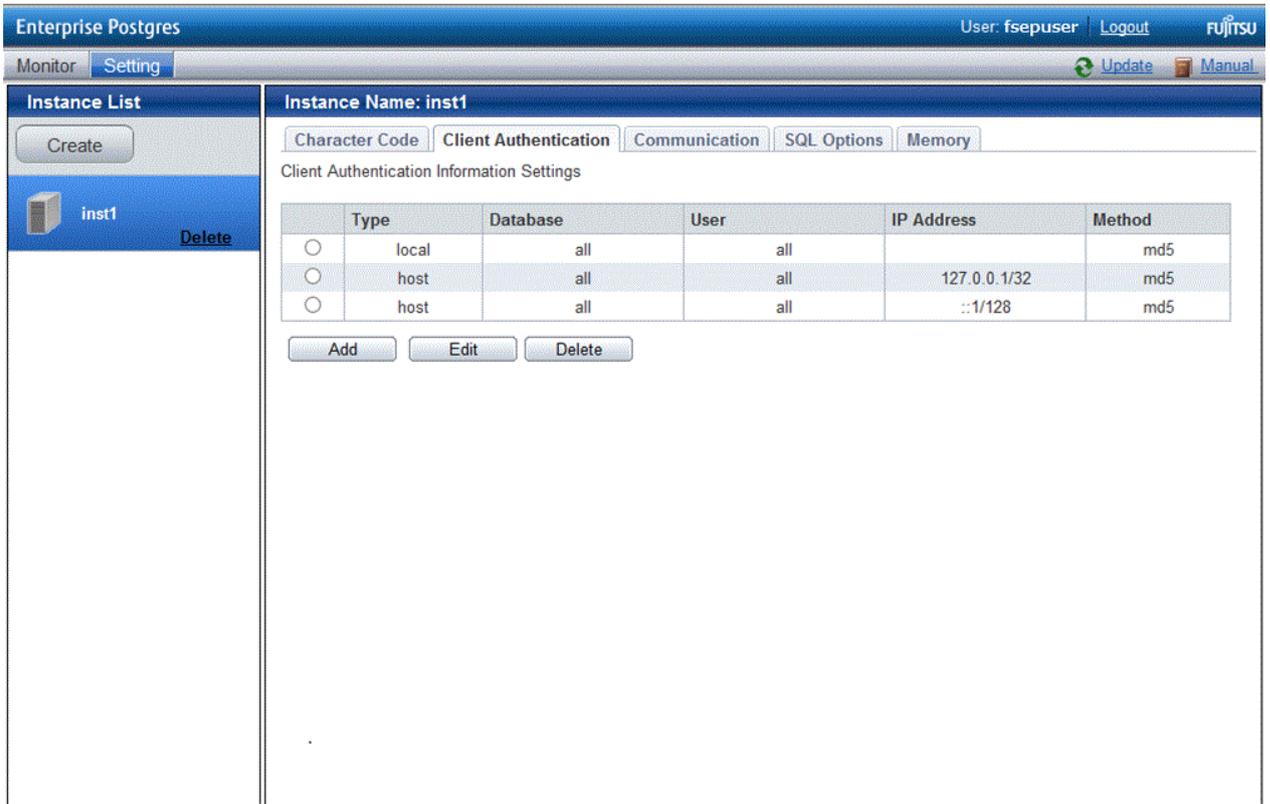
1. Start WebAdmin and log in to the database server.
2. Select [Client Authentication] from the [Setting] menu.
Click [Add] to register new authentication information.
To change authentication information, select the information, and then click [Edit].

To delete authentication information, select the information, and then click [Delete].

Note

When creating the instance, do not delete the entry below, because it is a connection required for WebAdmin to monitor the operational status of the database:

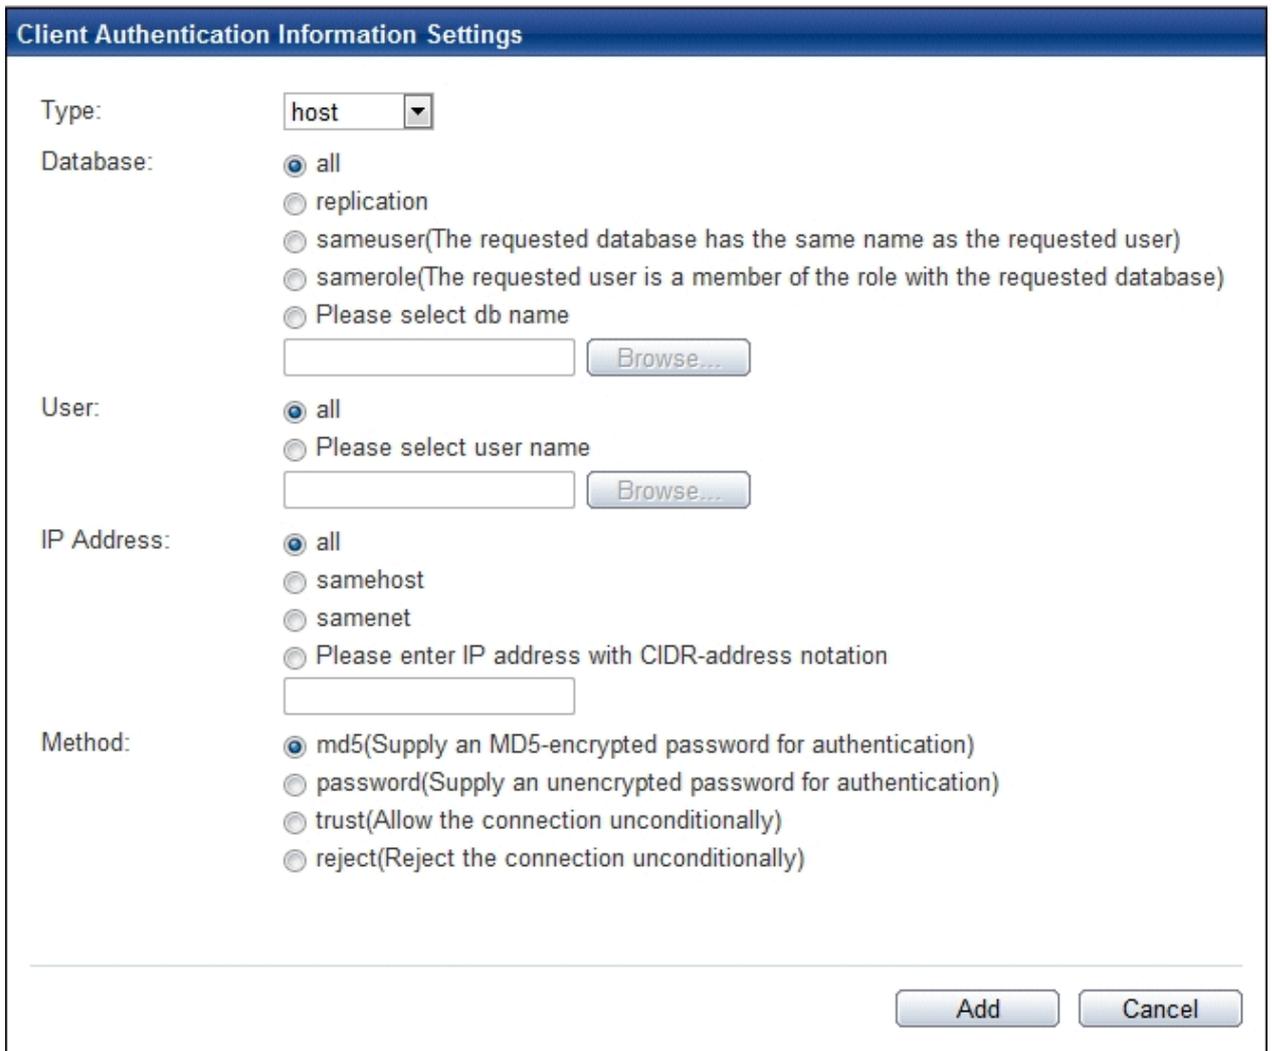
Type=local, Database=all, User=all, and Method=md5



The screenshot shows the Enterprise Postgres web interface. The top navigation bar includes 'Enterprise Postgres', 'User: fsepuser', 'Logout', and the 'FUJITSU' logo. Below the navigation bar are tabs for 'Monitor' and 'Setting', along with 'Update' and 'Manual' buttons. The main content area is split into two panels. The left panel, titled 'Instance List', contains a 'Create' button and a table with one entry for 'inst1' with a 'Delete' button. The right panel, titled 'Instance Name: inst1', has tabs for 'Character Code', 'Client Authentication', 'Communication', 'SQL Options', and 'Memory'. The 'Client Authentication' tab is active, showing 'Client Authentication Information Settings'. Below this is a table with columns for 'Type', 'Database', 'User', 'IP Address', and 'Method'. The table contains three rows: a local connection (Type: local, Database: all, User: all, Method: md5), a host connection (Type: host, Database: all, User: all, IP Address: 127.0.0.1/32, Method: md5), and another host connection (Type: host, Database: all, User: all, IP Address: ::1/128, Method: md5). Below the table are 'Add', 'Edit', and 'Delete' buttons.

	Type	Database	User	IP Address	Method
<input type="radio"/>	local	all	all		md5
<input type="radio"/>	host	all	all	127.0.0.1/32	md5
<input type="radio"/>	host	all	all	::1/128	md5

Click [Add] - the following window will be displayed.



The dialog box is titled "Client Authentication Information Settings". It contains several sections for configuration:

- Type:** A dropdown menu with "host" selected.
- Database:** Radio buttons for "all", "replication", "sameuser(The requested database has the same name as the requested user)", "samerole(The requested user is a member of the role with the requested database)", and "Please select db name". Below the last option is a text input field and a "Browse..." button.
- User:** Radio buttons for "all" and "Please select user name". Below the last option is a text input field and a "Browse..." button.
- IP Address:** Radio buttons for "all", "samehost", "samenet", and "Please enter IP address with CIDR-address notation". Below the last option is a text input field.
- Method:** Radio buttons for "md5(Supply an MD5-encrypted password for authentication)", "password(Supply an unencrypted password for authentication)", "trust(Allow the connection unconditionally)", and "reject(Reject the connection unconditionally)".

At the bottom right, there are two buttons: "Add" and "Cancel".

Changing communication

1. Start WebAdmin and log in to the database server.
2. Stop the instance.
3. Select [Communication] from the [Setting] menu.

- Click [Edit].

The screenshot shows the Enterprise Postgres WebAdmin interface. At the top, it says 'Enterprise Postgres' and 'User: fsepuser'. There are tabs for 'Monitor' and 'Setting'. The 'Instance List' on the left shows 'inst1' with a 'Delete' button. The main area is titled 'Instance Name: inst1' and has tabs for 'Character Code', 'Client Authentication', 'Communication', 'SQL Options', and 'Memory'. Under 'Communication Settings', there is a table with two rows: 'Port Number' with value '27500' and 'Max Connection' with value '100'. An 'Edit' button is located at the bottom right of this table.

- Edit [Port Number] and [Max Connection], and then click [Modify].

The screenshot shows a dialog box titled 'Communication Settings'. It has two input fields: 'Port Number' with the value '27500' and 'Max Connection' with the value '100'. Below the fields is a note: 'The value of max_connections should include the number of connections(3) used by WebAdmin.' At the bottom, there are 'Modify' and 'Cancel' buttons.

Note

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

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

*1: 3 is the default number of connections required by the system

Calculate the maximum number of connections using the following formula when changing either max_wal_senders (using streaming replication) or superuser_reserved_connections (connections reserved for use by the superuser) in postgresql.conf.

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

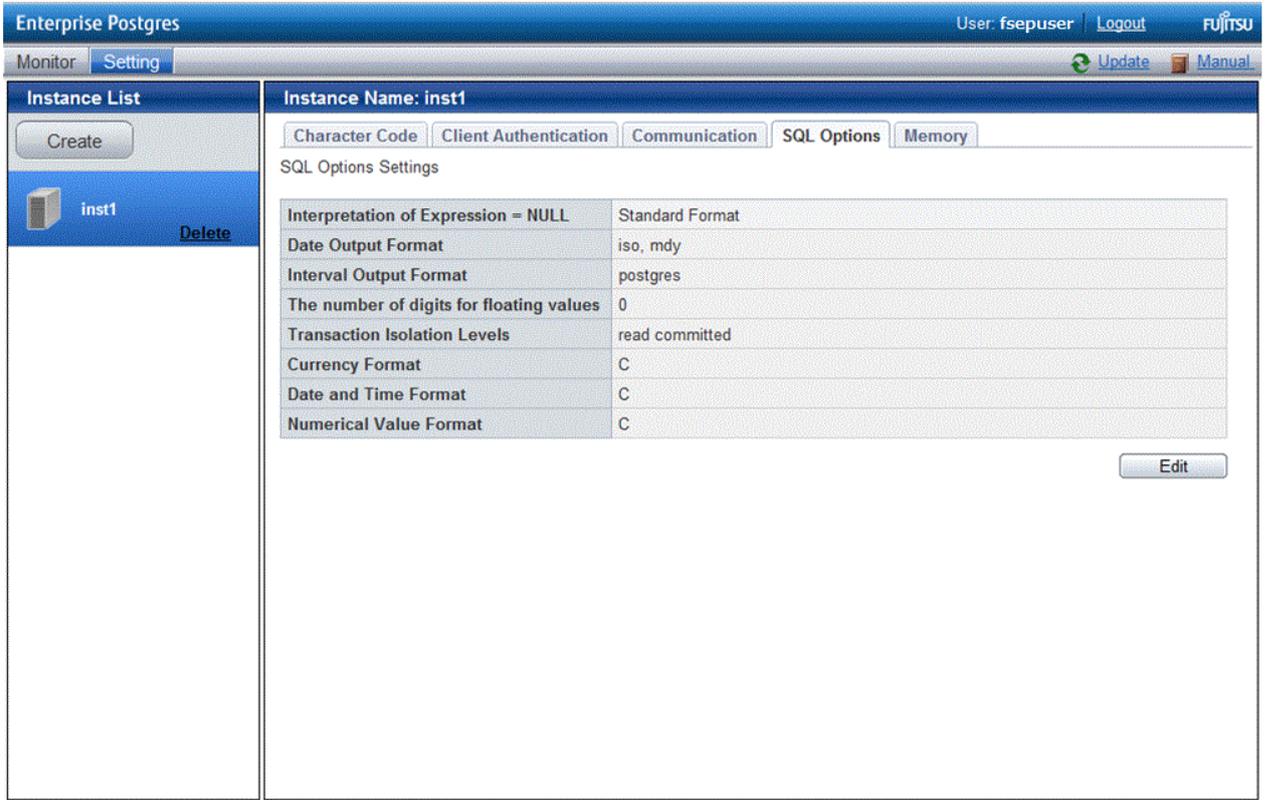
Refer to "[Appendix F Configuring Parameters](#)" for more information on postgresql.conf.

- Also check if the memory used exceeds the memory installed (refer to "[Parameters automatically set by WebAdmin according to the amount of memory](#)")

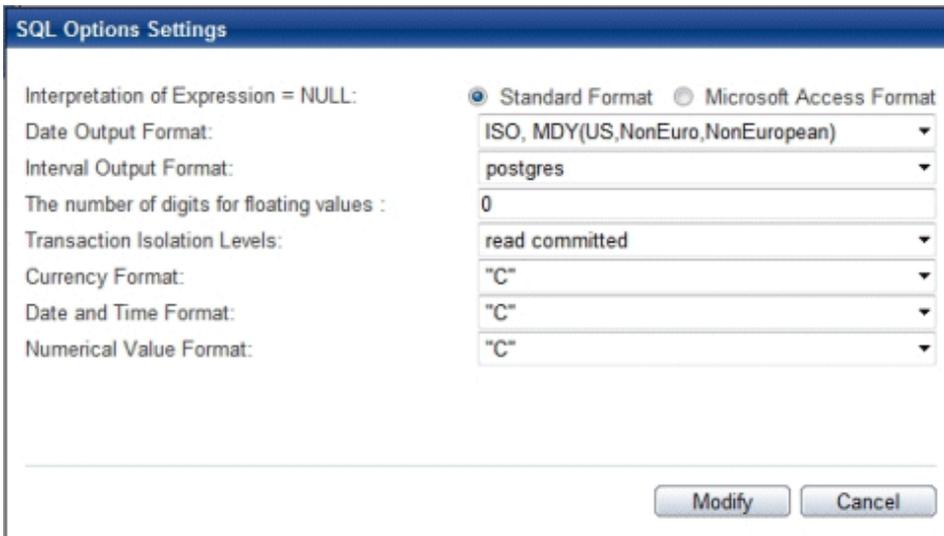
- To increase the maximum number of connections, edit the kernel parameter. Refer to "[Appendix L Configuring Kernel Parameters](#)", and "Managing Kernel Resources" in "Server Administration" in the PostgreSQL Documentation for details.

Changing the SQL option

- Start WebAdmin and log in to the database server.
- Select [SQL Options] from the [Setting] menu.
- Click [Edit].



- Edit [Interpretation of Expression = NULL], [Date Output Format], [Interval Output Format], [The number of digits for floating values], [Transaction Isolation Levels], [Currency Format], [Date and Time Format], and [Numerical Value Format], and then click [Modify].

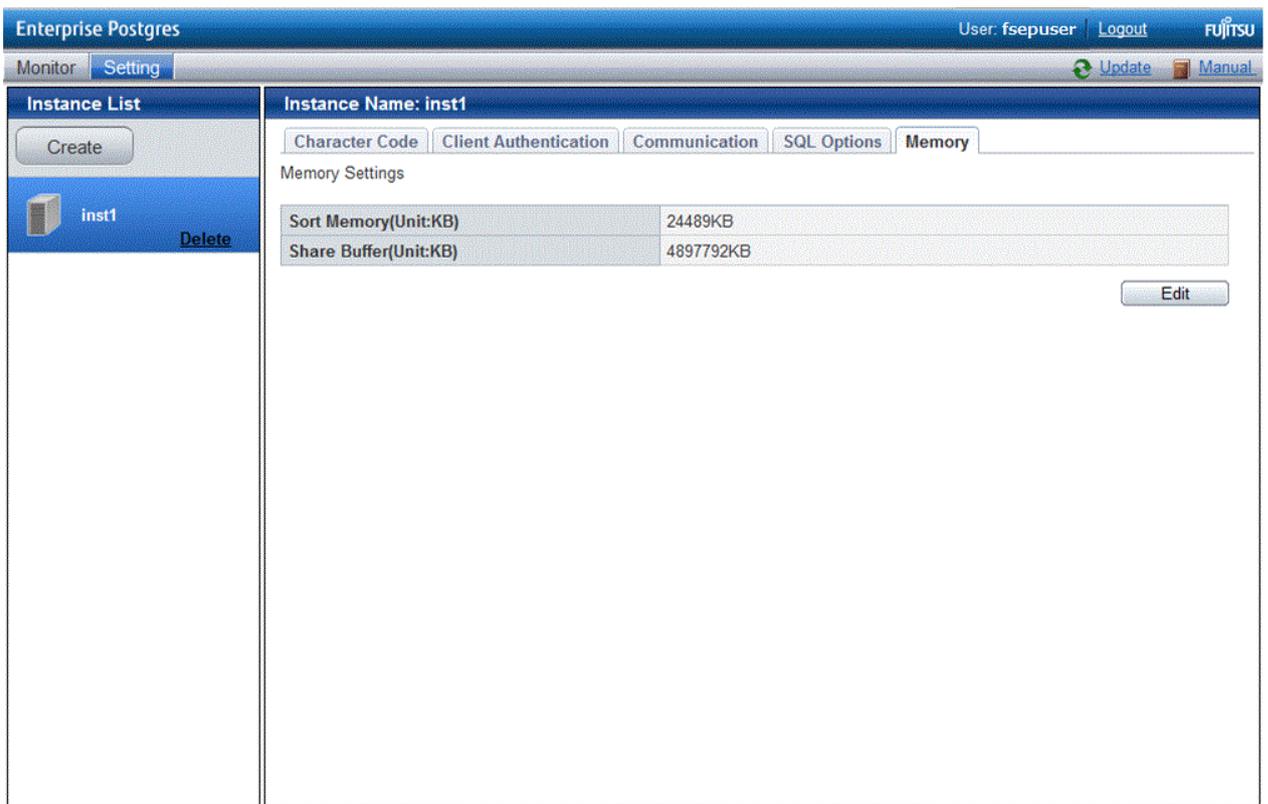


Note

If you select **postgres** in Date Output Format, dates will be output in the "12-17-1997" format, not the "Wed Dec 17 1997" format used in the PostgreSQL Documentation.

Changing allocated memory

1. Start WebAdmin and log in to the database server.
2. Stop the instance.
3. Select [Memory] from the [Setting] menu.
4. Click [Edit].



The screenshot shows the Enterprise Postgres WebAdmin interface. The top navigation bar includes "Enterprise Postgres", "User: fsepuser", "Logout", and the Fujitsu logo. Below the navigation bar, there are tabs for "Monitor" and "Setting". The main content area is divided into two panels. The left panel, titled "Instance List", contains a "Create" button and a table with one instance named "inst1" and a "Delete" button. The right panel, titled "Instance Name: inst1", has several tabs: "Character Code", "Client Authentication", "Communication", "SQL Options", and "Memory". The "Memory" tab is selected, showing "Memory Settings" with a table:

Sort Memory(Unit:KB)	24489KB
Share Buffer(Unit:KB)	4897792KB

An "Edit" button is located at the bottom right of the table.

5. Edit [Sort Memory(Unit:KB)] and [Share Buffer(Unit:KB)], and then click [Modify].



The screenshot shows a "Memory Setting" dialog box. It contains two input fields: "Sort Memory(Unit:KB)" with the value "1539" and "Share Buffer(Unit:KB)" with the value "307200". At the bottom of the dialog, there are "Modify" and "Cancel" buttons.

Note

Also check if the memory used exceeds the memory installed (refer to "[Parameters automatically set by WebAdmin according to the amount of memory](#)" for information on how to check).

6. To modify the buffer size of the database, edit the kernel parameter. Refer to "[Appendix L Configuring Kernel Parameters](#)", and "Managing Kernel Resources" in "Server Administration" in the PostgreSQL Documentation for details.

4.3.2 Using the initdb Command

This section describes the procedure to create an instance using the initdb command.

Note

If a port is blocked (access permissions have not been granted) by a firewall, enable use of the port by granting access. Refer to the vendor document for information on how to grant port access permissions.
Consider the security risks carefully when opening ports.

4.3.2.1 Editing Kernel Parameters

Refer to "[Appendix L Configuring Kernel Parameters](#)" prior to editing these settings.

After the settings are complete, check the command specifications of the relevant operating system and restart the system if required.

4.3.2.2 Creating an Instance

Create an instance, with the database cluster storage destination specified in the PGDATA environment variable or in the -D option. Furthermore, the user that executed the initdb command becomes the instance administrator.

Note

- Instances created using the initdb command cannot be managed using WebAdmin.
- If creating multiple instances, ensure that there is no duplication of port numbers or the directories that store database clusters.

See

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

The procedure to create an instance is described below.

1. Use the OS user account that you want as the instance administrator.

Connect with the server using the OS user account that you want as the instance administrator.

You cannot use the OS superuser (root).

The following example shows the OS superuser connected to the server being changed to the OS user account "fsepuser".

Example

```
# su fsepuser
```

2. Configure the environment variables

Configure the environment variables in the server with the newly created instance.

Set the following environment variables:

- PATH environment variables

Add the installation directory "/bin".

- MANPATH environment variables
Add the installation directory "/share/man".
- LD_LIBRARY_PATH environment variables
Add the installation directory "/lib".

Example

The following example configures environment variables when the installation directory is "/opt/fsepserver64".

sh, bash

```
$ PATH=/opt/fsepserver64/bin:$PATH ; export PATH
$ MANPATH=/opt/fsepserver64/share/man:$MANPATH ; export MANPATH
$ LD_LIBRARY_PATH=/opt/fsepserver64/lib:$LD_LIBRARY_PATH ; export LD_LIBRARY_PATH
```

csh, tcsh

```
$ setenv PATH /opt/fsepserver64/bin:$PATH
$ setenv MANPATH /opt/fsepserver64/share/man:$MANPATH
$ setenv LD_LIBRARY_PATH /opt/fsepserver64/lib:$LD_LIBRARY_PATH
```

3. Create a database cluster

Create the database cluster with the `initdb` command, specifying the storage destination directory.

Specify the transaction log storage destination and the locale setting option as required.

Example

```
$ initdb -D /database/inst1 --xlogdir=/transaction/inst1 --lc-collate="C" --lc-ctype="C" --
encoding=UTF8
```

Point

In some features, instance names are requested, and those names are required to uniquely identify the instance within the system. These features allow names that conform to WebAdmin naming conventions, so refer to the following points when determining the names:

- Maximum of 16 characters
- The first character must be ASCII alphabetic character
- The other characters must be ASCII alphanumeric characters

Note

- To balance I/O load, consider deploying the transaction log storage destination to a disk device other than the database cluster storage destination and the backup data storage destination.
- Specify "C" or "POSIX" for collation and character category. Performance deteriorates if you specify a value other than "C" or "POSIX", although the behavior will follow the rules for particular languages, countries and regions. Furthermore, this may need to be revised when running applications on systems with different locales.

For example, specify as follows:

```
initdb --locale="C" --lc-messages="C"
initdb --lc-collate="C" --lc-ctype="C"
```

- Specify the same string in the LANG environment variable of the terminal that starts Enterprise Postgres as was specified in `lc-messages` of `initdb` (`lc_messages` of `postgresql.conf`). If the same string is not specified, messages displayed on the terminal that was started, as well as messages output to the log file specified in the `-l` option of the `pg_ctl` command or the `postgres` command used for startup, may not be output correctly.

- Specify an encoding system other than SQL_ASCII for the database. If SQL_ASCII is used, there is no guarantee that the encryption system for data in the database will be consistent, depending on the application used to insert the data.



Refer to "Locale Support" in "Localization" in "Server Administration" in the PostgreSQL Documentation for information on locales.

4. Set port number.

Specify a port number in the port parameter of postgresql.conf. Ensure that the specified port number is not already used for other software. If a port number is not specified, "27500" is selected.

Register the specified port numbers in the /etc/services file if WebAdmin is used to create other instances. WebAdmin uses the /etc/services file to check if port numbers specified as available candidates have been duplicated.

Register any name as the service name.

5. Set the corefile output destination.

Specify the output destination of the corefile, which can later be used to collect information for investigation, by setting the core_directory and core_contents parameters of postgresql.conf.



Refer to "Parameters" in the Operation Guide for information on the settings for these parameters.

6. Set the backup storage destination.

Specify the backup data storage destination and other backup settings when backup is to be performed as a provision against database errors.



Refer to "Backup Methods" in the Operation Guide for information on specifying backup settings.

7. Start an instance.

Start with the start mode of the pg_ctl command.

It is recommended to specify the -w option, which causes the command to return after waiting for the instance to start. If the -w option is not specified, it may not be possible to determine if the starting of the instance completed successfully or if it failed.

If either of the following conditions are met, the message "FATAL:the database system is starting up(11189)" may be output.

- An application, command, or process connects to the database while the instance is starting
- An instance was started with the -w option specified

This message is output by the pg_ctl command to check if the instance has started successfully.

Therefore, ignore this message if there are no other applications, commands, or processes that connect to the database.

Example

```
$ pg_ctl start -w -D /database/inst1
```



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

4.4 Configuring Remote Connections

This section describes the settings required when connecting remotely to Enterprise Postgres from a database application or a client command.

4.4.1 When an Instance was Created with WebAdmin

Settings related to connection

The default is to accept connections from remote computers to the database.

Change "listen_addresses" in postgresql.conf to change this.

Refer to "[Appendix F Configuring Parameters](#)" for more information on postgresql.conf.

Client Authentication Information settings

The following content is set by default when WebAdmin is used to create an instance.

- Authentication of remote connections from local machines is performed.
- Only the instance administrator can perform connections in a UNIX domain socket.

When changing Client Authentication Information, select [Client Authentication] from [Setting], and then change the settings.

4.4.2 When an Instance was Created with the initdb Command

Connection settings

The default setting only permits local connections from the client to the database. Remote connections are not accepted.

Change "listen_addresses" in postgresql.conf to perform remote connection.

All remote connections will be allowed when changed as shown below.

Example

```
listen_addresses = '*'
```

Also, configure the parameters shown below in accordance with the applications and number of client command connections.

Parameter name	Parameter description
superuser_reserved_connections	Number of connections reserved for database maintenance, for example backup or index rebuilding. If you need to simultaneously perform a large number of processes that exceed the default value, change this value accordingly.
max_connections	Set the value as: <i>numberOfSimultaneousConnectionsToInstance</i> + superuser_reserved_connections

Client authentication information settings

When trying to connect from a client to a database, settings are required to determine whether the instance permits connections from the client - if it does, then it is possible to make settings to determine if authentication is required.



See

.....
Refer to "The pg_hba.conf File" in "Server Administration" in the PostgreSQL Documentation for details.
.....

4.5 Other Settings

This section describes settings that are useful for operations.

4.5.1 Error Log Settings

This section explains the settings necessary to monitor errors in applications and operations, and to make discovering the causes easier.

Make error log settings only when instances are created with the `initdb` command.

When creating instances with WebAdmin, these settings are already made and hence do not need to be set.

Furthermore, some parameters are used by WebAdmin, and if changed, may cause WebAdmin to no longer work properly. Refer to "[Appendix F Configuring Parameters](#)" for details.



Set the output destination for the system log to the server log so that it cannot be viewed by administrators of other instances.

Application errors are output to the system log or server log. The output destination directory for the system log and server log should have access permissions set so that they cannot be viewed by people other than the instance administrator.

Edit the following parameters in `postgresql.conf`:

Parameter name	Parameter description	How to enable the settings
<code>syslog_ident</code>	Used to specify labels to attach to messages, so that these can be identified when output to the system log if more than one Enterprise Postgres is used.	reload option of the <code>pg_ctl</code> mode
<code>logging_collector</code>	Specify "on" to ensure that messages are output by Enterprise Postgres to the server log file. The server log file is created in the <code>pg_log</code> directory in the database cluster.	restart option of the <code>pg_ctl</code> mode
<code>log_destination</code>	Specify "stderr,syslog" to output messages from Enterprise Postgres to the screen and either the system log or the event log.	reload option of the <code>pg_ctl</code> mode
<code>log_line_prefix</code>	Specify information to be added at the start of messages output by an instance. This information is useful for automatic monitoring of messages. You can output the SQLSTATE value, output time, executing host, application name, and user ID. Refer to "What To Log" in the PostgreSQL Documentation for details. Example: <code>log_line_prefix = '%e: %t [%p]: [%l-1] user = %u,db = %d,remote = %r app = %a '</code>	reload option of the <code>pg_ctl</code> mode



- If you want fewer application errors being output to the system log, refer to "When To Log" and "What To Log" in the PostgreSQL Documentation for information on how to reduce the output messages.
- If you want to separate errors output from other software, refer to "Where To Log" in the PostgreSQL Documentation to change the output destination to the server log file rather than the system log.

4.5.2 Configuring Automatic Start and Stop of an Instance

You can automatically start or stop an instance when the operating system on the database server is started or stopped.

Use the following procedure to configure automatic start and stop of an instance.

Note that, if an instance is activated in a failover operation, the cluster system will control the start or stop, therefore this feature should not be used. Also, when performing database multiplexing, refer to "Enabling Automatic Start and Stop of Mirroring Controller and Multiplexed Instances" in the Cluster Operation Guide.

If using an rc script

1. Set the OS user account of the instance administrator

If you logged in using the OS user account of the instance administrator, set the environment variables required for starting the instance.

2. Copy the shell script and set access permissions

Copy the sample shell script (stored in the location shown below) to /etc/rc.d/init.d, then set and register the access rights. After that, revise the shell script according to the target instance.

```
enterprisePostgresInstallDir/share/rc_fsepsvoi.sample
```

The content to be revised is as follows:

- In the "SYMDATA" variable, set the data storage destination.
- In the "SYMUSER" variable, set the OS user account of the instance administrator mentioned in step 1.
- Ensure that the file name in "/var/lock/subsys" of the "LOCKFILE" variable is the same as the shell script name.

Example

In the following example, the installation directory is "/opt/fsepserver64", and the shell script has been created as "rc_fsepsvoi_inst1".

```
# cp /opt/fsepserver64/share/rc_fsepsvoi.sample /etc/rc.d/init.d/rc_fsepsvoi_inst1
# chmod 755 /etc/rc.d/init.d/rc_fsepsvoi_inst1
```

3. Register and enable automatic start and stop

As the OS superuser, execute the chkconfig command to register and enable the script.

Execute "chkconfig --add" to register the script, and execute "chkconfig --level" to set the run level and enable the script.

Example

```
# chkconfig --add rc_fsepsvoi_inst1
# chkconfig --level 35 rc_fsepsvoi_inst1 on
```

If using systemd

1. Create a unit file

Copy the unit file sample stored in the directory below, and revise it to match the target instance.

```
enterprisePostgresInstallDir/share/fsepsvoi.service.sample
```

Example

In the following example, the installation directory is "/opt/fsepserver64", and the instance name is "inst1".

```
# cp /opt/fsepserver64/share/fsepsvoi.service.sample /usr/lib/systemd/system/
fsepsvoi_inst1.service
```

Revise the underlined portions of the options below in the unit file.

Section	Option	Specified value	Description
Unit	Description	Enterprise Postgres <u>instanceName</u>	Specifies the feature overview. Specifies the name of the target instance. (*1)

Section	Option	Specified value	Description
Service	ExecStart	<code>'/bin/bash -c '<i>installDir</i>/bin/pgx_symstd start <i>installDir</i> <i>dataStorageDestinationDir</i>'</code>	Command to be executed when the service is started.
	ExecStop	<code>'/bin/bash -c '<i>installDir</i>/bin/pgx_symstd stop <i>installDir</i> <i>dataStorageDestinationDir</i>'</code>	Command to be executed when the service is stopped.
	ExecReload	<code>'/bin/bash -c '<i>installDir</i>/bin/pgx_symstd reload <i>installDir</i> <i>dataStorageDestinationDir</i>'</code>	Command to be executed when the service is reloaded
	User	User	OS user account of the instance administrator.
	Group	Group	Group to which the instance administrator user belongs.

*1: The instance name should be as follows:

If WebAdmin is used to create the instance: *instanceName*

If the initdb command is used to create the instance: *nameThatIdentifiesTheInstance*

The naming conventions for the instance name or for identifying the instance are as follows:

- Up to 16 bytes
- The first character must be an ASCII alphabetic character
- The other characters must be ASCII alphanumeric characters

2. Enable automatic start and stop

As the OS superuser, use the systemctl command to enable automatic start and stop.

Example

```
# systemctl enable fsepsvoi_inst1.service
```

4.5.3 Settings when Using the features compatible with Oracle databases

The compatibility feature for Oracle databases enables Enterprise Postgres to be used without any special settings by creating a database instance. Note that settings are required only when using the SUBSTR function.



See

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

4.6 Integration with Message-Monitoring Software

To monitor messages output by Enterprise Postgres using software, configure the product to monitor SQLSTATE, instead of the message text - this is because the latter may change when Enterprise Postgres is upgraded.

Configure Enterprise Postgres to output messages in a format that can be read by the message-monitoring software by specifying "%e" in the log_line_prefix parameter of postgresql.conf to output the SQLSTATE value.

A setting example is shown below - it outputs the output time, executing host, application name, and user ID, in addition to the SQLSTATE value.

Example

```
log_line_prefix = '%e: %t [%p]: [%l-1] user = %u,db = %d,remote = %r app = %a '
```



See

Refer to "What To Log" in the PostgreSQL Documentation for information on how to configure the settings.

4.7 Deleting Instances

This section explains how to delete instances.

- [4.7.1 Using WebAdmin](#)
- [4.7.2 Using Server Commands](#)



Note

- Always use WebAdmin to delete instances that were created using it. Because WebAdmin management information cannot be deleted, WebAdmin will determine that the instance is abnormal.
- If you have set automatic start and stop of the instance, execute the following commands to disable the script and cancel registration.

If using systemd

```
chkconfig nameOfShellScriptForAutomaticStartAndStop off
chkconfig --del nameOfShellScriptForAutomaticStartAndStop
rm /etc/rc.d/init.d/nameOfShellScriptForAutomaticStartAndStop
```

Example

```
# chkconfig rc_fsepsvoi_inst1 off
# chkconfig --del rc_fsepsvoi_inst1
# rm /etc/rc.d/init.d/rc_fsepsvoi_inst1
```

If using systemd

```
systemctl disable nameOfUnitFileThatPerformsAutomaticStartAndStop
rm /usr/lib/systemd/system/nameOfUnitFileThatPerformsAutomaticStartAndStop
```

Example

```
# systemctl disable fsepsvoi_inst1.service
# rm /usr/lib/systemd/system/fsepsvoi_inst1.service
```

4.7.1 Using WebAdmin

This section explains how to delete instances using WebAdmin.

Use the following procedure to delete instances.

1. Stop the instance

In the [Monitor] window, click [Stop].

2. Back up files.

Before deleting the instance, back up any required files under the data storage destination, the backup data storage destination, and the transaction log storage destination.

3. Delete the instance

In the [Instance List] window, select the desired instance, and the click [Delete] - in the confirmation window, click [Delete].



Note

Deleting an instance deletes only the following lowest-level directories. If they are not required, delete them manually.

- Data storage destination
- Backup data storage destination
- Transaction log storage destination (if different from the data storage destination)

4.7.2 Using Server Commands

This section explains how to delete instances using server commands.

Use the following procedure to delete instances.

1. Stop the instance

Execute the stop mode of the `pg_ctl` command.

An example is shown below:

Example

```
$ pg_ctl stop -D /data/inst1
```

2. Back up files.

Before deleting the instance, back up any required files under the data storage destination, the backup data storage destination, and the transaction log storage destination.

3. Delete the instance

Use a standard UNIX tool (the `rm` command) to delete the following directories:

- Data storage destination
- Backup data storage destination
- Transaction log storage destination (if a directory different from the data storage directory was specified)

Chapter 5 Uninstallation

This chapter describes the procedure for uninstalling Enterprise Postgres.

5.1 Uninstallation in Interactive Mode

Uninstall according to the following procedure:



To reinstall Enterprise Postgres after it was uninstalled, and reuse an instance that was already created so that it can be managed from WebAdmin, back up the file shown below in which the WebAdmin instance management information had been defined before uninstalling Enterprise Postgres, and then restore the backed up file to its original location once Enterprise Postgres has been reinstalled.

```
installDir/gui/data/gui-instances
```



If an error occurs while the product is being uninstalled, refer to "[Appendix G Uninstall \(middleware\) Messages](#)" and take the required action.

1. Delete the operation information

If the Enterprise Postgres operation information has been registered in the operating system or another middleware product, for example, then it must be deleted. Cases in which deletion is required are as follows:

- If you have set automatic start and stop of the instance, execute the following commands to disable the script and cancel registration.

If using an rc script

```
chkconfig nameOfShellScriptForAutomaticStartAndStop off
chkconfig --del nameOfShellScriptForAutomaticStartAndStop
rm /etc/rc.d/init.d/nameOfShellScriptForAutomaticStartAndStop
```

Example

```
# chkconfig rc_fsepsvoi_inst1 off
# chkconfig --del rc_fsepsvoi_inst1
# rm /etc/rc.d/init.d/rc_fsepsvoi_inst1
```

If using systemd

```
systemctl disable nameOfUnitFileThatPerformsAutomaticStartAndStop

rm /usr/lib/systemd/system/nameOfUnitFileThatPerformsAutomaticStartAndStop
```

Example

```
# systemctl disable fsepsvoi_inst1.service
# rm /usr/lib/systemd/system/fsepsvoi_inst1.service
```

2. Stop applications and the client program

Before starting the uninstallation, stop the following:

- Applications that use the product

- Instance

Using WebAdmin

In the [Monitor] window, click [Stop].

Using server commands

Execute the pg_ctl command in stop mode.

```
$ /opt/fsepsserver64/bin/pg_ctl stop -D /database/inst1
```

- Web server feature of WebAdmin

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

Example

```
# cd /opt/fsepsserver64/gui/sbin  
# ./WebAdminStop
```

In the example above, /opt/fsepsserver64 is the installation directory.

3. Change to the superuser

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

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

4. Start the Uninstall (middleware)

Execute the following command:

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

5. Select the software

Type the number for the product to be uninstalled, and press Enter.

```
Loading Uninstaller...  
  
Currently installed products  
1. Enterprise Postgres Client(64bit) 9.4  
2. Enterprise Postgres Client(32bit) 9.4  
3. Enterprise Postgres Standard Edition(64bit) 9.4  
  
Type [number] to select the software you want to uninstall.  
[number,q]  
=>3
```

6. Start the uninstallation

To start the uninstallation, type "y" and press Enter.

To display the list of products again, type "b" and press Enter.

```
Enterprise Postgres Standard Edition(64bit)  
Description: Enterprise Postgres Standard Edition(64bit)  
Version: 9.4  
Manufacturer: Fujitsu Limited.  
Install directory: /opt/fsepsv64  
Date of install: 2015-6-2
```

```
Starting the uninstall of the software. Are you sure you want to continue?
[y,b,q]y
```

7. Finish the uninstallation

Upon successful completion, the window below is displayed.

The installation directory may remain after uninstallation. If it is not required, delete it.

```
Uninstalling...

Enterprise Postgres Standard Edition(64bit) is being uninstalled now.
 100% #####

The following products have been uninstalled successfully:
  Enterprise Postgres Standard Edition(64bit)

Uninstallation of "Enterprise Postgres Standard Edition (64bit) 9.4" has completed successfully.

Exiting Uninstaller.
```

5.2 Uninstallation in Silent Mode

Uninstall according to the following procedure:



To reinstall Enterprise Postgres after it was uninstalled, and reuse an instance that was already created so that it can be managed from WebAdmin, back up the file shown below in which the WebAdmin instance management information had been defined before uninstalling Enterprise Postgres, and then restore the backed up file to its original location once Enterprise Postgres has been reinstalled.

```
installDir/gui/data/gui-instances
```

1. Delete the operation information

If the Enterprise Postgres operation information has been registered in the operating system or another middleware product, for example, then it must be deleted. Cases in which deletion is required are as follows:

- If you have set automatic start and stop of the instance, execute the following commands to disable the script and cancel registration.

If using an rc script

```
chkconfig nameOfShellScriptForAutomaticStartAndStop off
chkconfig --del nameOfShellScriptForAutomaticStartAndStop
rm /etc/rc.d/init.d/nameOfShellScriptForAutomaticStartAndStop
```

Example

```
# chkconfig rc_fsepsvoi_inst1 off
# chkconfig --del rc_fsepsvoi_inst1
# rm /etc/rc.d/init.d/rc_fsepsvoi_inst1
```

If using systemd

```
systemctl disable nameOfUnitFileThatPerformsAutomaticStartAndStop

rm /usr/lib/systemd/system/nameOfUnitFileThatPerformsAutomaticStartAndStop
```

Example

```
# systemctl disable fsepsvoi_inst1.service
# rm /usr/lib/systemd/system/fsepsvoi_inst1.service
```

2. Stop applications and the client program

Before starting the uninstallation, stop the following:

- Applications that use the product
- Instance

Using WebAdmin

In the [Monitor] window, click [Stop].

Using server commands

Execute the `pg_ctl` command with the stop mode option specified.

```
$ /opt/fsepserver64/bin/pg_ctl stop -D /database/inst1
```

- Web server feature of WebAdmin

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

Example

```
# cd /opt/fsepserver64/gui/sbin
# ./WebAdminStop
```

In the example above, `/opt/fsepserver64` is the installation directory.

3. Change to the superuser

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

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

4. Run the uninstallation

Execute the command below.

The installation directory may remain after uninstallation. If it is not required, delete it.

Example

```
# /opt/fsepserver64/setup/suninst.sh
```

In the example above, `/opt/fsepserver64/setup` is the name of the installation directory in which the server product is installed.

5. Check the results

The uninstaller result is output to the log file.

Log file

`xyz` is the number part when the version is `x.y.z`.

64-bit products

```
Path name: /var/log/fsep_SERVER64_xyz_uninstall.log
```

32-bit products

```
Path name: /var/log/fsep_SERVER32_xyz_uninstall.log
```

Return values

The following return values are output:

Return values	Meaning
0	Uninstallation was successful.
13	Processing was interrupted.
15	The resources required for uninstallation do not exist.
19	A process is running.
60	An error occurred in Uninstall (middleware).
62	Uninstall (middleware) was started.
63	Uninstall (middleware) has not been installed, or there are no execution privileges.
99	A system error occurred.
100	The command argument is invalid.
106	The command was not executed with administrator privileges.

Appendix A Installation in Silent Mode

This appendix provides specifications for installation in silent mode.

A.1 Specification Format

The installation parameters CSV file, which is specified as the argument for the silent installer, has three columns per line in CSV format.

```
sectionName, parameterName, value
sectionName, parameterName, value
:
```

Enter the following settings in respective columns.

Item	Settings	Optional
<i>sectionName</i>	Specify the section name. There are two types of section names: "installinfo": Set the product information. "parameters": Set the parameter information for this product.	Mandatory
<i>parameterName</i>	Specify the parameter name. Each section has a valid parameter.	Mandatory
<i>value</i>	Specify the value.	Optional

Note

- Blank lines cannot be included.
- Section names and parameter names cannot be omitted.
- Undefined parameters cannot be set in lines where the section name is "installInfo". Also, note that the same parameter cannot be specified multiple times.
- Specify at least one line with the section name "parameters".
- Undefined parameters specified in lines within the "parameters" section will be ignored during execution. Note that when the same parameter is specified multiple times, the settings in the lowest line will be valid.
- The setting values for lines where the section name is "installInfo" may contain alphanumeric characters (at least one), and symbols, excluding double quotation marks (") and commas (,).
- The following characters can be used in the setting for rows with the section name "installInfo":
Alphanumeric characters and double quotation marks
- The following characters can be used in the setting for rows with the section name "parameters":
Alphanumeric characters, hyphens and forward slashes

Information

The template for the installation parameters CSV file is "mountpoint/sample/sample.csv".

A.2 List of Parameters

This section describes the parameters that can be set for each section.

installInfo section

The parameters that can be set in the installInfo section are shown below.

No.	Type	Parameter		Description
1	Software ID	Parameter name	Name	Mandatory. Specify the software identifier. For this software, specify the following value: "Enterprise Postgres"
		Value/Range	String	
		Default value	None	

Information

- In addition to "Name", the following parameters can be used in the "installInfo" section.
 - softwareName
 - OS
 - Version
 - Edition
- Note that specifying a parameter name other than "Name" and the parameters listed above will result in an error.
- The values set in these parameters do not affect silent installation.

Example

```
installInfo,Name,Enterprise Postgres
```

parameters section

The parameters that can be set in the parameters section are shown below.

No.	Type	Parameter		Description
1	Whether to install the server	Parameter name	ServerInstallExecute	Optional. Specify whether to run the installation of the server. Y: Install No: Do not install
		Value/Range	Y or N	
		Default value	Y	
2	Installation destination of the server	Parameter name	ServerInstallPath	Optional. Specify the installation destination of the software. The root directory(/) cannot be specified.
		Value/Range	Path name	
		Default value	/opt/fsepserver32, or /opt/fsepserver64	
3	Whether to install the client (32-bit)	Parameter name	Client32InstallExecute	Optional. Specify whether to run the installation of the client (32-bit). Y: Install N: Do not install
		Value/Range	Y or N	
		Default value	Y	
4	Installation destination of the client (32-bit)	Parameter name	Client32InstallPath	Optional.

No.	Type	Parameter		Description
		Value/Range	Path name	Specify the installation destination of the client (32-bit). The root directory(/) cannot be specified.
		Default value	/opt/fsepclient32	
5	Whether to install the client (64-bit)	Parameter name	Client64InstallExecute	Optional. Specify whether to run the installation of the client (64-bit). Y: Install N: Do not install This parameter is ignored on 32-bit operating systems.
		Value/Range	Y or N	
		Default value	Y	
6	Installation destination of the client (64-bit)	Parameter name	Client64InstallPath	Optional. Specify the installation destination of the client (64-bit). This parameter is ignored on 32-bit operating systems. The root directory(/) cannot be specified.
		Value/Range	Path name	
		Default value	/opt/fsepclient64	
7	Whether to install Mirroring Controller	Parameter name	MirroringControllerInstallExecute	Optional. Specify whether to run the installation of Mirroring Controller. Y: Install No: Do not install This will be ignored in all cases except EE.
		Value/Range	Y or N	
		Default value	N	
8	Whether to set up WebAdmin	Parameter name	WebSetupExecute	Optional. Specify whether to set up WebAdmin. Y: Install N: Do not install This will be ignored if the server is not installed.
		Value/Range	Y or N	
		Default value	Y	
9	Web server port number	Parameter name	WebPortNumber1	Optional. Specify the port number of the Web server. This will be ignored if the server is not installed.
		Value/Range	1024 to 32767	
		Default value	27515	
10	WebAdmin internal port number	Parameter name	WebPortNumber2	Optional. Specify the WebAdmin internal port number. This will be ignored if the server is not installed.
		Value/Range	1024 to 32767	
		Default value	27516	

Example

```
parameters,ServerInstallExecute,Y
parameters,ServerInstallPath,/opt/fsepserver64
parameters,Client32InstallExecute,Y
parameters,Client32InstallPath,/opt/fsepclient32
parameters,Client64InstallExecute,Y
parameters,Client64InstallPath,/opt/fsepclient64
```

```
parameters,MirroringControllerInstallExecute,N
parameters,WebSetupExecute,Y
parameters,WebPortNumber1,27515
parameters,WebPortNumber2,27516
```

A.3 Messages and Return Values

Messages are output when errors are detected during parametric analysis.

If an error is detected during installation of the product, a message is output to the log:

Log file

xyz is the number part when the version is x.y.z.

64-bit product

```
path name : /var/log/fsep_SERVER64_media_XYZ_install.log
```

32-bit product

```
Path name : /var/log/fsep_SERVER32_media_XYZ_install.log
```

Messages and return values

CSV file errors

The following messages are output if errors are detected while parsing CSV files.

Return value	Message	Explanation and actions
4	CSV file error:code = 1, Invalid CSV error.	There is an error in the specification format of the CSV file.
4	CSV file error:code = 2, installInfo/Name is required.	Either installInfo or the Name parameter has not been specified.
4	CSV file error:code = 3, Invalid installInfo key.	There is an error in the installInfo specification. Or the section name is invalid.
4	CSV file error:code = 4, Duplicated installInfo key.	The same parameter has been defined more than once in installInfo.
4	CSV file error:code = 5, Invalid character length.	No setting value is specified, or the specified string is too long.
4	CSV file error:code = 6, Invalid character format or encoding.	An invalid character has been specified in the installation parameters CSV file.
4	CSV file error:code = 8, Parameter is required.	There is no line in the "parameters" section.
20	The input file does not exist.	The input file does not exist.
21	The value of @1@ is incorrect, the value is @2@.	The value is incorrect. Specify the correct value. The parameter name is displayed in @1@. The specified parameter is displayed in @2@.
22	The value same at @1@ and @2@ is specified.	The same value is specified in different parameters. Specify different values. The parameter name is displayed in @1@ and @2@.
23	The @2@ of @1@ already exists.	The path already exists. Specify a different path. The parameter name is displayed in @1@.

Return value	Message	Explanation and actions
		The specified value is displayed in @2@.
26	Port number @1@ is already used in the service file(/etc/services).	The port number is already being used. Specify an unused port number. The port number is displayed in @1@
28	There is no product for installation.	No product was specified for installation. Specify "Y" for at least one of the following parameters: - ServerInstallExecute - Client32InstallExecute - Client64InstallExecute
29	USAGE : silent.sh inputfile	The argument specified in the command is incorrect. Specify the correct argument.

Product installer errors

The following return values are returned when errors occur while the product installer is running.

Return value	Explanation
0	Completed successfully.
11	A product that cannot coexist with Enterprise Postgres has been installed.
12	This OS is not supported.
13	Processing was interrupted.
14	The resources required for installation do not exist on the media.
15	The resources required for installation do not exist.
16	Unable to change the permission because the directory does not exist.
19	A process is running.
30	Failed to copy the installation material.
31	Upgrade installation not possible.
32	The product is already installed.
60	An error occurred in Uninstall (middleware).
61	Failed to install Uninstall (middleware).
62	Uninstall (middleware) was started.
68	Failed to set up WebAdmin.
99	A system error occurred.
106	The command was not executed with administrator privileges.

A.4 CSV File Format

The format of CSV files is based on RFC4180, with the following specifications.

Records

- Separate each record with a "CRLF" newline (operation is not guaranteed with only a "CR" or "LF" newline).
- Specify a newline at the end of a file.

- Separate each field within a record with a halfwidth comma ",".

Format	Record		
aaa,bbb,ccc	aaa	bbb	ccc

- If several commas are entered in succession, or if a comma precedes a newline, the data following the comma is regarded as empty.

Format	Record		
aaa,,ccc	aaa		ccc
aaa,bbb,	aaa	bbb	

- Headers cannot be specified.

Format	Record		
	field1	field2	field 3
aaa,bbb,ccc	aaa	bbb	ccc

Appendix B Enterprise Postgres Migration Procedures

The tasks shown in the following table are required to migrate from an older version/level:

Table B.1 Reference destination for migration tasks by product

Product before migration	Product after migration	Operation mode	Refer to:
Enterprise Postgres Standard Edition V9.2	Enterprise Postgres Standard Edition 9.4	Basic operations	B.1
		Failover operations using the standby feature	B.2

B.1 Migration Tasks for Basic Operations

If the Enterprise Postgres version that was used to create the instance is V9.2, execute "CREATE EXTENSION" in order to reference Information added by Enterprise Postgres.



Refer to "Information that can be Collected" in the Operation Guide for information on Information added by Enterprise Postgres.

Follow the procedure below to install any missing extensions:

1. Check if CREATE EXTENSION is required for the target database.

Use the psql command to connect to the target database (in the example below, "postgres"), and then execute the following SELECT statements:

```
postgres=# SELECT * FROM pgx_stat_sql;
ERROR: Relation "pgx_stat_sql" does not exist (10292)
postgres=# SELECT * FROM pgx_stat_lwlock;
ERROR: Relation "pgx_stat_lwlock" does not exist (10292)
postgres=# SELECT * FROM pgx_stat_latch;
ERROR: Relation "pgx_stat_latch" does not exist (10292)
postgres=# SELECT * FROM pgx_stat_walwriter;
ERROR: Relation "pgx_stat_walwriter" does not exist (10292)
```

CREATE EXTENSION is required if the view specified in the SELECT statement does not exist.

2. Execute the statements below to install the statistics views independently added by Enterprise Postgres and that were not displayed in step 1. If there are multiple databases, execute the statements for all of them.

If the connected database is not a registration target, these views cannot be browsed. If the database has been registered for template1, there is no need to consider this point for databases added thereafter.

The statistics are reset when the view is installed.

```
postgres=# CREATE EXTENSION pgx_stat_sql;
CREATE EXTENSION
postgres=# CREATE EXTENSION pgx_stat_lwlock;
CREATE EXTENSION
postgres=# CREATE EXTENSION pgx_stat_latch;
CREATE EXTENSION
postgres=# CREATE EXTENSION pgx_stat_walwriter;
CREATE EXTENSION
postgres=# \q
```

B.2 Migration Procedures when Performing Failover Operations using the Standby Feature

This section describes the migration procedure for migration from an older version/level for failover operations using the standby feature.

Perform the migration procedure as described in the following table:

Step	Task	
	Active node	Standby node
1	Install PRIMECLUSTER. (*1)	Install PRIMECLUSTER. (*1)
2	Stop the standby cluster application. (*2)	
3	Stop RMS. (*3)	
4	Start the GDS volume.	
5	Mount the file system.	
6	Modify the configuration of the standby cluster application.	
7	Delete the resource information for the Enterprise Postgres database cluster.	
8	Upgrade Enterprise Postgres.	Upgrade Enterprise Postgres.
9	Register the resource information for the Enterprise Postgres database cluster.	
10	Modify the configuration of the standby cluster application.	
11	Unmount the file system.	
12	Stop the GDS volume.	
13	Start RMS.	
14	Start the standby cluster application.	
15	Check the behavior.	

*1: If a supported version of PRIMECLUSTER is not been installed

*2: If the standby cluster application is running

*3: If RMS is running.

The details for each step are described below.

1. Install PRIMECLUSTER.

On all nodes on that comprise the cluster system, check the current version of PRIMECLUSTER. If the required version is not installed, then install it.



See

- Refer to "[2.2 Related Software](#)" for information on the PRIMECLUSTER versions.
- Refer to the PRIMECLUSTER Installation Guide for information on installing PRIMECLUSTER.

2. Stop the standby cluster application.

Check the standby cluster application status. If it is running, stop it.

3. Stop RMS.

Check the RMS status on all nodes that comprise the cluster system. If RMS is running, stop it.

4. Start the GDS volume.

Start the GDS volume on the active node.

5. Mount the file system.

Mount the file system where the data storage destination directory was stored on the active node.

6. Modify the configuration of the standby cluster application.

Delete the database cluster resources from the standby cluster application.

7. Delete the resource information for the Enterprise Postgres database cluster.fs

Use the `pgx_pclrsc` command to delete the database cluster resource information on the active node.

8. Upgrade Enterprise Postgres.

Refer to "[Chapter 3 Installation](#)", and then upgrade Enterprise Postgres on all nodes that comprise the cluster system.

9. Register the resource information for the Enterprise Postgres database cluster.

Use the `pgx_pclrsc` command to register the database cluster resource information that was deleted in step 7 on the active node.

10. Modify the configuration of the standby cluster application.

Register the database cluster resources in the standby cluster application.

11. Unmount the file system.

Unmount the file system where the data storage destination directory was stored on the active node.

12. Stop the GDS volume.

Stop the GDS volume on the active node.

13. Start RMS.

Start RMS on all nodes that comprise the cluster system.

14. Start the standby cluster application.

Start the standby cluster application.

15. Check the behavior.

Execute `switch` or `rebuild` to check if the migration was successful.



See

.....
Refer to the Cluster Operation Guide for information on failover operations.
.....

Appendix C Using WebAdmin in a Multi-version Installation Environment

To use WebAdmin with multiple Enterprise Postgres versions in a multi-version installation environment, ensure that the "[Requirements for supporting the use of WebAdmin in a multi-version installation environment](#)" are matched.

If the requirements are not matched, perform the steps described in "[How to support WebAdmin in a multi-version installation environment](#)", and then create the instance in the multi-version installation environment.

Requirements for supporting the use of WebAdmin in a multi-version installation environment

1. Enterprise Postgres is upgraded after the instance was created using WebAdmin, and
2. The Enterprise Postgres version installed as the multi-version installation is the one before the upgrade mentioned in item 1, and
3. WebAdmin is used in the Enterprise Postgres multi-version installation environment mentioned in item 2, and
4. The steps described in "[How to support WebAdmin in a multi-version installation environment](#)" are not performed for the instance mentioned in item 1.

How to support WebAdmin in a multi-version installation environment

Perform this procedure in the Enterprise Postgres environment (the environment in item 1 above) after upgrading.

1. Stop the existing instance.

If the WebAdmin instance created before the upgrade is running, then stop it.

2. Modify the settings information.

- Change the following parameters in the "postgresql.conf" parameter file:
 - log_directory
 - core_directory
 - syslog_ident

How to modify the parameter

Delete the "fsep_@@@_XX_YY" part, and modify the name so that it is unique in the server.

Example)

Pre-upgrade version and level: Enterprise Postgres V9.2 Standard Edition 64-bit

User name: fsepuser

Instance name: inst1

Unique name: package1

[log_directory][core_directory]

Before: /var/tmp/fsep_0920_SE_64/fsep_inst1/log

After: /var/tmp/fsep_package1/fsep_inst1/log

[syslog_ident]

Before: fsep_0920_SE_64_fsep_inst1

After: fsep_package1_fsep_inst1

- Create the directories for resource deployment using the modified name.

If the directory specified in "log_directory" and "core_directory" does not exist, create it.

Additionally, the directory access privileges must meet the following requirements:

- The directory owner is the OS user account that will become the instance administrator.

- There are permissions to write to the directory.
- The directory is empty.

3. Start the instance

Start the instance stopped in step 1.

 Note

.....
If operations management middleware (for example, Systemwalker Centric Manager) is used to monitor instance ID information for the service name or server log, modify the operations management middleware settings together with these changes. If multiple methods were specified for the same table combination, the method specified last will be effective.
.....

Appendix D Recommended WebAdmin Environments

This appendix describes the recommended WebAdmin environment.

D.1 Recommended Browser Settings

- Use a display resolution to be 1024 x 768 or higher, and 256 colors or more.
- Select [View] >> [Text size] >> [Medium].
- Click [Tools] >> [Internet options] >> [General] >> [Fonts], and then:
 - Set [Webpage font] to [Times New Roman]
 - Set [Plain text font] to be [Courier New]

D.2 How to Set Up the Pop-up Blocker

If the Pop-up Blocker is enabled, use the procedure below to configure settings to allow pop-ups from the server where Enterprise Postgres is installed.

1. Click [Tools] >> [Internet options], and then select the [Privacy] tab.
If [Turn on Pop-up Blocker] is not selected, the Pop-up Blocker feature will not operate, and therefore steps below are not required.
2. Click [Settings].
3. In the [Pop-up Blocker Settings] window, enter in the [Address of website to allow] the URL of the server where Enterprise Postgres is installed, and then click [Add].
4. Click [Close].
5. In the [Internet Options] window, click [OK].

Appendix E Setting Up and Removing WebAdmin

This appendix describes how to set up and remove WebAdmin.

E.1 Setting Up WebAdmin

This section explains how to set up WebAdmin.

E.1.1 Setting Up WebAdmin

Follow the procedure below to set up WebAdmin.

1. Change to the superuser

Acquire superuser privileges on the system.

Example

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

2. Set up WebAdmin

Set up WebAdmin.

Example

If Enterprise Postgres is installed in "/opt/fsepserver64":

```
# cd /opt/fsepserver64/gui/sbin
# ./WebAdminSetup
```

3. Specify the port number

Specify the following port numbers to be used in WebAdmin.

Refer to the "/etc/services" file and only change to a different port number if there is overlap with a port number from another service.

Make a note of the port number for the Web server, because it will be required for activating the WebAdmin window.

Item	Value (recommended value)
Web server port number enter port number of Web Server (default: 27515):	27515
WebAdmin internal port number enter Internal port number for WebAdmin (default: 27516):	27516
WebAdmin automatic start Start WebAdmin automatically when system starting ? [y,n] (default: y)	y

Web server port number

Specify a numeric value from 1024 to 32767 for the port number to be used for communication between the Web browser and the Web server.

The Web server port number will be registered as a port number with the following service name in the "/etc/services" file.

64-bit product:

fsep_94_64_WebAdmin_Port1

32-bit product:

fsep_94_32_WebAdmin_Port1

WebAdmin internal port number

Specify a numeric value from 1024 to 32767 for the port number to be used for communication between the Web server and the WebAdmin runtime environment.

The WebAdmin internal port number will be registered as a port number with the following service name in the /etc/services file:

64-bit product:

```
fsep_94_64_WebAdmin_Port2
```

32-bit product:

```
fsep_94_32_WebAdmin_Port2
```

WebAdmin automatic start

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



- Unused port numbers

Irrespective of the information specified in the "/etc/services" file, unused port numbers in the OS and other products can sometimes be automatically numbered and then used, or port numbers specified in environment files within products may also be used. Check the port numbers used by the OS and other products, and ensure that these are not duplicated.

- Access restrictions

Prevent unauthorized access and maintain security by using a firewall product, or the packet filtering feature of a router device, to restrict access to the server IP address and the various specified port numbers.

- Port access permissions

If a port is blocked (access permissions have not been granted) by a firewall, enable use of the port by granting access. Refer to the vendor document for information on how to grant port access permissions.

Consider the security risks carefully when opening ports.

E.1.2 Activating the Web Server Feature of WebAdmin

Follow the procedure below to activate the Web server feature of WebAdmin.

1. Change to the superuser

Acquire superuser privileges on the system.

Example

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

2. Activate the Web server feature of WebAdmin

Execute the WebAdminStart command to activate the Web server feature of WebAdmin.

Example

If Enterprise Postgres is installed in "/opt/fsepserver64":

```
# cd /opt/fsepserver64/gui/sbin  
# ./WebAdminStart
```

E.1.3 Stopping the Web Server Feature of WebAdmin

Follow the procedure below to stop the Web server feature of WebAdmin.

1. Change to the superuser

Acquire superuser privileges on the system.

Example

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

2. Stop the Web server feature of WebAdmin

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

Example

If Enterprise Postgres is installed in "/opt/fsepserver64":

```
# cd /opt/fsepserver64/gui/sbin  
# ./WebAdminStop
```

E.2 Removing WebAdmin

This section explains how to remove WebAdmin.

This removal procedure stops WebAdmin and ensures that it no longer starts automatically when the machine is restarted.

1. Change to the superuser

Acquire superuser privileges on the system.

Example

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

2. Remove WebAdmin setup

Execute the WebAdminSetup command to remove WebAdmin setup.

Example

If Enterprise Postgres is installed in "/opt/fsepserver64":

```
# cd /opt/fsepserver64/gui/sbin  
# ./WebAdminSetup -d
```

Appendix F Configuring Parameters

WebAdmin operates and manages databases according to the contents of the following configuration files:

- [postgresql.conf](#)

Contains various items of information that define the operating environment of Enterprise Postgres.

- [pg_hba.conf](#)

Contains various items of information related to client authentication.

These configuration files are deployed to a data storage destination. Data is written to them when the instance is created by WebAdmin and when settings are changed, and data is read from them when the instance is started and when information from the [Setting] menu is displayed.

Direct editing of each configuration file is possible with a text editor.



See

Refer to "Server Configuration" and "Client Authentication" in "Server Administration" in the PostgreSQL Documentation for information on the parameters.

postgresql.conf

Parameters that can be changed in WebAdmin

The postgresql.conf parameters that can be changed in WebAdmin are shown below:

Tab	WebAdmin item	Parameter
Character Code	Character Set	client_encoding
	Message Locale	lc_messages
Communication	Port Number	port
	Max Connection	max_connections
SQL Options	Interpretation of Expression = NULL	transform_null_equals
	Date Output Format	DateStyle (*1)
	Interval Oputput Format	IntervalStyle
	The number of digits for floating values	extra_float_digits
	Transaction Isolation Levels	default_transaction_isolation
	Currency Format	lc_monetary
	Date and Time Format	lc_time
Memory	Sort Memory(Unit:KB)	work_mem
	Share Buffer(Unit:KB)	shared_buffers

*1: If you specify "Postgres" as the output format, dates will be output in the "12-17-1997" format, not the "Wed Dec 17 1997" format used in the PostgreSQL Documentation.

Parameters set by WebAdmin

Parameters set by WebAdmin during instance startup are shown below (they will be ignored even if specified in postgresql.conf):

Parameter	Value
listen_addresses	*

Parameter	Value
log_destination	'stderr,syslog'
logging_collector	on
log_line_prefix	'%e: %t [%p]: [%l-1] user = %u,db = %d,remote = %r app = %a '
log_directory	'/var/tmp/fsep_version/instanceAdmin_instanceName/log'
log_filename (*1) (*2)	'logfile-%a.log'
log_file_mode	0600
log_truncate_on_rotation	on
log_rotation_age	1d

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

*2) If the date changes while the instance is stopped, old logs are not deleted and continue to exist. Manually delete old logs that are no longer required to release disk space.

Parameters automatically set by WebAdmin according to the amount of memory

The postgresql.conf parameters automatically set according to the amount of installed memory, during the creation of instances by WebAdmin, are shown below:

Parameter	Value
shared_buffers	30% of the machine's installed memory
work_mem	30% of the machine's installed memory / max_connections / 2
effective_cache_size	75% of the machine's installed memory
maintenance_work_mem	10% of the machine's installed memory / (1 + autovacuum_max_workers)

When determining the values to be configured in the above parameters, you must take into account any anticipated increases in access volume or effects on performance during business operations, such as the number of applications and commands that will access the instance, and the content of processes. Also, note that in addition to Enterprise Postgres, other software may be running on the actual database server. You will need to determine the degree of priority for the database and other software, as well as the memory allocation size.

WebAdmin automatically configures complex parameter settings such as those mentioned above, based on the size of the internal memory of the machine. This enables maximum leverage of the machine memory to facilitate resistance against fluctuations during business operations.

Accordingly, the effects of the above-mentioned factors must be estimated and taken into account when determining and configuring parameter values, so that memory resources can be effectively allocated among other software or instances, and so that adverse effects can be mutually avoided. Refer to "Memory" in "Resource Consumption", and "Planner Cost Constants" in "Query Planning", under "Server Administration" in the PostgreSQL Documentation for information on parameter values and required considerations.

Parameter values can be modified using the WebAdmin [Setting] menu, or edited directly using a text editor.

If adding an instance, determine the parameter values, including for existing instances, and make changes accordingly.



Kernel parameters need to be tuned according to the parameters being changed. Refer to "[Appendix L Configuring Kernel Parameters](#)", and "Managing Kernel Resources" in "Server Administration" in the PostgreSQL Documentation for information on tuning kernel parameters.

Note

- Do not directly edit the following postgresql.conf parameters with a text editor, otherwise WebAdmin may not work properly if you make a mistake):
 - port
 - archive_mode
 - archive_command
 - wal_level
 - log_line_prefix
 - log_destination
 - logging_collector
 - log_directory
 - log_file_mode
 - log_filename
 - log_truncate_on_rotation
 - log_rotation_age
 - backup_destination
 - You must take care with the following parameter:
 - superuser_reserved_connections

Set it to a number that includes the 3 connections required in WebAdmin (the default is 3).
-

pg_hba.conf

Refer to "Client Authentication" in "Server Administration" in the PostgreSQL Documentation for information on content that can be configured in pg_hba.conf.

Note

- Configure the instance administrator permissions in the "local" connection format settings. WebAdmin may not work properly if permissions are not configured.
 - If you specify an item or value that cannot be set by WebAdmin when editing the pg_hba.conf file with a text editor, it will not be possible to reference that line from WebAdmin.
-

Appendix G Uninstall (middleware) Messages

G.1 Messages output by FJSVcir

FSP_FJSVCIR_CIRINST: ERROR: 101: CIRINST101: Administrator privilege is required.

Description

The user does not have administrator privileges.

Action method

Login with administrator privileges and run the command again.

FSP_FJSVCIR_CIRINST: ERROR: 102: CIRINST102: Failed to create %s

Parameters

%s: Directory name

Description

Failed to create a directory during the installation.

Action method

Creating the directory might have failed for some reason. Run the installer again, or run cirinst command again. If the problem persists, check the system log.

FSP_FJSVCIR_CIRINST: ERROR: 103: CIRINST103: It is required at least %dMB

Parameters

%d: Disk spaces

Description

Insufficient disk space.

Action method

Confirm that the system has sufficient disk spaces.

[Windows] FSP_FJSVCIR_CIRINST: ERROR: 104: CIRINST104: Failed to copy CIR files.

[Linux] FSP_FJSVCIR_CIRINST: ERROR: 104: CIRINST104: Failed to copy CIR installer files.

Description

Failed to copy files required for the installation.

Action method

Copying the files failed for some reason. Run the installer again, or run cirinst command again. If the problem persists, check the system log.

FSP_FJSVCIR_CIRINST: ERROR: 105: CIRINST105: Failed to copy JRE.

Description

Failed to copy JRE required for installation.

Action method

Copying the files failed for some reason. Run the installer again, or run cirinst command again. If the problem persists, check the system log.

FSP_FJSVCIR_CIRINST: ERROR: 106: CIRINST106: Failed to copy CIR**Description**

Failed to copy files during the installation.

Action method

Copying the files failed for some reason. Run the installer again, or run cirinst command again. If the problem persists, check the system log.

FSP_FJSVCIR_CIRINST: ERROR: 107: CIRINST107: invalid option.**Description**

An invalid option was specified.

Action method

Specify a valid option. [-c]: CUI

FSP_FJSVCIR_CIRINST: ERROR: 108: CIRINST108: JRE which should be deployed is not bundled.**Description**

An appropriate JRE for this system is not bundled.

Action method

Confirm that the product supports the platform.

FSP_FJSVCIR_CIRINST: ERROR: 109: CIRINST109: Unsupported OS.**Description**

Unsupported operating system.

Action method

Confirm that the product supports the operating system.

FSP_FJSVCIR_CIRINST: ERROR: 185: CIRINST185: lsb_release was not found. Please install the package below.

- redhat-lsb

Description

lsb_release command was not found.

Action method

Add redhat-lsb package and then execute install again.

FSP_FJSVCIR_CIRINST: ERROR: 186: CIRINST186: lsb_release was not found.**Description**

lsb_release command was not found.

Action method

Add the relevant package and then execute install again.

FSP_FJSVCIR_CIRINST: ERROR: 201: CIRINST201: Administrator privilege is required.**Description**

The user does not have administrator privileges.

Action method

Login with administrator privileges and run the command again.

FSP_FJSVCIR_CIRINST: ERROR: 203: CIRINST203: Failed to delete files.

Description

Failed to delete FJSVcir files.

Action method

Deleting the files failed for some reason. Run cirremove command again. If the problem persists, check the system log.

FSP_FJSVCIR_CIRINST: ERROR: 205: CIRINST205: invalid option.

Description

Invalid option is specified.

Action method

No option is supported. Retry without any options.

Specified installation parameters file was not found.

Description

Specified installation parameters file was not found. (silent mode)

Action method

Specify the installation parameters file.

It failed to install. See log for details

Description

An error occurred during install. (silent mode)

Action method

Collect the files stored in the following directories.

[UNIX]

/var/opt/FJSVcir/cir/logs/cirlog0.log.0

[Windows]

%ProgramData%\Fujitsu\FujitsuF4CR\cir\logs\cirlog0.log.0

Parameters are invalid. Please specify valid parameters.

Description

Invalid option is specified. (silent mode)

Action method

Specify a valid option.

Internal data is broken.

Description

Internal data used by this product is corrupt.

Action method

Collect the files stored in the following directories and contact Fujitsu technical support:

[UNIX]

/var/opt/FJSVCIR/cir/

/etc/opt/FJSVCIR/cir/CIR.properties

[Windows]

%ProgramData%\Fujitsu\FujitsuF4CR\cir\

G.2 Messages output by Uninstall (middleware)

It has already been running in another process.

Description

An attempt was made to start Uninstall (middleware) more than once. Alternatively, an attempt was made to start Uninstall (middleware) while the installer was starting.

Action method

End any Uninstall (middleware) that have already started. Alternatively, start Uninstall (middleware) after the installer has ended.

Specified option is invalid.

Description

An invalid option was specified.

Action method

Specify valid options and run the command again.

Please enter a valid option.

Description

An invalid character was entered. Valid characters are y, n, or item numbers.

Action method

Enter a valid character.

Failed to uninstall of the following software packages:

Description

Failed to uninstall the software listed in this message.

Action method

Confirm the information shown after this message and take the appropriate action.

Unable to uninstall the software as it is referred from more than one software.

Description

Failed to uninstall the software because it is used by other software.

Action method

No specific action is required.

Unable to uninstall the software as the software is mandatory.

Description

Failed to uninstall the software because it is required by another program.

Action method

No specific action is required.

Internal data is broken.

Description

Internal data used by this product is corrupt.

Action method

Collect the files stored in the following directories and contact Fujitsu technical support:

[UNIX]

/var/opt/FJSVCIR/cir/

/etc/opt/FJSVCIR/cir/CIR.properties

[Windows]

%ProgramData%\Fujitsu\FujitsuF4CR\cir\

Unable to uninstall the software you were about to uninstall as the existence of the software is a precondition for the operation of another software.

Description

Failed to uninstall the software because it is required by another program.

Action method

No specific action is required.

The program terminated abnormally.

Description

The program has terminated abnormally.

Action method

Collect the files stored in the following directories and contact Fujitsu technical support:

[UNIX]

/var/opt/FJSVCIR/cir/

/etc/opt/FJSVCIR/cir/CIR.properties

[Windows]

%ProgramData%\Fujitsu\FujitsuF4CR\cir\

An unexpected error has occurred during uninstall.

Description

An error occurred during uninstall.

Action method

Collect the files stored in the following directories and contact Fujitsu technical support:

[UNIX]

/var/opt/FJSVCIR/cir/

/etc/opt/FJSVCIR/cir/CIR.properties

[Windows]

%ProgramData%\Fujitsu\FujitsuF4CR\cir\

It failed to uninstall. See log for details.

Description

An error occurred during uninstall. (silent mode)

Action method

Collect the files stored in the following directories.

[UNIX]

/var/opt/FJSVCir/cir/logs/cirlog0.log.0

[Windows]

%ProgramData%\Fujitsu\FujitsuF4CR\cir\logs\cirlog0.log.0

Failed to initialize the temp directory.

Description

Can not start Uninstall (middleware) because failed to initialize the temp directory.

Action method

Run Uninstall (middleware) again. If the problem persists, check whether other processes have accessed the files in the following directories.

[UNIX]

/var/opt/FJSVCir/cir/temp/meta_db

[Windows]

%ProgramData%\Fujitsu\FujitsuF4CR\cir\temp\meta_db

[Notice]Need to restart for uninstall completion.

Description

Uninstallation was completed. (silent mode)

Action method

Restart the system.

Appendix H Uninstall (middleware)

H.1 Features that are Installed

[Windows]

Feature	Package name	Component name	Remarks	Selectively installed?
Common tool	FJSVcir	CIRuntime Application	Controls the installation and uninstallation of Fujitsu middleware products. Manages the installation information of Fujitsu middleware products and includes a management function for viewing installation information and launching the uninstaller of each product. ("Uninstall (middleware)")	No

[Linux]

Feature	Package name	Component name	Remarks	Selectively installed?
Common tool	FJSVcir	CIRuntime Application	Controls the installation and uninstallation of Fujitsu middleware products. Manages the installation information of Fujitsu middleware products and includes a management function for viewing installation information and launching the uninstaller of each product. ("Uninstall (middleware)")	No

H.2 Notes

H.2.1 Notes on the Uninstall (middleware) Tool

"Uninstall (middleware)" is a tool used by Fujitsu middleware products. It performs information management for middleware products, and also, launching the uninstaller for removing these products from the installed systems.



Note

- To uninstall Enterprise Postgres, uninstall from "Uninstall (middleware)".

This tool manages information about other Fujitsu middleware products even if Enterprise Postgres. For this reason, do not uninstall it unless absolutely necessary. If you have accidentally uninstalled this tool, re-install it as described below.

[Windows]

1. Log on to the machine to be installed using a user name belonging to the Administrators group, or switch to an account with administrator privileges.
2. Insert the server program DVD in the drive device.
3. Execute the installation command.

```
z:\CIR\Windows\cirinst.exe
```

In the example above, z is the drive on which the DVD is mounted.

[Linux]

1. Log in as a superuser to the installation target machine, or switch to a user with administrative privileges.
2. Insert the product media in the drive.
3. Execute the installation command.

```
# /media/dvd/CIR/Linux/cirinst.sh
```

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

- If you uninstall this tool as described below.

1. Check that there are no Fujitsu middleware products installed in the system.
You can check this by using "Uninstall (middleware)", which is launched by the command below.

[Windows]

Select [All Programs] > [Fujitsu] > [Uninstall (middleware)]

[Linux]

```
/opt/FJSVcir/cir/bin/cimanager.sh -c
```

-c : CUI

2. If no Fujitsu middleware products have been installed, then issue the command below.

[Windows]

```
%SystemDrive%\FujitsuF4CR\bin\cirremove.exe
```

[Linux]

```
# /opt/FJSVcir/bin/cirremove.sh
```

3. At the uninstallation confirmation prompt, enter "y".

```
This software is a common tool of Fujitsu products.Are you sure you want to remove it?[y/n]:
```

Uninstallation will start.

Appendix I Estimating Database Disk Space Requirements

This appendix describes how to estimate database disk space requirements.

I.1 Estimating Table Size Requirements

The following tables provide the formulas for estimating table size requirements.

Table I.1 Estimation formula when the record length is 2032 bytes or less

Item	Estimation formula (bytes)
(1) Record length	<p>$27(*1) + \text{NULL map} + \text{OID} + \text{column data}$</p> <p>NULL map: Number of columns / 8 (*2) OID: 4 for a table "WITH OID". The default is 0. Column data: Sum of column lengths</p> <p>*1: Record head section *2: Round the result up to the next integer.</p> <ul style="list-style-type: none"> - Because the column data is placed in boundaries of 8 bytes, you need to make an adjustment so that the sum of the record head section, NULL map and OID is a multiple of 8. For example, if the calculated length is $27 + 1 / 8$ (rounded up) + 0 = 28 bytes, add 4 to make the length 32 bytes. - Because the data of each column is placed in boundaries of the defined data type, take the boundary of each data type into account for the length of the column data. For example, the length of the column data in the table below will not be the sum of the data types, which is 37 bytes, but will instead be 64 bytes following boundary adjustment. Definition: create table tb1(c1 char(1), c2 long, c3 int, c4 box) Estimation: CHAR type 1 byte + boundary adjustment of 7 bytes for LONG type 8 bytes + LONG type 8 bytes + INT type 4 bytes + boundary adjustment of 12 bytes for BOX type 32 bytes + BOX type 32 bytes = 64 bytes - Because each record is placed in boundaries of 8 bytes, you need to make an adjustment so that the length of the column data is a multiple of 8. - If the calculated record length exceeds 2,032 bytes, the variable length data in the record might be compressed automatically. If so, use the estimation formulas in "Table I.2 Estimation formula when the record length exceeds 2032 bytes" to estimate the table size.
(2) Number of records per page	<p>$8168 (*1) / ((1) \text{ record length} + 4 (*2))$</p> <p>*1: Page length (8192) - page head (24) *2: Pointer length (4)</p> <ul style="list-style-type: none"> - The result will be rounded down to the next integer.
(3) Number of pages required for storing records	<p>Total number of records / (2) number of records per page</p> <ul style="list-style-type: none"> - The result will be rounded up to the next integer.
(4) Amount of space	<p>(3) Number of pages required for storing records x page length x safety factor (*1)</p> <p>*1: Specify 2.0 or higher.</p> <ul style="list-style-type: none"> - This is the safety factor assumed if vacuuming is performed for garbage collection in tables and indexes.

Table I.2 Estimation formula when the record length exceeds 2032 bytes

Item	Estimation formula (bytes)
(4) Amount of space	Total number of records x (1) record length x safety factor (*1) *1: Specify 2.0 or higher. - This is the safety factor assumed if vacuuming is performed for garbage collection in tables and indexes.

I.2 Estimating Index Size Requirements

This section provides the formulas for estimating index size requirements.

Enterprise Postgres provides five index types: B-tree, Hash, GiST, GIN, and SP-GiST (the same as PostgreSQL). If you do not specify the index type in the CREATE INDEX statement, a B-tree index is generated.

The following describes how to estimate a B-tree index.

A B-tree index is saved as a fixed-size page of 8 KB. The page types are meta, root, leaf, internal, deleted, and empty. Since leaf pages usually account for the highest proportion of space required, you need to calculate the requirements for these only.

Table I.3 Estimation formula when the key data length is 512 bytes or less

Item	Estimation formula (bytes)
(1) Entry length	8 (*1) + key data length (*2) *1: Entry head *2: The key data length depends on its data type (refer to "I.3 Sizes of Data Types" for details). Because each entry is placed in boundaries of 8 bytes, you need to make an adjustment so that the length of the key data is a multiple of 8. For example, if the calculated length is 28 bytes, add 4 to make the length 32 bytes. - If the key data length exceeds 512 bytes, key data may be automatically compressed. In this case, use the estimation formula given in "Table I.4 Estimation formula when the key data length exceeds 512 bytes" to estimate the key data length.
(2) Page size requirement	8192 (*1) *1: Page length (8192) - page header (24) - special data (16) = 8152
(3) Number of entries per page	(2) Page size requirement / ((1) entry length + 4 (*1)) *1: Pointer length - Result of (3) number of entries per page will be rounded down to the nearest integer.
(4) Number of pages required for storing indexes	Total number of records / (3) number of entries per page - Result of (4) number of pages required for storing indexes will be rounded up to the nearest integer.
(5) Space requirement	(4) Number of pages required for storing indexes x 8192 (*1) / usage rate (*2) *1: Page length *2: Specify 0.7 or lower.

Table I.4 Estimation formula when the key data length exceeds 512 bytes

Item	Estimation formula (bytes)
(5) Space requirement	Total number of records x key data length x compression ratio (*1) / usage rate (*2) *1: The compression ratio depends on the data value, so specify 1. *2: Specify 0.7 or lower as the usage rate.

I.3 Sizes of Data Types

This section lists the sizes of the data types.

I.3.1 Sizes of Fixed-Length Data Types

The following table lists the sizes of fixed-length data types.

Data type	Size (bytes)
SMALLINT (INT2)	2
INTEGER (INT4)	4
BIGINT (INT8)	8
REAL	4
DOUBLE PRECISION	8
SERIAL (SERIAL4)	4
BIGSERIAL (SERIAL8)	8
MONEY	8
FLOAT	8
FLOAT (1-24)	4
FLOAT (25-53)	8
TIMESTAMP WITHOUT TIME ZONE	8
TIMESTAMP WITH TIME ZONE	8
DATE	4
TIME WITHOUT TIME ZONE	8
TIME WITH TIME ZONE	12
INTERVAL	12
BOOLEAN	1
CIDR	IPv4: 7 IPv6: 19
INET	IPv4: 7 IPv6: 19
MACADDR	6
POINT	16
LINE	32
LSEG	32
BOX	32
CIRCLE	24

I.3.2 Sizes of Variable-Length Data Types

The following table lists the sizes of variable-length data types.

Data type	Size (bytes)	Remarks
path	Length of size portion + 12 + 16 x number of vertices	1) When carrying out division, round to the next integer.

Data type	Size (bytes)	Remarks
polygon	Length of size portion + 36 + 16 x number of vertices	2) If the real data length is less than 127, then the length of the size portion is 1 byte, otherwise it is 4 bytes. 3) The number of bytes per character depends on the character set (refer to " I.3.4 Number of Bytes per Character " for details).
decimal	Length of size portion + 2 + (integer precision / 4 + decimal precision / 4) x 2	
numeric		
bytea	Length of size portion + real data length	
character varying(n), varchar(n)	Length of size portion + number of characters x number of bytes per character	
character(n), char(n)	Length of size portion + n x number of bytes per character	
text	Length of size portion + number of characters x number of bytes per character	

I.3.3 Sizes of Array Data Types

The following table lists the sizes of array data types.

Data type	Size (bytes)	Remarks
Array	Length of size portion + 12 + 8 x number of dimensions + data size of each item	If the real data length is less than 127, then the length of the size portion is 1 byte, otherwise it is 4 bytes. - Example of estimation when array data is "ARRAY[[1,2,3],[1,2,3]]" Number of dimensions: 2 INTEGER data size: 4 Total size = 1+12+8x2+6x4 = 53

I.3.4 Number of Bytes per Character

The following table lists the number of bytes per character.

The given values relate to the common character sets EUC-JP and UTF8.

Character type	Character set	Number of bytes per character
ASCII	EUC_JP	1
Halfwidth katakana	EUC_JP	2
JIS X 0208 kanji characters	EUC_JP	2
JIS X 0212 kanji characters	EUC_JP	3
ASCII	UTF8	1
Halfwidth katakana	UTF8	3
JIS X 0208 kanji characters	UTF8	3
JIS X 0212 kanji characters	UTF8	3

I.4 Estimating Transaction Log Space Requirements

This section provides the formula for estimating transaction log space requirements.

```
Transaction log space requirements = ( checkpoint_segments x 3 + 1 ) x 16 MB
```

However, if the update volume is extremely high (for example, due to a large data load and batch processing), disk writing at a checkpoint may not be able to keep up with the load, and a higher number of transaction logs than indicated here may temporarily be accumulated.

I.5 Estimating Archive Log Space Requirements

This section explains how to estimate archive log space requirements.

The archive log is an archive of the transaction logs from the time of a previous backup to the present, so it fluctuates depending on the backup period and the content of update transactions.

The longer the backup period and the more update transactions, the greater the space required for the archive log.

Therefore, measure the actual archive log space by using a test environment to simulate backup scheduling and database update in a real operating environment.

I.6 Estimating Backup Disk Space Requirements

This section provides the formula for estimating backup disk space requirements.

```
Backup disk space requirements = size of the database cluster + transaction log space requirements  
+ archive log space requirements
```

Appendix J Estimating Memory Requirements

This appendix explains how to estimate the memory.

J.1 Enterprise Postgres Memory Requirements

This section describes the formulas for estimating Enterprise Postgres memory requirements.

Use the following formula to obtain a rough estimate of memory required for Enterprise Postgres:

$$enterprisePostgresRequiredMemory = sharedMemoryAmount + LocalMemoryAmount$$

Shared memory amount

Refer to "Shared Memory and Semaphores" under "Server Administration" in the PostgreSQL Documentation for information on shared memory.

However, note that if instances have been created using WebAdmin, the parameters below will be configured automatically when the instances are created. Take this into account when calculating the shared memory size.

Parameter name	Set value
shared_buffers	30 percent of the internal memory of the machine.
max_connections	100
max_prepared_transactions	100

Local memory amount

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

Process stack area

$$\begin{aligned}
processStackArea \\
= & max_stack_depth \times (max_connections + autovacuum_max_workers + 9)
\end{aligned}$$

This formula evaluates to the maximum value.

Actually it is used according to the growth of the stack.

In the formula above, 9 is the number of processes that perform roles specific to servers.

Memory used in database sessions that use temporary tables

$$\begin{aligned}
memoryUsedInDbSessionsThatUseTempTables \\
= & temp_buffers \times max_connections
\end{aligned}$$

This formula evaluates to the maximum value.

Memory is gradually used as temporary buffers are used, and is released when the session ends.

Memory used in database sessions that perform sort and hash table operations

$$\begin{aligned}
memoryUsedInDbSessionsThatPerformSortAndHashTableOperations \\
= & work_mem \times max_connections
\end{aligned}$$

This formula evaluates to the maximum value.

Memory is gradually used as operations such as sort are performed, and is released when the query ends.

Memory used in maintenance operations

$$\begin{aligned} & \text{memoryUsedInMaintenanceOperations} \\ & = \text{maintenance_work_mem} \times (\text{numOfSessionsPerformingMaintenance} + \text{autovacuum_max_workers}) \end{aligned}$$

Note that 'maintenance operations' are operations such as VACUUM, CREATE INDEX, and ALTER TABLE ADD FOREIGN KEY.

Base memory used in each process

$$\begin{aligned} & \text{baseMemoryUsedInEachProcess} \\ & = 3\text{MB} \times (\text{max_connections} + \text{autovacuum_max_workers} + 9) \end{aligned}$$

This formula evaluates to the memory used when server processes are running.

In the formula above, 9 is the number of processes that perform roles specific to servers.

Memory used preparing for data access

$$\begin{aligned} & \text{memoryUsedPreparingForDataAccess} \\ & = \text{variationAmount} \times (\text{max_connections} + \text{autovacuum_max_workers} + 4) \end{aligned}$$

where $\text{variationAmount} = \text{shared_buffers} / 8\text{KB} \times 4 \text{ bytes}$

(note that 8KB is the page length, and 4 bytes is the size of page management data)

This formula evaluates to the memory required to access the database cache in the shared memory.

In the formula above, among the processes that perform roles specific to servers, 4 is the number of processes that access the database.

J.2 Database Multiplexing Memory Requirements

This section describes the formula for estimating database multiplexing memory requirements.

Use the following formula to obtain a rough estimate of memory required for database multiplexing:

$$\begin{aligned} & \text{Memory usage of the database multiplexing feature} \\ & = \text{Peak memory usage of the Mirroring Controller processes} \\ & \quad + \text{Peak memory usage of the Mirroring Controller commands} \\ & \text{Peak memory usage of the Mirroring Controller processes} = 150 \text{ MB} \\ & \text{Peak memory usage of the Mirroring Controller commands} = 50 \text{ MB} \times \text{Number of commands executed} \\ & \text{simultaneously} \end{aligned}$$

J.3 PL/extJava Memory Requirements

This section describes the formula for estimating PL/extJava memory requirements.

Use the following formula to obtain a rough estimate of memory required for PL/extJava:

$$\text{memoryUsedByPLExtJavaInMegabytes} = 960 + (\text{numOfServerInstances(JavaVMs)} \times 47)$$



Note

- Creating a container will create one server instance (Java VM).
- The number of server instances (Java VMs) may fluctuate based on the following actions:
 - Adding a container (one server instance (Java VM) will be added)
 - Adding a server instance (Java VM) to a container
 - Deleting a server instance (Java VM) from a container

Appendix K Quantitative Limits

This appendix lists the quantitative limits of Enterprise Postgres.

Table K.1 Length of identifier

Item	Limit
Database name	Up to 63 bytes (*1)(*2)(*3)
Schema name	Up to 63 bytes (*1)(*2)
Table name	Up to 63 bytes (*1)(*2)
View name	Up to 63 bytes (*1)(*2)
Index name	Up to 63 bytes (*1)(*2)
Table space name	Up to 63 bytes (*1)(*2)
Cursor name	Up to 63 bytes (*1)(*2)
Function name	Up to 63 bytes (*1)(*2)
Aggregate function name	Up to 63 bytes (*1)(*2)
Trigger name	Up to 63 bytes (*1)(*2)
Constraint name	Up to 63 bytes (*1)(*2)
Conversion name	Up to 63 bytes (*1)(*2)
Role name	Up to 63 bytes (*1)(*2)(*4)
Cast name	Up to 63 bytes (*1)(*2)
Collation sequence name	Up to 63 bytes (*1)(*2)
Encoding method conversion name	Up to 63 bytes (*1)(*2)
Domain name	Up to 63 bytes (*1)(*2)
Extension name	Up to 63 bytes (*1)(*2)
Operator name	Up to 63 bytes (*1)(*2)
Operator class name	Up to 63 bytes (*1)(*2)
Operator family name	Up to 63 bytes (*1)(*2)
Rewrite rule name	Up to 63 bytes (*1)(*2)
Sequence name	Up to 63 bytes (*1)(*2)
Text search settings name	Up to 63 bytes (*1)(*2)
Text search dictionary name	Up to 63 bytes (*1)(*2)
Text search parser name	Up to 63 bytes (*1)(*2)
Text search template name	Up to 63 bytes (*1)(*2)
Data type name	Up to 63 bytes (*1)(*2)
Enumerator type label	Up to 63 bytes (*1)(*2)

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

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

*3: Names of databases that use PL/extJava must be 28 bytes or less.

*4: Names of roles for connecting from the PL/extJava application server must be 8 bytes or less.

Table K.2 Database object

Item	Limit
Number of databases	Less than 4,294,967,296 (*1)
Number of schemas	Less than 4,294,967,296 (*1)
Number of tables	Less than 4,294,967,296 (*1)
Number of views	Less than 4,294,967,296 (*1)
Number of indexes	Less than 4,294,967,296 (*1)
Number of table spaces	Less than 4,294,967,296 (*1)
Number of functions	Less than 4,294,967,296 (*1)
Number of aggregate functions	Less than 4,294,967,296 (*1)
Number of triggers	Less than 4,294,967,296 (*1)
Number of constraints	Less than 4,294,967,296 (*1)
Number of conversion	Less than 4,294,967,296 (*1)
Number of roles	Less than 4,294,967,296 (*1)
Number of casts	Less than 4,294,967,296 (*1)
Number of collation sequences	Less than 4,294,967,296 (*1)
Number of encoding method conversions	Less than 4,294,967,296 (*1)
Number of domains	Less than 4,294,967,296 (*1)
Number of extensions	Less than 4,294,967,296 (*1)
Number of operators	Less than 4,294,967,296 (*1)
Number of operator classes	Less than 4,294,967,296 (*1)
Number of operator families	Less than 4,294,967,296 (*1)
Number of rewrite rules	Less than 4,294,967,296 (*1)
Number of sequences	Less than 4,294,967,296 (*1)
Number of text search settings	Less than 4,294,967,296 (*1)
Number of text search dictionaries	Less than 4,294,967,296 (*1)
Number of text search parsers	Less than 4,294,967,296 (*1)
Number of text search templates	Less than 4,294,967,296 (*1)
Number of data types	Less than 4,294,967,296 (*1)
Number of enumerator type labels	Less than 4,294,967,296 (*1)
Number of default access privileges defined in the ALTER DEFAULT PRIVILEGES statement	Less than 4,294,967,296 (*1)
Number of large objects	Less than 4,294,967,296 (*1)
Number of rows in tables defined by WITH OIDS	Less than 4,294,967,296 (*1)
Number of index access methods	Less than 4,294,967,296 (*1)

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

Table K.3 Schema element

Item	Limit
Number of columns that can be defined in one table	From 250 to 1600 (according to the data type)
Table row length	Up to 400 gigabytes

Item	Limit
Number of columns comprising a unique constraint	Up to 32 columns
Data length comprising a unique constraint	Less than 2,000 bytes (*1)(*2)
Table size	Up to 32 terabyte
Search condition character string length in a trigger definition statement	Up to 800 megabytes (*1)(*2)
Item size	Up to 1 gigabyte

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

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

Table K.4 Index

Item	Limit
Number of columns comprising a key (B-tree, GiST, GIN index)	Up to 32 columns
Key length	Less than 2,000 bytes (*1)

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

Table K.5 Data types and attributes that can be handled

Item	Limit		
Character	Data length	Data types and attributes that can be handled (*1)	
	Specification length (n)	Up to 10,485,760 characters (*1)	
Numeric	External decimal expression	Up to 131,072 digits before the decimal point, and up to 16,383 digits after the decimal point	
	Internal binary expression	2 bytes	From -32,768 to 32,767
		4 bytes	From -2,147,483,648 to 2,147,483,647
		8 bytes	From -9,223,372,036,854,775,808 to 9,223,372,036,854,775,807
	Internal decimal expression		Up to 13,1072 digits before the decimal point, and up to 16,383 digits after the decimal point
	Floating point expression	4 bytes	From -3.4E+38 to -7.1E-46, 0, or from 7.1E-46 to 3.4E+38
		8 bytes	From -1.7E+308 to -2.5E-324, 0, or from 2.5E-324 to 1.7E+308
bytea		Up to one gigabyte minus 53 bytes	
Large object		Up to two gigabytes	

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

Table K.6 Function definition

Item	Limit
Number of arguments that can be specified	Up to 100
Number of variable names that can be specified in the declarations section	No limit

Item	Limit
Number of SQL statements or control statements that can be specified in a function processing implementation	No limit

Table K.7 Data operation statement

Item	Limit
Maximum number of connections for one process in an application (remote access)	4,000 connections
Number of expressions that can be specified in a selection list	Up to 1,664
Number of tables that can be specified in a FROM clause	No limit
Number of unique expressions that can be specified in a selection list/DISTINCT clause/ORDER BY clause/GROUP BY clause within one SELECT statement	Up to 1,664
Number of expressions that can be specified in a GROUP BY clause	No limit
Number of expressions that can be specified in an ORDER BY clause	No limit
Number of SELECT statements that can be specified in a UNION clause/INTERSECT clause/EXCEPT clause	Up to 4,000 (*1)
Number of nestings in joined tables that can be specified in one view	Up to 4,000 (*1)
Number of functions or operator expressions that can be specified in one expression	Up to 4,000 (*1)
Number of expressions that can be specified in one row constructor	Up to 1,664
Number of expressions that can be specified in an UPDATE statement SET clause	Up to 1,664
Number of expressions that can be specified in one row of a VALUES list	Up to 1,664
Number of expressions that can be specified in a RETURNING clause	Up to 1,664
Total expression length that can be specified in the argument list of one function specification	Up to 800 megabytes (*2)
Number of cursors that can be processed simultaneously by one session	No limit
Character string length of one SQL statement	Up to 800 megabytes (*1) (*3)
Number of input parameter specifications that can be specified in one dynamic SQL statement	No limit
Number of tokens that can be specified in one SQL statement	Up to 10,000
Number of values that can be specified as a list in a WHERE clause IN syntax	No limit
Number of expressions that can be specified in a USING clause	No limit
Number of JOINS that can be specified in a joined table	Up to 4,000 (*1)
Number of expressions that can be specified in COALESCE	No limit
Number of WHEN clauses that can be specified for CASE in a simple format or a searched format	No limit
Data size per record that can be updated or inserted by one SQL statement	Up to one gigabyte minus 53 bytes
Number of objects that can share a lock simultaneously	Up to 256,000 (*1)

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

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

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

Table K.8 Data size

Item	Limit
Data size per record for input data files (COPY statement, psql command \copy meta command)	Up to 800 megabytes (*1)
Data size per record for output data files (COPY statement, psql command \copy meta command)	Up to 800 megabytes (*1)

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

Appendix L Configuring Kernel Parameters

Use the "System V IPC Parameters" table in "Managing Kernel Resources" in the PostgreSQL Documentation for the relationship between configuration parameters and kernel parameters, as well as calculation formulas.

Refer to the "Managing Kernel Resources" in the PostgreSQL Documentation to calculate shared memory usage.

For multiple instances, the kernel parameters should be evaluated for all instances. For example, in the case of the maximum number of shared memory segments for the entire system (SHMMNI), the total number of segments obtained by all instances should be added to the kernel parameters. In the case of the maximum number of semaphores for each process (SEMMSL), the largest of all sizes obtained by all instances should be compared to the current value prior to configuring the settings.



Note

If there is insufficient shared memory due to miscalculation of SHMMAX, a message will be output indicating that the shmget system call failed at "errno=22 (EINVAL)". Review the calculation, and reconfigure.

The relationship between System V IPC parameters and kernel parameters in various operating systems is shown below.

System	V IPC parameter	Kernel parameter action
SHMMAX	kernel.shmmax	If <i>currentValue</i> < <i>calculatedValue</i> , configure the calculated value
SHMMIN	No compatible parameter	
SHMALL	kernel.shmall	Specify <i>currentValue</i> + <i>calculatedValue</i>
SHMSEG	No compatible parameter	
SHMMNI	kernel.shmmni	Specify <i>currentValue</i> + <i>calculatedValue</i>
SEMNI	Fourth parameter of kernel.sem	Specify <i>currentValue</i> + <i>calculatedValue</i>
SEMNS	Second parameter of kernel.sem	Specify <i>currentValue</i> + <i>calculatedValue</i>
SEMMSL	First parameter of kernel.sem	If <i>currentValue</i> < <i>calculatedValue</i> , configure the calculated value
SEMAP	No compatible parameter	
SEMVMX	No compatible parameter	

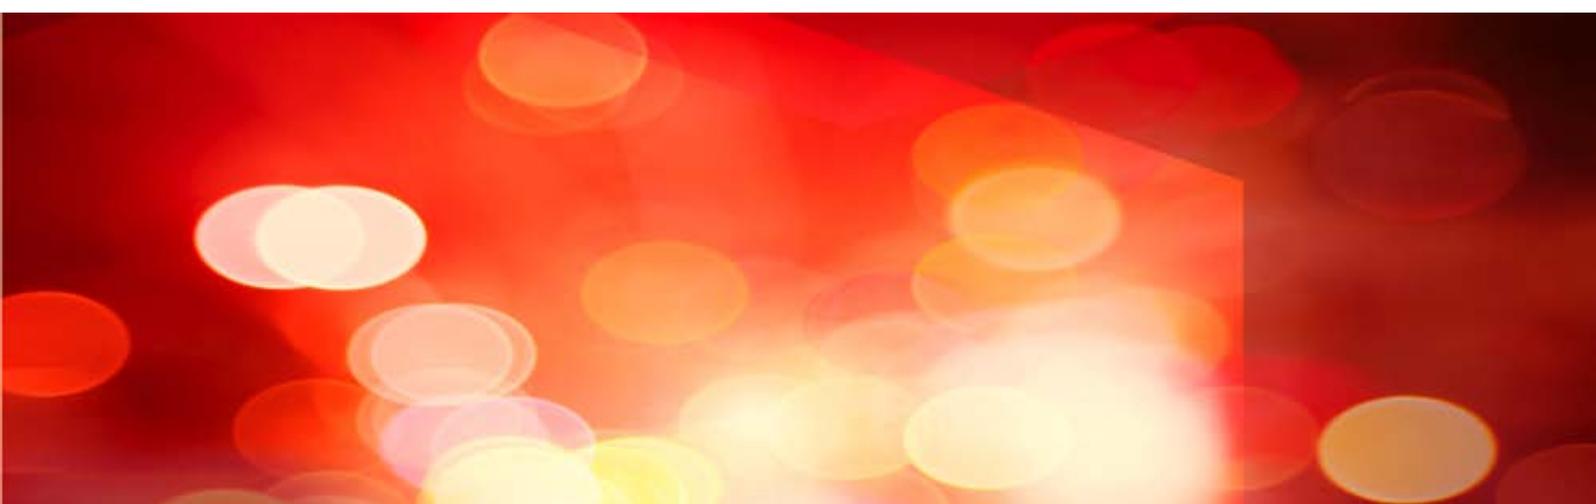
Remark 1: kernel.shmall is the page specification.

Remark 2: Specify all four parameters for kernel.sem. At this time, the value specified in the third parameter should be the same value as before configuration.

Index

	[C]		[W]
Changing allocated memory.....	31	WebAdmin automatic start.....	59
Changing client authentication.....	26	Web server port number.....	58
Changing communication.....	28	When an Instance was Created with the initdb Command.....	35
Changing the character set.....	25	When an Instance was Created with WebAdmin.....	35
Changing the settings.....	25		
Changing the SQL option.....	30		
Check the disk space.....	7		
Check the installed product and determine the installation method	7		
Client Authentication Information settings.....	35		
Creating an Instance Administrator.....	15		
	[D]		
Disk Space Required for Installation.....	5		
	[E]		
Excluded Software.....	4		
	[H]		
Hardware Environment.....	5		
How to Set Up the Pop-up Blocker	57		
	[I]		
Installation Procedure.....	1		
Installation Types.....	1		
	[M]		
Migration Procedures when Performing Failover Operations using the Standby Feature.....	52		
	[N]		
New Installation.....	1		
	[O]		
Operating Environment.....	3		
Operating Method Types and Selection.....	14		
	[P]		
Port number to use when Tomcat is stopped.....	59		
postgresql.conf.....	61		
Pre-installation Tasks.....	7		
Preparations for Setup.....	15		
	[R]		
Recommended Browser Settings.....	57		
Related Software.....	4		
Remove applied updates.....	7		
Required Operating System.....	3		
Required Patches.....	4		
	[S]		
Settings related to connection.....	35		
	[T]		
TCP/IP Protocol.....	5		

FUJITSU Enterprise Postgres 9.4



Installation and Setup Guide for Server

Windows

B1WS-1221-01ENZO(00)
September 2015

Preface

Purpose of This Document

The Enterprise Postgres database system extends the PostgreSQL features and runs on the Windows platform.

This document describes how to install and setup "Enterprise Postgres".

Intended Readers

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

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

- PostgreSQL
- SQL
- Windows

Structure of This Document

This document is structured as follows:

[Chapter 1 Overview of Installation](#)

Describes the installation types and procedures

[Chapter 2 Operating Environment](#)

Describes the operating environment required to use Enterprise Postgres

[Chapter 3 Installation](#)

Describes how to perform a new installation of Enterprise Postgres

[Chapter 4 Setup](#)

Describes the setup to be performed after installation

[Chapter 5 Uninstallation](#)

Describes how to uninstall Enterprise Postgres

[Appendix A Installation in Silent Mode](#)

Provides specifications for installation in silent mode

[Appendix B Enterprise Postgres Migration Procedures](#)

Describes how to migrate Enterprise Postgres from an old version/level.

[Appendix C Using WebAdmin After an Upgrade](#)

Describes the required operations for using WebAdmin after an upgrade.

[Appendix D Recommended WebAdmin Environments](#)

Describes the recommended WebAdmin environment.

[Appendix E Setting Up and Removing WebAdmin](#)

Describes how to set up and remove WebAdmin

[Appendix F Configuring Parameters](#)

Describes Enterprise Postgres parameters.

[Appendix G Uninstall \(middleware\)](#)

Describes the Uninstall (middleware)

[Appendix H Uninstall \(middleware\) Messages](#)

Explains the messages output by the Uninstall (middleware) tool.

[Appendix I Estimating Database Disk Space Requirements](#)

Describes how to estimate database disk space requirements

[Appendix J Estimating Memory Requirements](#)

Describes the formulas for estimating memory requirements

[Appendix K Quantitative Limits](#)

Describes the quantity range

Export Restrictions

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

Issue Date and Version

First Edition: September 2015

Copyright Notice

Copyright 2015 FUJITSU LIMITED

Contents

Chapter 1 Overview of Installation.....	1
1.1 Features that can be Installed.....	1
1.2 Installation Types.....	1
1.2.1 New Installation.....	1
1.2.2 Reinstallation.....	1
1.2.3 Upgrade Installation.....	1
1.2.4 Multi-Version Installation.....	1
1.3 Installation Procedure.....	1
1.3.1 Installation in Interactive Mode.....	2
1.3.2 Installation in Silent Mode.....	2
1.4 Uninstallation.....	2
Chapter 2 Operating Environment.....	3
2.1 Required Operating System.....	3
2.2 Related Software.....	4
2.3 Excluded Software.....	5
2.4 Required Patches.....	5
2.5 Hardware Environment.....	5
2.6 Supported System Environment.....	5
2.6.1 TCP/IP Protocol.....	5
2.6.2 File System.....	6
2.7 Enterprise Postgres-based PostgreSQL Versions.....	6
2.8 Notes on Using Streaming Replication.....	6
Chapter 3 Installation.....	7
3.1 Pre-installation Tasks.....	9
3.2 Installation in Interactive Mode.....	11
3.3 Installation in Silent Mode.....	17
3.4 Post-installation Tasks.....	18
Chapter 4 Setup.....	19
4.1 Operating Method Types and Selection.....	19
4.2 Preparations for Setup.....	20
4.2.1 Creating an Instance Administrator.....	20
4.2.1.1 Security policy settings.....	21
4.2.2 Preparing Directories for Resource Deployment.....	21
4.2.3 Estimating Resources.....	26
4.2.4 Windows Firewall Settings.....	26
4.2.5 Preparing for Output to the Event Log.....	27
4.3 Creating an Instance.....	28
4.3.1 Using WebAdmin.....	29
4.3.1.1 Logging into WebAdmin.....	29
4.3.1.2 Creating an Instance.....	30
4.3.1.3 Changing the Settings.....	35
4.3.2 Using the initdb Command.....	43
4.3.2.1 Creating an Instance.....	44
4.4 Configuring Remote Connections.....	46
4.4.1 When an Instance was Created with WebAdmin.....	46
4.4.2 When an Instance was Created with the initdb Command.....	47
4.4.3 Windows Firewall Settings.....	47
4.5 Other Settings.....	48
4.5.1 Error Log Settings.....	49
4.5.2 Configuring Automatic Start and Stop of an Instance.....	49
4.5.3 Settings when Using the features compatible with Oracle databases.....	50
4.6 Integration with Message-Monitoring Software.....	50

4.7 Deleting Instances.....	51
4.7.1 Using WebAdmin.....	51
4.7.2 Using Server Commands.....	51
Chapter 5 Uninstallation.....	53
5.1 Uninstallation in Interactive Mode.....	53
5.2 Uninstallation in Silent Mode.....	58
Appendix A Installation in Silent Mode.....	61
A.1 Specification Format.....	61
A.2 List of Parameters.....	61
A.3 Messages and Return Values.....	64
A.4 CSV File Format.....	65
Appendix B Enterprise Postgres Migration Procedures.....	67
B.1 Migration Tasks for Basic Operations.....	67
Appendix C Using WebAdmin After an Upgrade.....	68
Appendix D Recommended WebAdmin Environments.....	70
D.1 Recommended Browser Settings	70
D.2 How to Set Up the Pop-up Blocker.....	70
Appendix E Setting Up and Removing WebAdmin.....	71
E.1 Setting Up WebAdmin.....	71
E.1.1 Setting Up WebAdmin.....	71
E.1.2 Activating the Web Server Feature of WebAdmin.....	73
E.1.3 Stopping the Web Server Feature of WebAdmin.....	73
E.2 Removing WebAdmin.....	73
Appendix F Configuring Parameters.....	75
Appendix G Uninstall (middleware).....	78
G.1 Features that are Installed.....	78
G.2 Notes.....	78
G.2.1 Notes on the Uninstall (middleware) Tool.....	78
Appendix H Uninstall (middleware) Messages.....	80
H.1 Messages output by FJsvcir.....	80
H.2 Messages output by Uninstall (middleware).....	83
Appendix I Estimating Database Disk Space Requirements.....	86
I.1 Estimating Table Size Requirements.....	86
I.2 Estimating Index Size Requirements.....	87
I.3 Sizes of Data Types.....	88
I.3.1 Sizes of Fixed-Length Data Types.....	88
I.3.2 Sizes of Variable-Length Data Types.....	88
I.3.3 Sizes of Array Data Types.....	89
I.3.4 Number of Bytes per Character.....	89
I.4 Estimating Transaction Log Space Requirements.....	90
I.5 Estimating Archive Log Space Requirements.....	90
I.6 Estimating Backup Disk Space Requirements.....	90
Appendix J Estimating Memory Requirements.....	91
J.1 Enterprise Postgres Memory Requirements.....	91
J.2 Database Multiplexing Memory Requirements.....	92
J.3 PL/extJava Memory Requirements.....	92
Appendix K Quantitative Limits.....	93
Index.....	98

Chapter 1 Overview of Installation

This chapter provides an overview of Enterprise Postgres installation.

1.1 Features that can be Installed

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

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

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

Y: Can be installed

N: Cannot be installed

1.2 Installation Types

The following four installation types are available for Enterprise Postgres:

- New installation
- Reinstallation

1.2.1 New Installation

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

1.2.2 Reinstallation

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

1.2.3 Upgrade Installation

Perform upgrade installation to upgrade installed program files to the latest version.



Note

When performing an upgrade installation, caution regarding WebAdmin is advised. Refer to "[Appendix C Using WebAdmin After an Upgrade](#)" for details.

1.2.4 Multi-Version Installation

Perform multi-version installation to install different versions to the installed program files separately.

1.3 Installation Procedure

The following installation procedures are available for Enterprise Postgres:

- Installation in interactive mode
- Installation in silent mode

Select the installation procedure that corresponds to your environment.



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

1.3.1 Installation in Interactive Mode

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

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

- New installation
- Reinstallation

1.3.2 Installation in Silent Mode

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

New installations can be performed in silent mode.

1.4 Uninstallation

Uninstallation removes the system files of the installed Enterprise Postgres.

Chapter 2 Operating Environment

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



See

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

2.1 Required Operating System

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

Table 2.1 Operating systems

Operating system name	64-bit product	32-bit product
- Microsoft(R) Windows Server(R) 2008 Standard	Y	Y
- Microsoft(R) Windows Server(R) 2008 Enterprise		
- Microsoft(R) Windows Server(R) 2008 Datacenter		
- Microsoft(R) Windows Server(R) 2008 Standard without Hyper-V		
- Microsoft(R) Windows Server(R) 2008 Enterprise without Hyper-V		
- Microsoft(R) Windows Server(R) 2008 Datacenter without Hyper-V		
- Microsoft(R) Windows Server(R) 2008 R2 Datacenter	Y	Y (*2)
- Microsoft(R) Windows Server(R) 2008 R2 Enterprise		
- Microsoft(R) Windows Server(R) 2008 R2 Standard		
- Microsoft(R) Windows Server(R) 2008 R2 Foundation		
- Microsoft(R) Windows Server(R) 2012 Datacenter	Y	Y (*2)
- Microsoft(R) Windows Server(R) 2012 Standard		
- Microsoft(R) Windows Server(R) 2012 Essentials		
- Microsoft(R) Windows Server(R) 2012 Foundation		
- Microsoft(R) Windows Server(R) 2012 R2 Datacenter	Y	Y (*2)
- Microsoft(R) Windows Server(R) 2012 R2 Standard		
- Microsoft(R) Windows Server(R) 2012 R2 Foundation		

Y: Can be used

N: Cannot be used

*1: If using database multiplexing mode, use Windows Server(R) 2008 or later.

*2: Enterprise Postgres runs in a 32-bit environment and also can run on the Server Core installation option supported for Windows Server(R) 2008 or later.



Note

- The TCP/IP protocol must be installed.

- One of the following basic software products is required if performing failover that is linked to the cluster software of the operating system:
 - Microsoft(R) Windows Server(R) 2008 Enterprise
 - Microsoft(R) Windows Server(R) 2008 Datacenter
 - Microsoft(R) Windows Server(R) 2008 Enterprise without Hyper-V
 - Microsoft(R) Windows Server(R) 2008 Datacenter without Hyper-V
 - Microsoft(R) Windows Server(R) 2008 R2 Datacenter
 - Microsoft(R) Windows Server(R) 2008 R2 Enterprise
 - Microsoft(R) Windows Server(R) 2012 Datacenter
 - Microsoft(R) Windows Server(R) 2012 Standard
 - Microsoft(R) Windows Server(R) 2012 Essentials
 - Microsoft(R) Windows Server(R) 2012 Foundation
 - Microsoft(R) Windows Server(R) 2012 R2 Datacenter
 - Microsoft(R) Windows Server(R) 2012 R2 Standard
 - Microsoft(R) Windows Server(R) 2012 R2 Foundation

2.2 Related Software

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

Table 2.2 Related software

No.	Software name	Version	Product name of Enterprise Postgres	Remarks
			SE	
1	Enterprise Postgres Mirroring Controller	9.4	N	[Database multiplexing enhancement feature] Required for automation of the switch/disconnection of the database multiplexing feature.
2	Visual Studio	2008 2010 2012 2013	Y	-
3	.NET Framework	3.5 SP1 or later 4 4.5	Y	-
4	JDK or JRE	Java SE 7 or later	Y	[When using PL/extJava] Required for creation of Java applications used in PL/extJava. Note that Java SE 7 must be used.

Y: Can be used

N: Cannot be used



Note

The following programs are installed during installation of Enterprise Postgres:

- Microsoft Visual C++ 2005 Redistributable Package version 8.0.50727.42
- Microsoft Visual C++ 2010 Redistributable Package version 10.0.40219

Do not uninstall the above programs as they are required for running Enterprise Postgres.

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

Table 2.3 Connectable servers

OS	Product name
Windows	FUJITSU Software Enterprise Postgres Standard
Linux	Edition 9.4 or later

2.3 Excluded Software

This section describes excluded software.

Enterprise Postgres

If all the following conditions are met, Enterprise Postgres cannot be installed:

- The products are all 32-bit, or all 64-bit
- The product generations are the same
- The editions are different

Example

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

- The product to be installed is Enterprise Postgres Standard Edition (64-bit) 9.4

Other products

There are no exclusive products.

2.4 Required Patches

There are no required patches.

2.5 Hardware Environment

The following hardware is required to use Enterprise Postgres.

Memory

256 MB or more is recommended (At least 128 MB is required).

2.6 Supported System Environment

This section describes the supported system environment.

2.6.1 TCP/IP Protocol

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



Note

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

2.6.2 File System

You can install Enterprise Postgres only if the system folder is an NTFS volume.

2.7 Enterprise Postgres-based PostgreSQL Versions

The Enterprise Postgres-based PostgreSQL version is PostgreSQL 9.4.4.

2.8 Notes on Using Streaming Replication

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

*1: The product version is indicated by "Vxx" in the notation "Vxx.y.z".



Note

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

Chapter 3 Installation

This chapter describes the procedures for the installation of Enterprise Postgres.

Note

- The installation must be performed by a user with administrator privileges (a user ID that belongs to the Administrators group).
- Stop all applications before starting the installation.
- The Windows Installer service must be running.
- In either of the following cases, it is necessary to switch to install mode by executing the command shown below before installation. Also, after the installation is completed, execute the command shown below to switch back to execute mode.
 - In Windows Server(R) 2008, the terminal server is installed in application server mode.
 - In Windows Server(R) 2008 R2, Windows Server(R) 2012 or Windows Server(R) 2012 R2, the remote desktop service is installed in application server mode.

[Before the installation]

```
CHANGE USER /INSTALL
```

[After the installation]

```
CHANGE USER /EXECUTE
```

- The following window may be displayed when executing the installation program:



If this window is displayed, perform the following operations:

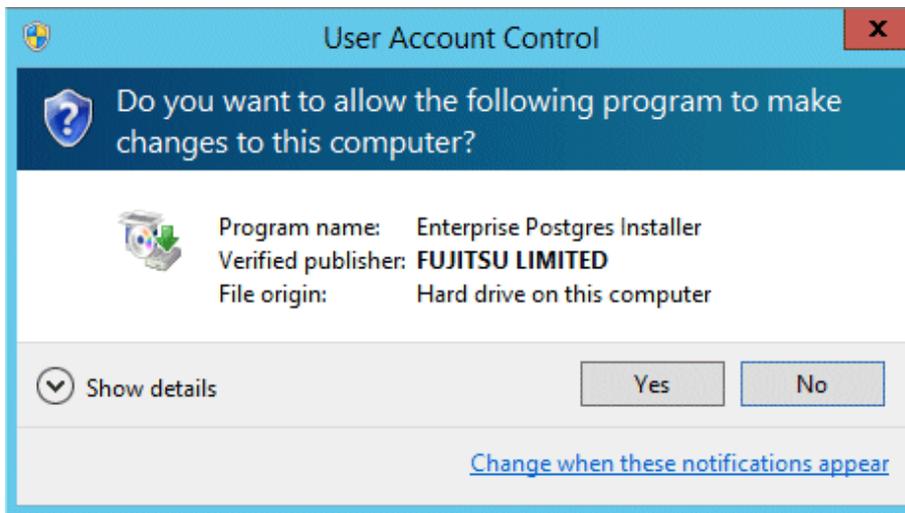
1. Perform the installation steps until the [InstallShield Wizard Complete] window is displayed.
2. At the window shown above, click [Next].

3. The window shown below is displayed. Click [Finish].



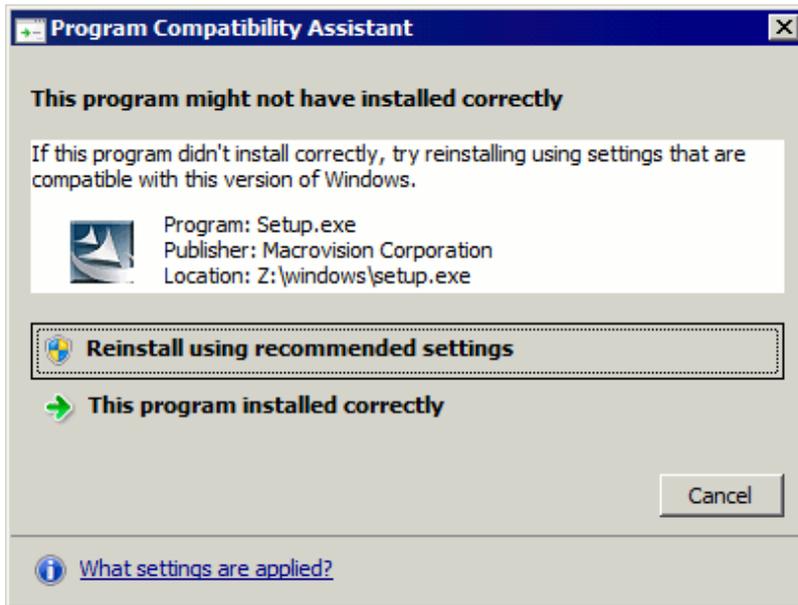
Information

- If a [User Account Control] dialog box such as the following is displayed at the start of the installation, click [Yes] to continue processing:



If [No] is clicked, permission to continue is denied and an [Error] dialog box will be displayed. To continue the installation, click [Retry] at the [Error] dialog box. To end the installation, click [Cancel].

- If installation is suspended or processing terminates abnormally, a dialog box of the program compatibility assistant similar to the one shown below may be displayed. Click [This program installed correctly] and continue operation.



3.1 Pre-installation Tasks

Check the system environment below before installing Enterprise Postgres.

Check the disk space

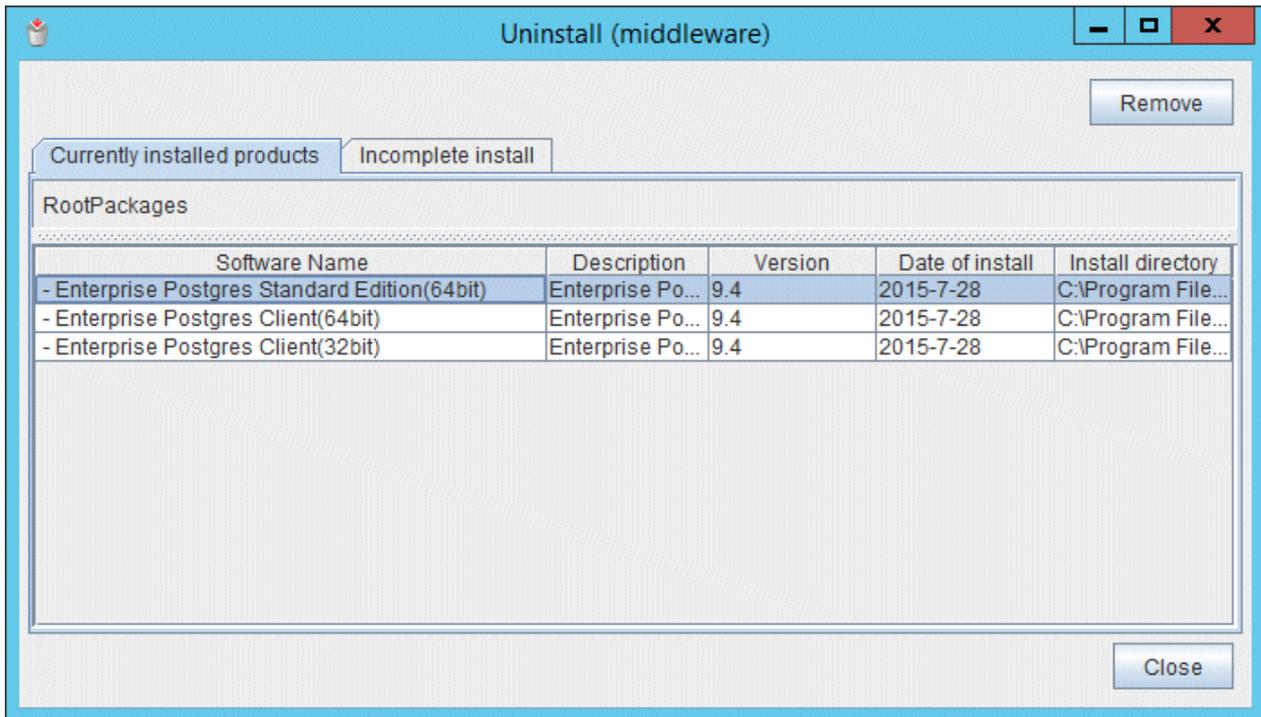
Ensure that there is sufficient disk space to install Enterprise Postgres.

Check the installed product and determine the installation method

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

The installed products are displayed in the window below.

In Windows, click [All Programs] or [All apps], then [Fujitsu], and then [Uninstall (middleware)].



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

- Reinstallation
- Upgrade installation
- Multi-version installation

Remove applied updates

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



Note

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

- Performing reinstallation
 - If an update with the same update and version number is applied, an error informing you that the update has already been applied is displayed.
- Performing upgrade installation
 - Unnecessary resources (area for backout use) remain on the disk.

1. Display the applied updates

Execute the following command to display the applied updates:

Example

```
C:\>uam showup
```

2. Remove the updates

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

Example

```
C:\>uam remove -i update-number
```

Prepare when performing migration

To perform migration from an old version/level, refer to "[Appendix B Enterprise Postgres Migration Procedures](#)", check the procedure, and then start the migration tasks.

3.2 Installation in Interactive Mode

The installation must be performed by a user with administrator privileges (a user ID that belongs to the Administrators group).

The installation procedure is described below.

1. Stop applications and the client program

If the installation method is one of the following, the applications and the program must be stopped:

- Reinstallation
- Upgrade installation

Before starting the installation, stop the following:

- Applications that use the product
- pgAdmin
- Instance

Using WebAdmin

In the [Monitor] menu window, click [stop].

Using server commands

Execute the net command or sc command of the operating system.

Example 1

```
> net stop windowsServiceName
```

Example 2

```
> sc stop windowsServiceName
```

- Web server feature of WebAdmin

If you are using WebAdmin, stop WebAdmin.

Refer to "[E.1.3 Stopping the Web Server Feature of WebAdmin](#)" for details.

- Mirroring Controller

Execute the mc_ctl command with the stop mode option specified and stop the Mirroring Controller.

Example

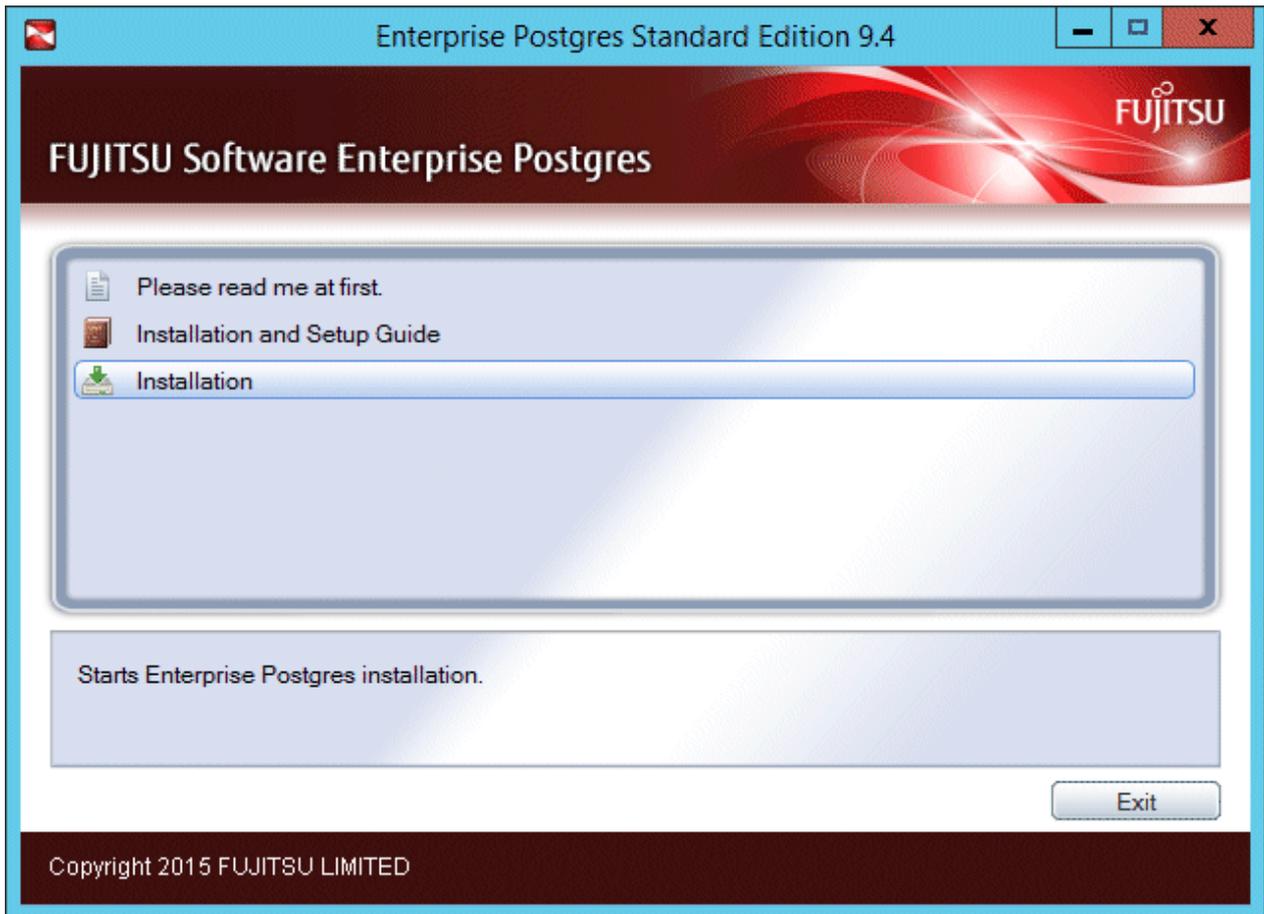
```
> mc_ctl stop -M D:\mdir\inst1
```

2. Inserting the DVD

Insert the Enterprise Postgres DVD into the drive.

3. Run the installation

The installation menu will be displayed. Click [Installation].



Note

If the Autorun feature of Windows is disabled, or if the operating system is a remote desktop service (terminal service), the installation program is not automatically started. Execute the following file using [Run] or Windows Explorer.

```
Z:\autorun.exe
```

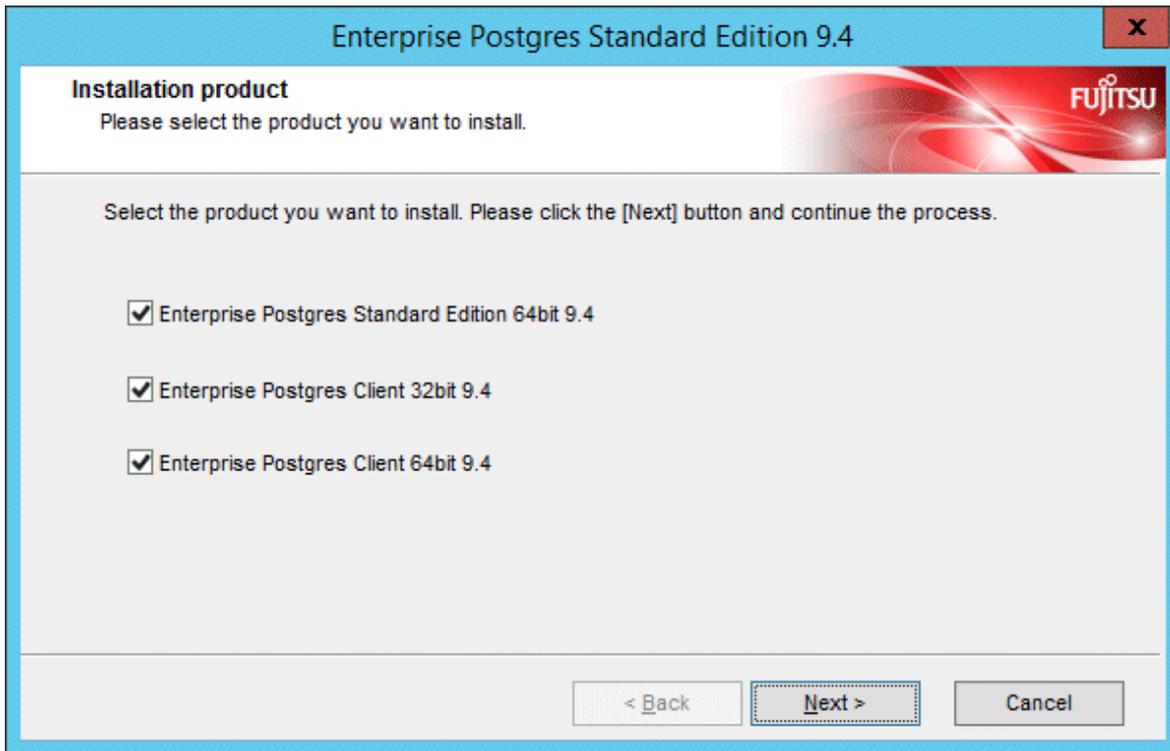
For Z, specify the drive name of the drive where the DVD has been inserted.

4. Select the product to install

The [Installation Product] window will be displayed.

Select the product to install, and then click [Next].

If the selected product can only be reinstalled, refer to "[8. Check the settings](#)".



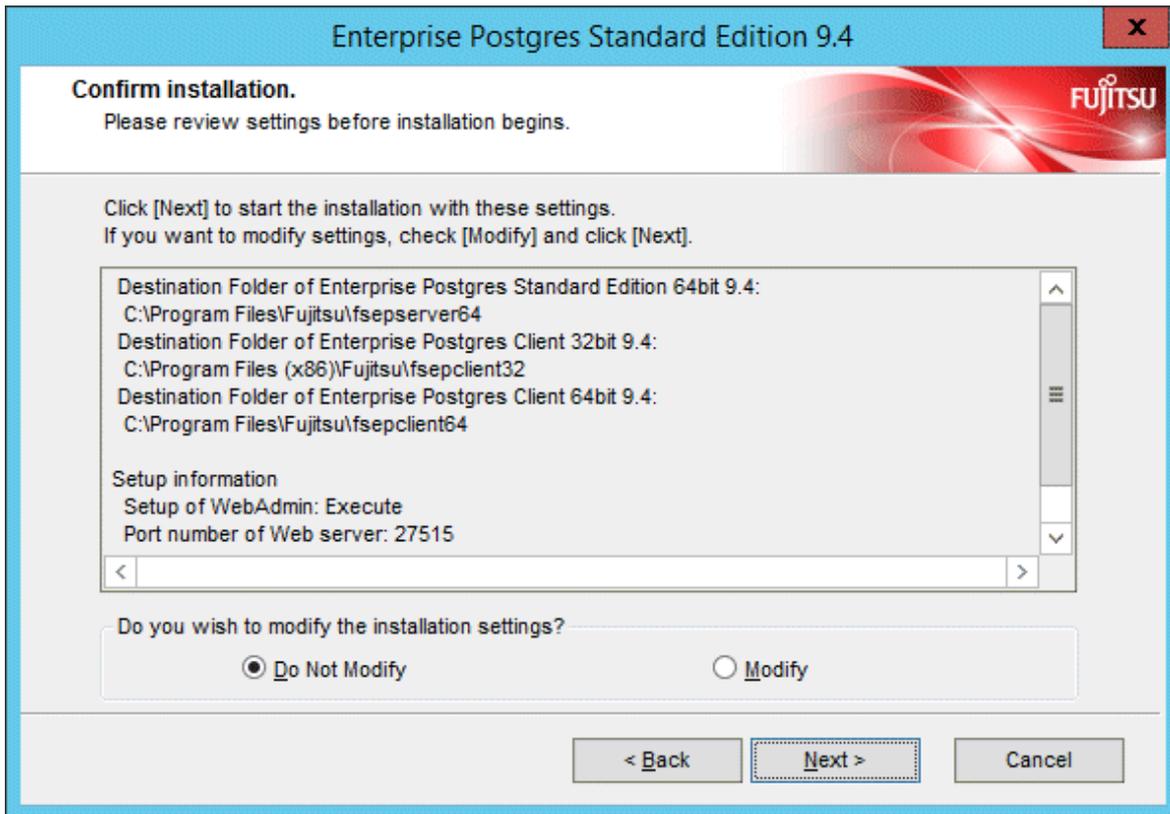
5. Checking the installation content

The [Confirm Setup Information] window will be displayed.

If there are no issues with the settings content, click [Next].

The installation process will then start. Proceed to "[9. Completing installation](#)".

To modify the settings content, select [Modify], and then click [Next].



Note

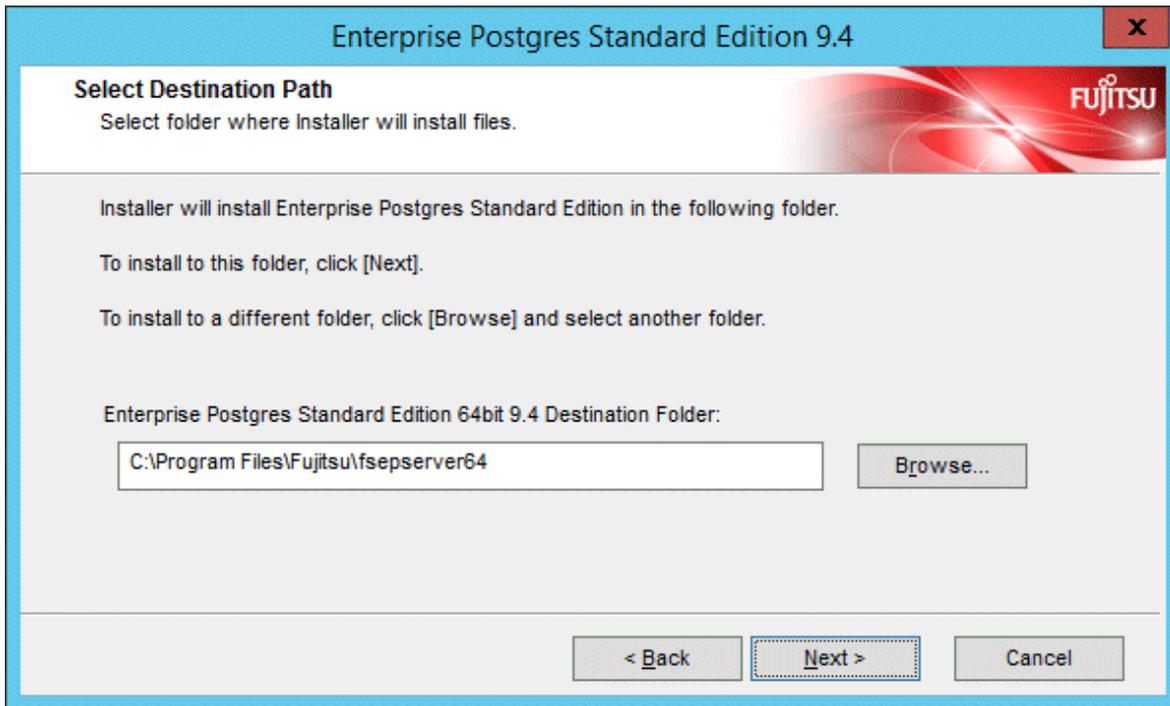
If using WebAdmin for operation, make a note of the port number of the Web server displayed in the settings, for use in the Windows firewall settings.

6. Enter the installation destination

If the installation method is one of the following, the [Select Destination Path] window is displayed for each product:

- New installation
- Multi-version installation

Enter the installation destination, and then click [Next].



Note

- It is necessary to specify a local disk as the installation destination of Enterprise Postgres.
- If using PL/extJava, specify the directory defined in [Enterprise Postgres Destination Folder] using up to 50 bytes.

7. Setting up WebAdmin

The [The Setup of WebAdmin] window is displayed. To perform setup, confirm that the port numbers are unused port numbers in the following ranges, and then click [Next]:

- 1024 to 49151

Enterprise Postgres Standard Edition 9.4

The setup of WebAdmin
Please select whether to run the installation and the setup of WebAdmin simultaneously.

Is the setup of WebAdmin of Enterprise Postgres executed?

No

Yes

Port number of Web server[1024-49151]

Internal Port number for WebAdmin[1024-49151]

< Back Next > Cancel

 **Note**

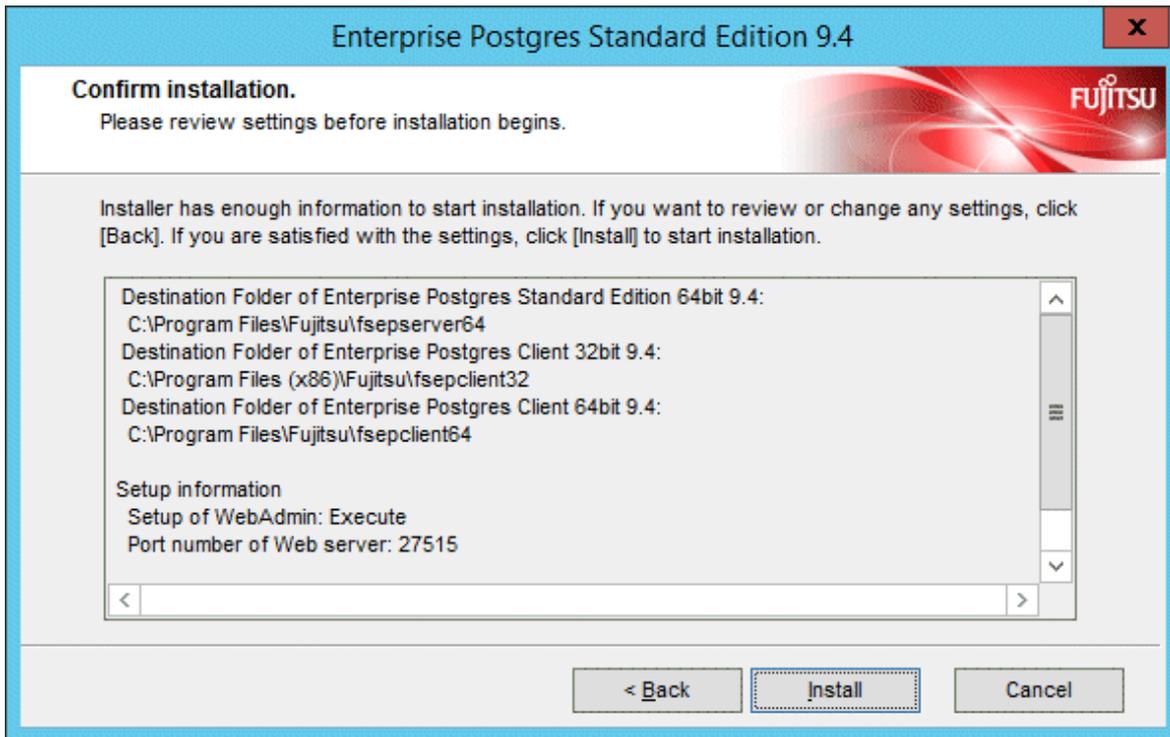
Make a note of the Web server port number for use in the Windows firewall settings.

8. Check the settings

The [Confirm installation] window will be displayed.

If there are no issues with the settings content, click [Install].

To modify the settings content again, click [Back].



9. Completing installation

The completion window is displayed. Click [Finish].



Installation is successful if the following has been added to [All Programs] or [All apps] in Windows when you finish the installation:

- Enterprise Postgres Standard Edition V9.4

3.3 Installation in Silent Mode

Installation in silent mode can be performed only when the installation method is one of the following:

- New installation
- Multi-version installation

The installation procedure is described below.

1. Insert the DVD

Insert the "server program" DVD in the DVD drive.

The [Install Menu] window will be displayed automatically. Click [Finish].

2. Create an installation parameters CSV file

Consider the features that will be required for system operations, and then create an installation parameters CSV file that uses the following specification format.

```
sectionName, parameterName, value
sectionName, parameterName, value
:
```

Refer to "[Appendix A Installation in Silent Mode](#)" for information on installation parameters CSV files.



Note

If using WebAdmin for operation, make a note of the port number of the Web server displayed in the settings (the port number defined in WebPortNumber1), for use in the Windows firewall settings.



Information

The template for the installation parameters CSV file is "*Z:\sample\sample.csv*" (for Z, specify the drive name of the drive where the DVD has been inserted).

3. Start the command prompt

In Windows, right-click [Command Prompt] and then select [Run as administrator].

4. Run the installation

Execute the command below.

```
Z:\>silent.bat c:\temp\inspara.csv
```

For Z, specify the drive name of the drive where the DVD has been inserted.

Also in the example above, c:\temp\inspara.csv is the installation parameter CSV file name.

If the installer ends in an error, a message is output to the log file and return values are returned. Refer to "[Appendix A Installation in Silent Mode](#)" for details.

3.4 Post-installation Tasks

There are no post-installation tasks.

Chapter 4 Setup

This chapter describes the setup procedures to be performed after installation completes.

4.1 Operating Method Types and Selection

This section describes how to operate Enterprise Postgres.

There are two methods of managing Enterprise Postgres operations - select one that suits your purposes:

The Operation Guide describes the operating method using WebAdmin, and the equivalent operating method using the server commands.

Simple operation management using a web-based GUI tool (WebAdmin)

Suitable when using frequently used basic settings and operations for operation management.

This method allows you to perform simple daily tasks such as starting the system before beginning business, and stopping the system when business is over, using an intuitive operation.

Usage method

Usage is started by using WebAdmin to create the instance.

By using an external scheduler and the `pgx_dmpall` command, periodic backups can be performed, which can then be used in recovery using WebAdmin.



- Do not use a server command other than `pgx_dmpall` or a server application. Operation modes that use server commands and server applications cannot be used in conjunction with WebAdmin. If used, WebAdmin will not be able to manage the instances correctly.
Refer to Reference and the PostgreSQL Documentation for information on server commands and server applications.
- An instance that was created using WebAdmin can be set up to operate database multiplexing, however once the setup is complete, it will no longer be possible to select that instance from the WebAdmin management window. Refer to the Cluster Operation Guide for information on how to perform database multiplexing operations.

Advanced operation management using server commands

When operating in a system that is automated by operation management middleware (Systemwalker Centric Manager, for example), this method allows you to use more detailed settings and operations and perform higher level operation management.

An overview of the operating method using the GUI, and its relationship with the operating method using the server commands, are shown below.

Refer to the Operation Guide for details.

Operation		Operation with the GUI	Operation with commands
Setup	Creating an instance	WebAdmin is used. The server machine capacity, and the optimum parameter for operations using WebAdmin, are set automatically.	The configuration file is edited directly using the <code>initdb</code> command.
	Changing the configuration files	WebAdmin is used.	The configuration file is edited directly.
Activating and stopping an instance		WebAdmin is used.	The <code>net</code> command or <code>sc</code> command of the operating system is used.
Creating a database		This is defined using <code>pgAdmin</code> of the GUI tool, or using the <code>psql</code> command or the application after specifying the DDL statement.	

Operation		Operation with the GUI	Operation with commands
Backing up the database		WebAdmin, or the <code>pgx_dmpall</code> command, is used.	It is recommended that the <code>pgx_dmpall</code> command be used. Recovery to the latest database can be performed.
Database recovery		WebAdmin is used.	To use the backup that was performed using the <code>pgx_dmpall</code> command, the <code>pgx_rcvall</code> command is used.
Monitoring	Database errors	The status in the WebAdmin window can be checked. (*1)	The messages that are output to the database server log are monitored (*1)
	Disk space	The status in the WebAdmin window can be checked. A warning will be displayed if the free space falls below 20%. (*1)	This is monitored using the <code>fsutil</code> command (check free space), and the <code>dir</code> command (check used space), of the operating system, for example. (*1)
	Connection status	This can be checked using <code>pgAdmin</code> of the GUI tool, or referencing <code>pg_stat_activity</code> of the standard statistics view from <code>psql</code> or the application.	

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



See

Refer to "Periodic Operations" and "Actions when an Error Occurs" in the Operation Guide for information on monitoring and database recovery.

4.2 Preparations for Setup

This section describes the preparation required before setting up Enterprise Postgres.

4.2.1 Creating an Instance Administrator

Decide the OS user account that will become the instance administrator. Use either a new user, or a user that already exists.

To create a user in Windows, select [Administrative Tools], [Computer Management], and then create the user in [Local Users and Groups]. Refer to "Help and Support" for details.

The following characters can be used for user names:

- - (hyphen)
- _ (underscore)
- Space
- A-Z, a-z, 0-9 (alphanumeric)



Note

The following notes apply if using WebAdmin for operations:

- After creating the user account of the instance administrator, log in to the operating system. A profile directory is created for the user when logging in to the operating system for the first time. This directory will be used by WebAdmin.

- If changing the password for the user account of the instance administrator, always ensure to stop the instance and log out of WebAdmin before making the change. If you mistakenly change the password while logged in to WebAdmin or while the instance is running, log out from WebAdmin, and then log in again, and stop and start the instance.

4.2.1.1 Security policy settings

If using commands for operation, security settings that allow logon as a service are required for the operating system user account of the instance administrator in order to start and stop an instance using a Windows service.

Information

If using WebAdmin for operation, these settings are not required as WebAdmin performs the settings automatically for the user ID (operating system user account) that logged in to the database server.

The following explains how to perform the security settings to allow logon as a service:

1. Displaying the Local Security Policy window

In Windows, select [Administrative Tools], and then click [Local Security Policy].

2. Setting up security

1. In the [Local Security Policy] window, select [Security Settings], select [Local Policies], and then click [User Rights Assignment].
2. Under [Policy] in the [User Rights Assignment] window, double-click [Log on as a service].
3. In the [Log on as a service Properties] window, set the following:
 - a. Select the [Local Security Setting] tab.
 - b. On the [Local Security Setting] tab, click [Add User or Group].
 - c. In the [Select Users or Groups] window, enter the operating system user account of the instance administrator in [Enter the object names to select].
 - d. Click [OK].
4. In the [Log on as a service Properties] window, click [OK].

4.2.2 Preparing Directories for Resource Deployment

Prepare the directories required when creating instances.

Considerations when deploying resources

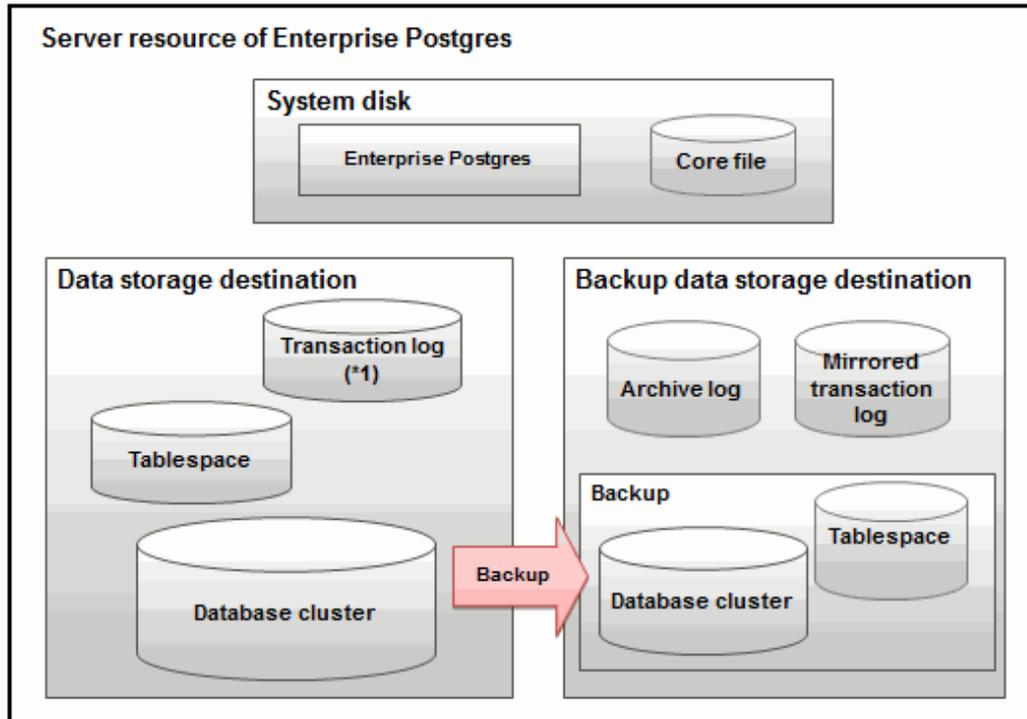
The disk configuration on the resource deployment destination is important, because it affects not only recovery following disk corruption, but normal operation as well. The points for determining the disk configuration are as follows:

1. If the backup data storage destination and the data storage destination are both lost, it will not be possible to recover the data, so deploy them to separate disks.
2. To shorten the recovery time following a single disk fault, deploy the system disk and data storage destination to separate disks.
3. The backup data storage destination requires at least double the capacity of the data storage destination, so deploy it to the disk with the most space available.
4. When large amounts of data are updated, the write-to load for the data storage destination, transaction log storage destination, and backup data storage destination (mirrored transaction log) will also be great. For this reason, deploy them to separate disks, out of consideration for performance.

Note

- When using the volume manager provided by the operating system, be aware of which physical disk the file system has been created on, for example, by deploying the data storage destination and the backup data storage destination to separate disks.

- If using WebAdmin, specify an NTFS volume for the data storage destination and backup data storage destination. A network drive cannot be specified.



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

Resource	Role
Database cluster	The area where the database is stored. It is a collection of databases managed by an instance.
Tablespace	Stores table files and index files in a separate area from the database cluster.
Transaction log	Stores log information in preparation for a crash recovery or rollback.
Archive log	Stores log information for recovery
Corefile	Enterprise Postgres process corefile output when an error occurs with a Enterprise Postgres process.

Examples of disk deployment

The following are examples of disk deployment:

Number of disks	Disk	Deployment
3	System disk	Enterprise Postgres program
		Corefile
	Connected physical disk	Data storage destination, transaction log storage destination
	Connected physical disk	Backup data storage destination
2	System disk	Enterprise Postgres program
		Corefile

Number of disks	Disk	Deployment
		Data storage destination, transaction log storage destination
	Connected physical disk	Backup data storage destination

Proposal for disk deployment using WebAdmin

To generate an instance using WebAdmin, we propose an optimum deployment that takes into account the status of all disks at the time of instance is generation, and item 1 to 3 in the "Considerations when deploying resources" subheading above, based on the criteria below (note that a different deployment can also be specified).

- The instance administrator has read and write privileges for the volumes.

Information

- Even if the mount point of each file system was deployed to a different partition or slice, it will still be verified if they are on the same physical disk.
- If the operating system volume manager was used, it will not be verified if the file systems were created on the same physical disk.

Preparing directories

The directories to be prepared depend on the way that you create the instances.

The following shows the directories that need to be prepared:

Directory to be prepared	Using WebAdmin	Using the initdb command
Data storage destination	Y (*1)	Y
Backup data storage destination	Y (*1)	O
Transaction log storage destination	O (*1) (*2)	O
Corefile output destination	N (*3)	O

Y: Required

O: Optional

N: Not required

*1: WebAdmin automatically creates a directory

*2: The default is to create in a directory in the data storage destination. When it is necessary to distribute the I/O load for the database data and the transaction log, consider putting the transaction log storage destination on a different disk from the data storage destination

*3: The default is to output to "*userProfileFolder\localSettingsFolder\Fujitsu\fssep_version\instanceName\core*". For example, if using Windows Server(R) 2008 R2, *userProfileFolder\localSettingsFolder* will be "C:\Users*userName*\AppData\Local".

To change the output destination, specify in the *core_directory* parameter and *core_contents* parameter in *postgresql.conf*. Refer to "Parameters" in the Operation Guide for information on the settings for these parameters.

Note

- The directories must meet the following conditions:
 - The directory owner must be the OS user account that you want to be the instance administrator
 - The directory must have write permission
 - The directory must be empty
- If using WebAdmin, network drives cannot be used.

- If using WebAdmin, the following halfwidth characters can be used for directory names:
 - \ (backslash)
 - - (hyphen)
 - _ (underscore)
 - : (colon)
 - Space
 - A-Z, a-z, 0-9 (alphanumeric)
 - If anti-virus software is used, set scan exception settings for folders so that none of the server resources that comprise Enterprise Postgres are scanned for viruses. Alternatively, if the server resources that comprise Enterprise Postgres are to be scanned for viruses, stop the instance and perform the scan when tasks that use Enterprise Postgres are not operating.
-

Confirm and configure directory access permissions

If the instance administrator user has "Administrator" permissions (user ID belonging to the Administrators group), it is necessary to configure the settings so that each directory inherits the file and directory access permissions for the instance administrator user.

Therefore, ensure that the setting to inherit permissions has been configured.

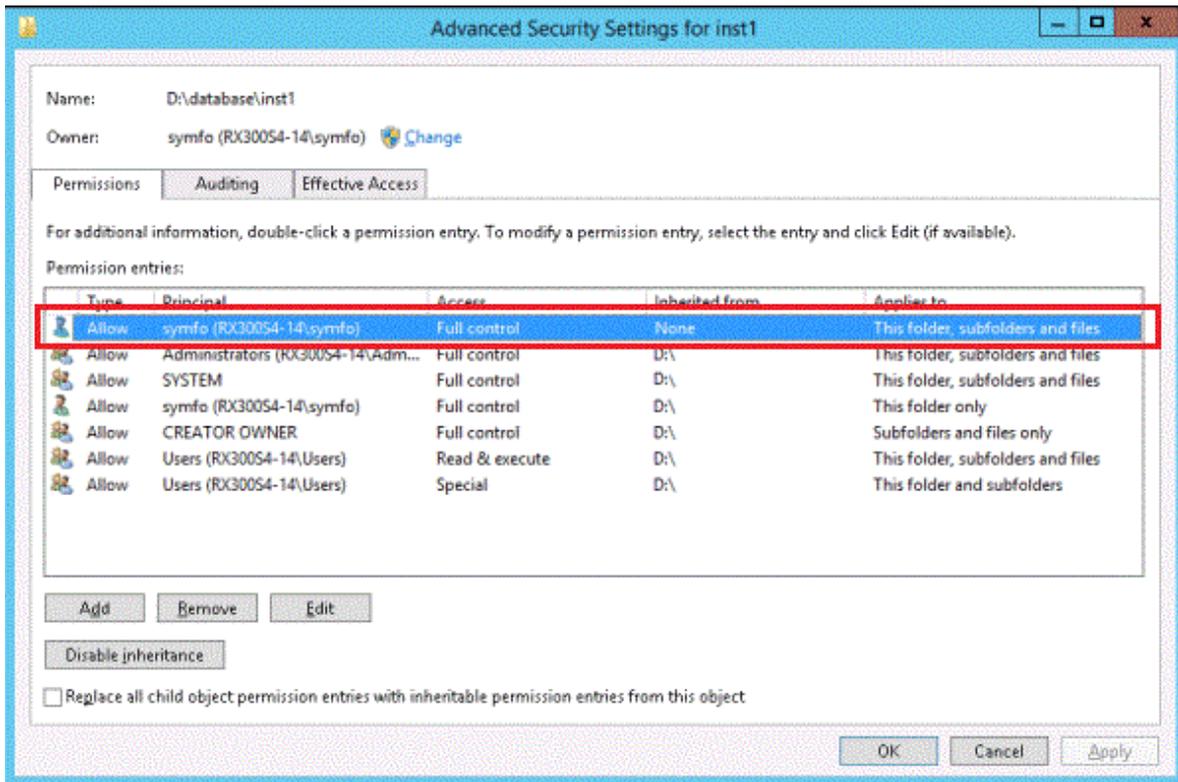
The following is an explanation on how to confirm and configure the settings.

How to confirm access permissions

Perform the following operations in Windows Explorer on the directories to be prepared in advance:

1. Right-click on the applicable directory, and then click [Properties] from the menu that is displayed.
2. In the [*applicableDir* Properties] window, select [Security] >> [Advanced].
3. In the [Advanced Security Settings for *applicableDir*] window, and in the [Permission entries] list under the [Permissions] tab, confirm that the application destination of the instance administrator user is "This folder, subfolders and files".
4. Click [OK].

A confirmation example using Windows Server(R) 2012 is shown below.



How to configure the access permissions

Perform the following operations in Windows Explorer if there are any directories that have not been configured for the access permissions to be inherited.

Windows Server(R) 2008 or Windows Server(R) 2008 R2:

1. Right-click on the applicable directory, and then click [Properties] from the menu that is displayed.
2. In the [*applicableDir* Properties] window, select [Security] >> [Advanced].
3. In the [Advanced Security Settings for *applicableDir*] window, under the [Permissions] tab, click [Change Permissions].
4. Click [Add].
5. In the [Select User or Group] window, enter the instance administrator user name as the object name to select, and then click [OK].
6. In the [Permission Entry for *applicableDir*] window, set [This folder, subfolders and files] for [Apply to], and under [Permissions], allow read and write permissions, and then click [OK].
7. In the [Advanced Security Settings for *applicableDir*] window, confirm that the instance administrator user has been added, with [This folder, subfolders and files] set for [Apply to] in the [Permission entries] list.
8. Click [OK].

Windows Server(R) 2012 or Windows Server(R) 2012 R2:

1. Right-click on the applicable directory, and then click [Properties] from the menu that is displayed.
2. In the [*applicableDir* Properties] window, select [Security] >> [Advanced].
3. In the [Advanced Security Settings for *applicableDir*] window, click [Add].
4. In the [Permission Entry for *applicableDir*] window, click [Select a principal].
5. In the [Select User or Group] window, enter the instance administrator user name as the object name to select, and then click [OK].
6. In the [Permission Entry for *applicableDir*] window, set "This folder, subfolders and files" for "Apply to:", and under "Basic permissions", allow read and write permissions, and then click [OK].

7. In the [Advanced Security Settings for *applicableDir*] window, confirm that the instance administrator user has been added, with "This folder, subfolders and files" set for "Apply to" in the "Permission entries" list.
8. Click [OK].

Information

The access permissions can also be configured using the `icacls` command provided by the operating system.

The following is an execution example in which the application destination is set to "(OI)(CI)" and the access permissions are set to "(F)(Full access permissions)" when the data storage destination is "D:\database\inst1" and the instance administrator user is "fsepuser":

```
>icacls D:\database\inst1 /grant fsepuser:(OI)(CI)(F)
processed file: D:\database\inst1
Successfully processed 1 files; Failed processing 0 files
```

4.2.3 Estimating Resources

Estimate the resources to be used on the Enterprise Postgres.

Refer to "[Appendix I Estimating Database Disk Space Requirements](#)" for information on estimating database disk space requirements.

Refer to "[Parameters automatically set by WebAdmin according to the amount of memory](#)" when creating multiple instances with WebAdmin.

Refer to "[Appendix J Estimating Memory Requirements](#)" when creating instances with the `initdb` command, to estimate memory usage.

4.2.4 Windows Firewall Settings

This section explains the Windows firewall settings required if using WebAdmin for operation.

These settings are not required if using server commands for operation.

If the Windows firewall feature is to be enabled, you should enable a port number on the Web server. The following explains how to enable a port number:

Windows Server(R) 2008:

1. In the [Start] menu, click [Control Panel].
2. Click [Security] and then click [Windows Firewall].
3. In the [Windows Firewall] window, click [Change settings].
4. On the [Exceptions] tab, click [Add port].
5. In the [Add a Port] window, set the following:
 - a. In [Name], specify the desired name.
 - b. In [Port number], specify the Web server port number that was specified during the WebAdmin setup.
 - c. Select [TCP].
6. Click [OK].
7. On the [Exceptions] tab, in the [Program or port] list, check if the added port is enabled.
8. In the [Windows Firewall] window, click [OK].

Windows Server(R) 2008 R2:

1. In the [Start] menu, click [Control Panel].
2. Click [System and Security], and then click [Windows Firewall].
3. In the [Windows Firewall] window, click [Advanced settings] on the left side of the window.

4. In the [Windows Firewall with Advanced Security] window, click [Inbound Rules] on the left side of the window.
5. Click [New Rule] on the right side of the window.
6. In the [New Inbound Rule Wizard] window, select [Port], and then click [Next].
7. Select [TCP] and [Specific local ports], then specify the Web server port number specified during the WebAdmin setup, and click [Next].
8. Select [Allow the connection], and then click [Next].
9. Select the profiles for which this rule applies, and then click [Next].
10. In [Name], specify the desired name, and then click [Finish].
11. In the [Windows Firewall with Advanced Security] window, check if the added rule is enabled under [Inbound Rules] in the center of the window.

Windows Server(R) 2012 or Windows Server(R) 2012 R2:

1. Right-click the [Start] screen, and then click [All apps] that is displayed in the bottom-right of the screen. In the [Apps] window, click [Control Panel].
2. Click [System and Security], and then click [Windows Firewall].
3. In the [Windows Firewall] window, click [Advanced settings] on the left side of the window.
4. In the [Windows Firewall with Advanced Security] window, click [Inbound Rules] on the left side of the window.
5. Click [New Rule] on the right side of the window.
6. In the [New Inbound Rule Wizard] window, select [Port], and then click [Next].
7. Select [TCP] and [Specific local ports], then specify the Web server port number specified during the WebAdmin setup, and click [Next].
8. Select [Allow the connection], and then click [Next].
9. Select the profiles for which this rule applies, and then click [Next].
10. In [Name], specify the desired name, and then click [Finish].
11. In the [Windows Firewall with Advanced Security] window, check if the added rule is enabled under [Inbound Rules] in the center of the window.

4.2.5 Preparing for Output to the Event Log

This section provides an explanation on the preparation to be carried out if you are outputting error logs to the event log.

If outputting error logs to the event log, you should register an event source name beforehand.

If you do not register an event source name, the message content output to the event log may be incomplete.

Due to the default event source name "Enterprise Postgres" being output to the event log when using the following commands, you should register this default event source name beforehand:

- pg_ctl command
- pgx_dmpall command
- pgx_rcvall command

The following is an example in which the DLL of a 64-bit product is registered under the default event source name:

```
regsvr32 "C:\Program Files\Fujitsu\fsepserver64\lib\pgevent.dll"
```

If using multiple instances

You can output messages corresponding the event source name assigned by the user, so that messages output to the event log can be identified by instance.

The following is an example in which the DLL of a 64-bit product is registered under the event source name "Enterprise Postgres inst1":

```
regsvr32 /n /i:"Enterprise Postgres inst1" "C:\Program Files\Fujitsu\fsepserver64\lib\pgevent.dll"
```

You will need to edit the parameters for each instance, therefore, after creating an instance, refer to ["4.5.1 Error Log Settings"](#) when performing this setting.

Note that this step is not required if using WebAdmin to create an instance.

If installing multiple versions

If Enterprise Postgres is already installed on the same machine, search for the key below in Registry Editor, and make a note of the path of the registered DLL. Afterwards, register a new DLL under the default event source name.

Use the DLL path that you made a note of in the above step when re-registering the default event source name during an uninstall.

Enterprise Postgres



See

Refer to "Registering Event Log on Windows" in "Server Setup and Operation" in the PostgreSQL Documentation for information on how to register event source names.

4.3 Creating an Instance

There are two methods that can be used to create an instance:

- [4.3.1 Using WebAdmin](#)
- [4.3.2 Using the initdb Command](#)

Creating multiple instances

Multiple instances can be created.

The memory allocated needs to be adjusted when multiple instances are created with WebAdmin (refer to ["Parameters automatically set by WebAdmin according to the amount of memory"](#) for details).

Features that cannot be set up using WebAdmin

The feature below cannot be set up using WebAdmin. After creating the instance using WebAdmin, perform the additional setup tasks according to the manual for each feature described in the Operation Guide.

- Storage data protection using transparent data encryption



Note

- Instances created using the initdb command cannot be managed using WebAdmin.
- Always use WebAdmin to delete instances that were created using it. Because WebAdmin management information cannot be deleted, WebAdmin will determine that the instance is abnormal.
- Databases with the names 'template0' and 'template1' are automatically created when an instance is created. These databases are used as the templates for databases created later. Furthermore, a default database with the name 'postgres' is automatically created, which will be used with Enterprise Postgres commands. It is important that you do not delete these databases created by default.
- When an instance that uses WebAdmin is created successfully, the following Windows service is registered:

fsep_version_edition_architecture_username_instanceName

The account and password of the instance administrator are registered in the Windows service. If the password for this account is changed, you must also change the password registered in the service. Change this at the Properties window registered in the Windows service.

- Refer to "[4.5.2 Configuring Automatic Start and Stop of an Instance](#)" for information on how to start and stop the operating system of the database server, and how to start and stop linked instances.

4.3.1 Using WebAdmin

This section describes how to create an instance using WebAdmin.

If WebAdmin is not set up when the installation is performed, refer to "[E.1 Setting Up WebAdmin](#)" and then set up WebAdmin.

Use Windows(R) Internet Explorer 8.0, 9.0, 10.0, or 11.0 as the browser, and allow cookies and pop-up requests from the server on which Enterprise Postgres is installed.

Refer to "[Appendix D Recommended WebAdmin Environments](#)" for information on how to change the pop-up request settings and other recommended settings.



- WebAdmin does not run in Windows(R) safe mode.
- If the same instance is operated from multiple WebAdmin windows, it will not work correctly.
- The browser [Back] button will not work correctly. If you need to navigate to a previous page, close the browser and then log in again.
- If the browser was not operated for a fixed period (about 15 minutes), the login window will be displayed again for the next operation. However, this will not happen if the [Monitor] window was displayed.
- Port access permissions
If a port is blocked (access permissions have not been granted) by a firewall, enable use of the port by granting access. Refer to the vendor document for information on how to grant port access permissions.
Consider the security risks carefully when opening ports.

4.3.1.1 Logging into WebAdmin

This section describes how to log in to WebAdmin.

Activation URL for WebAdmin

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

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

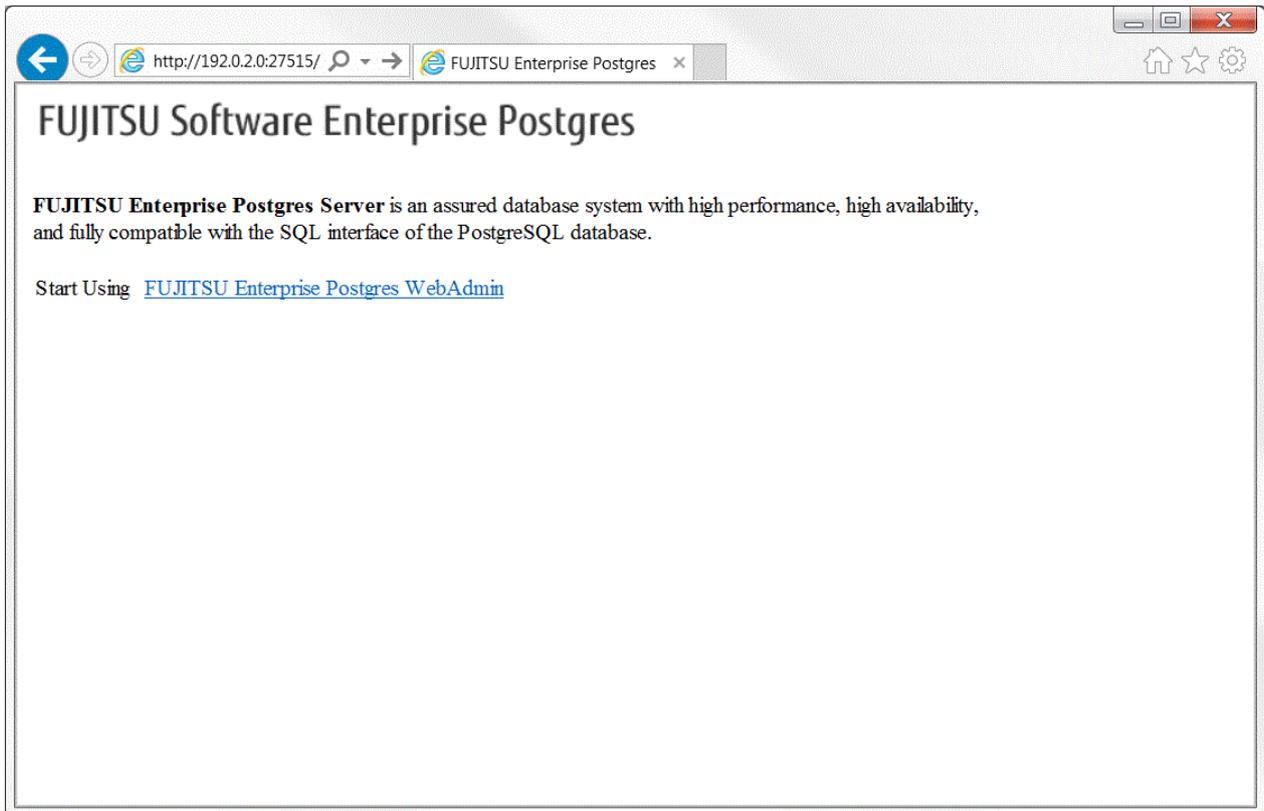
- *hostNameOrIpAddress*: Host name or IP address of the server where Enterprise Postgres is installed
- *portNumber*: Port number of WebAdmin. The default port number is 27515.

Example

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

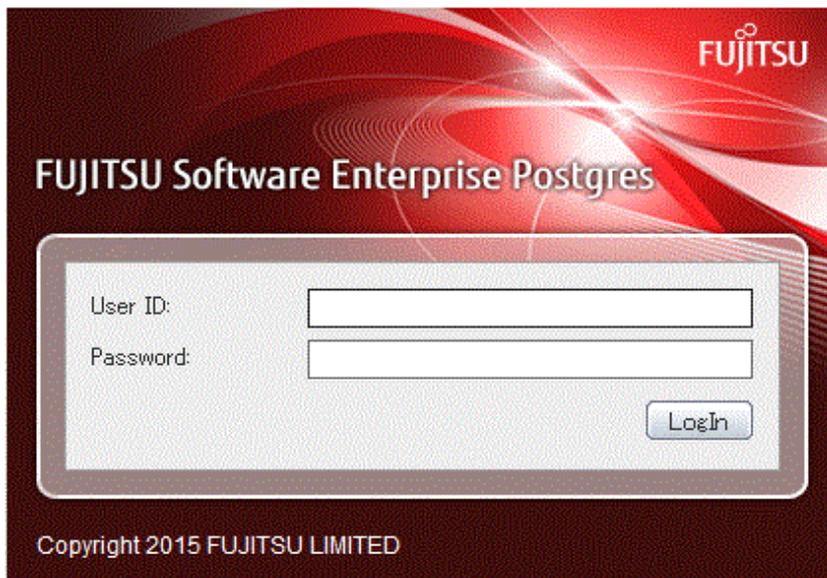
```
http://192.0.2.0:27515/
```

The activation URL window shown below is displayed.



Logging in to the database server

Clicking "FUJITSU Enterprise Postgres WebAdmin" on the startup window displays the WebAdmin [Login] window. Enter the instance administrator user ID (operating system user account name) and password, and log in to Enterprise Postgres.

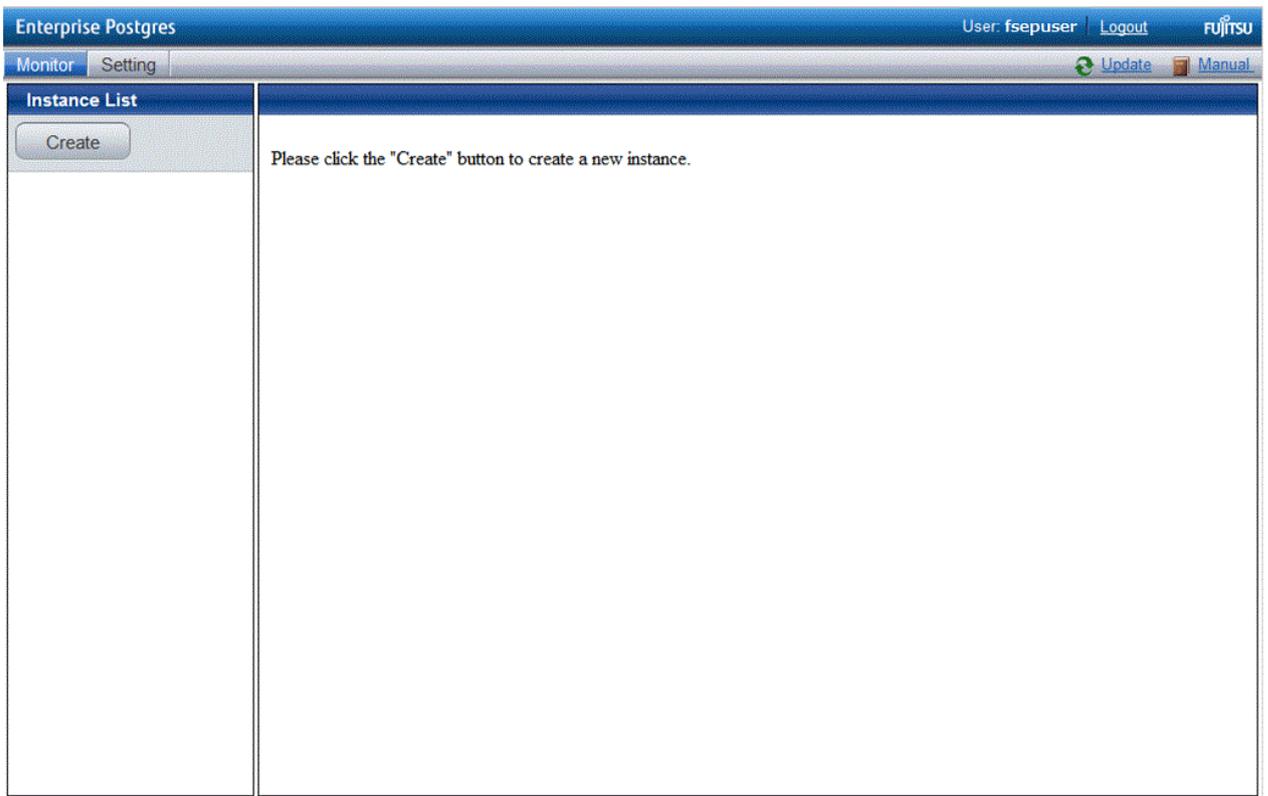


4.3.1.2 Creating an Instance

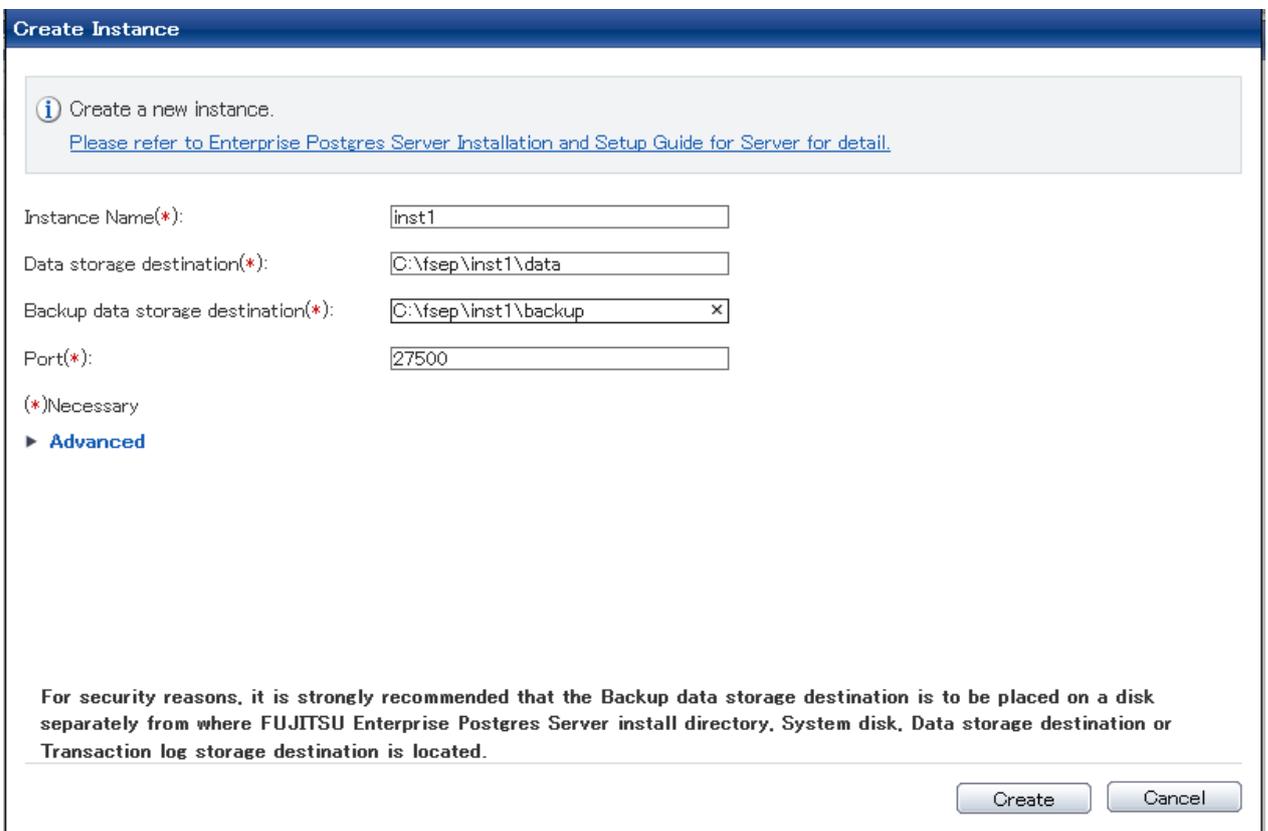
This section describes how to create an instance.

1. Activate WebAdmin, and log in to the database server.

2. Click [Create] in the [Instance List] window of the [Monitor].



3. Enter the information for the instance to be created.



Enter the following items:

- [Instance Name]: Name of the database instance to manage
The name must meet the conditions below:
 - Maximum of 16 characters
 - The first character must be an ASCII alphabetic character
 - The other characters must be ASCII alphanumeric characters
- [Data storage destination]: Directory where the database data will be stored
- [Backup data storage destination]: Directory where the database backup will be stored
- [Port]: Port number of the database server

Note

- Refer to "[4.2.2 Preparing Directories for Resource Deployment](#)" - "Considerations when deploying resources" for information on points to consider when determining the data storage destination, backup data storage destination, and transaction log storage destination.
 - Only the port number can be modified after the instance has been created.
 - Do not specify shortcuts for the data storage destination or backup data storage destination.
 - The following characters can be used when specifying the data storage destination and backup data storage destination:
 - \ (backslash)
 - - (hyphen)
 - _ (underscore)
 - : (colon)
 - Space
 - A-Z, a-z, 0-9 (alphanumeric)
 - Instance administrator read and write permissions are required for the data storage destination and backup data storage destination.
 - For the port number, specify an unused port number in the following range:
 - Windows Server(R) 2008, Windows Server(R) 2008 R2, Windows Server(R) 2012 or Windows Server(R) 2012 R2: 1024 to 49151
 - Make a note of the port number for use in the Windows firewall settings.
-

- Click [Advanced] to change the transaction log storage destination, database encoding and Auto Start.

Create Instance

 Create a new instance.
[Please refer to Enterprise Postgres Server Installation and Setup Guide for Server for detail.](#)

Instance Name(*):	<input type="text" value="inst1"/>
Data storage destination(*):	<input type="text" value="C:\fsep\inst1\data"/>
Backup data storage destination(*):	<input type="text" value="C:\fsep\inst1\backup"/>
Port(*):	<input type="text" value="27500"/>
(*)Necessary	
▼ Advanced	
Transaction log storage destination:	<input type="text" value="C:\fsep\inst1\tranlog"/>
Encoding:	<input type="text" value="UTF8"/>
Auto Start:	<input checked="" type="radio"/> Yes <input type="radio"/> No

For security reasons, it is strongly recommended that the Backup data storage destination is to be placed on a disk separately from where FUJITSU Enterprise Postgres Server install directory, System disk, Data storage destination or Transaction log storage destination is located.

Enter the following items:

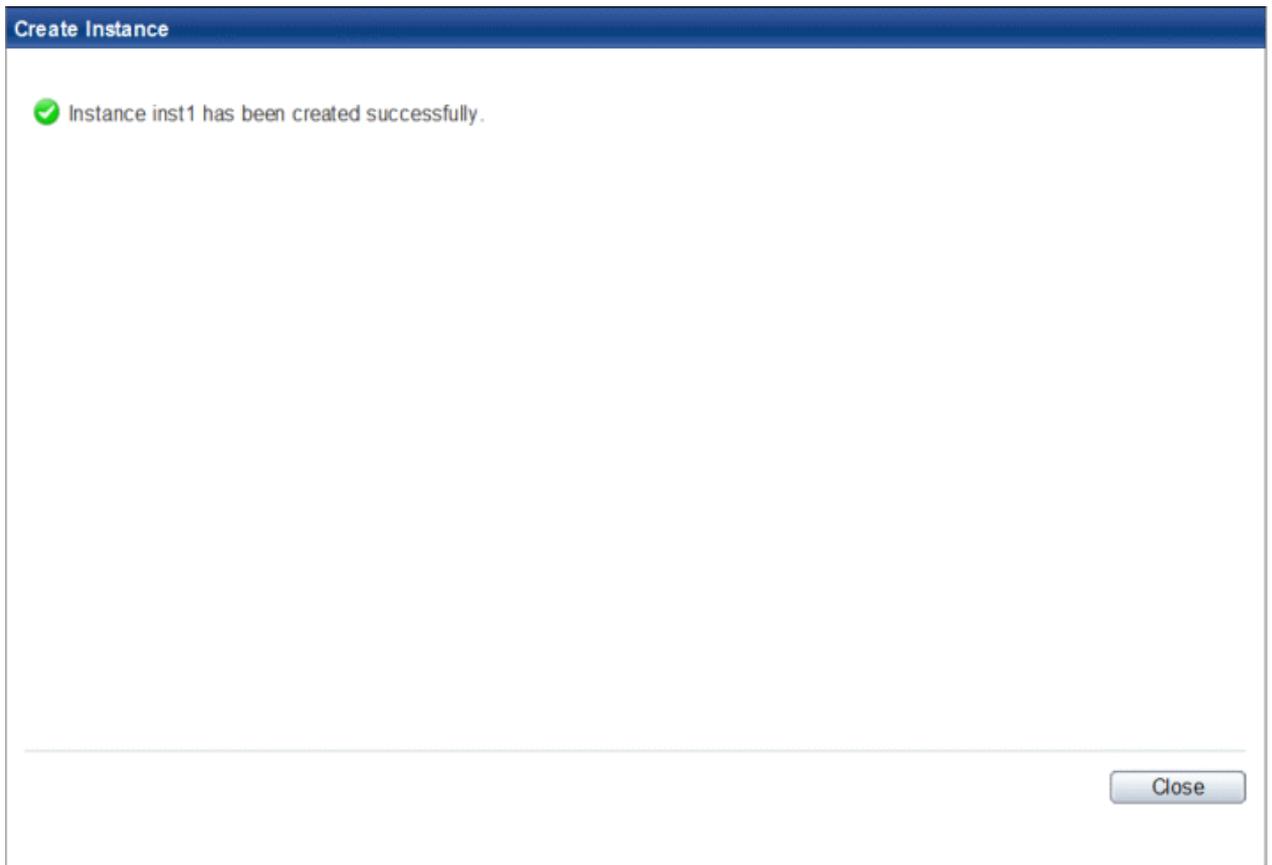
- [Transaction log storage destination]: Directory where the transaction log will be stored
- [Encoding]: Database encoding system
- [Auto Start]: Specify whether to automatically start or stop an instance when the operating system starts or stops.

Note

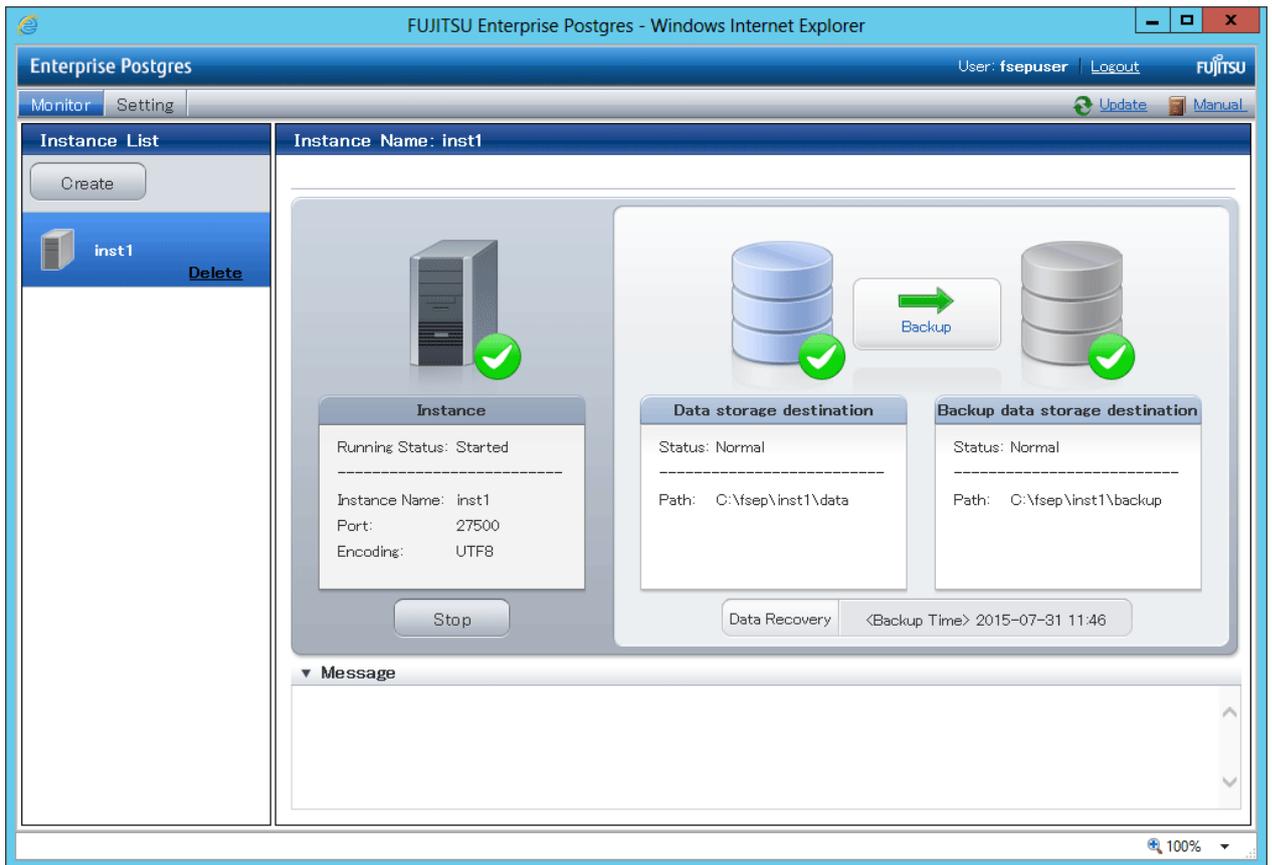
- Refer to "4.2.2 Preparing Directories for Resource Deployment" - "**Considerations when deploying resources**" for information on points to consider when determining the transaction log storage destination.
- Do not specify a shortcut for the transaction log storage destination.
- In the instance that was created using WebAdmin, the locale of the character set to be used in the database, and the locale of the collating sequence, are fixed using C.
- The following characters can be used when specifying the transaction log storage destination:
 - \ (backslash)
 - - (hyphen)
 - _ (underscore)
 - : (colon)
 - Space
 - A-Z, a-z, 0-9 (alphanumeric)
- Instance administrator read and write permissions are required for the data storage destination and backup data storage destination.

5. Click [Create] to create an instance.

When the creation of the instance is complete, the following window will be displayed:



6. The instance will be started when it is created successfully.



7. Back up the basic information that was set

Perform a backup of the file shown below. The file to be backed up contains the basic information specified in this screen, and this information will be needed to recover the WebAdmin management information when a fault occurs on the system disk.

```
installDir/gui/data/gui-instances
```

Note

The following message is output during startup of an instance when the startup process is operating normally, therefore, the user does not need to be aware of this message.

```
FATAL: the database system is starting up (XXXXX)
```

4.3.1.3 Changing the Settings

You can change the following information that is set when instances are created.

You can change the character set and maximum number of connections, for example, to suit the operating and management environment for Enterprise Postgres.

- [Character set](#)
- [Client authentication](#)
- [Communication](#)
- [SQL option](#)
- [Used memory](#)

Information

These settings are the same as the parameters that can be set in the files shown below. Refer to "[Appendix F Configuring Parameters](#)" for information on the equivalence relationship between the item name and the parameter.

- postgresql.conf
- pg_hba.conf

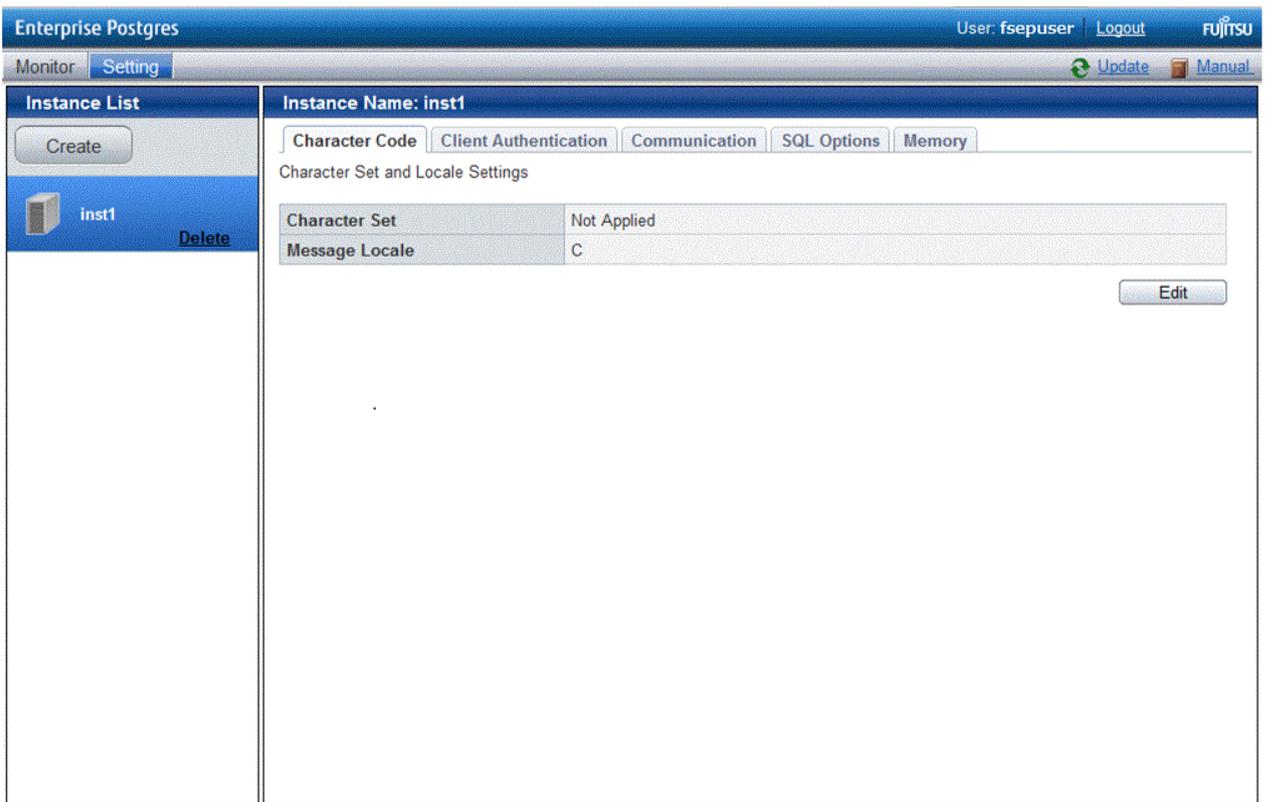
Note

The files shown below can also be modified directly, however if a parameter not described in "[Appendix F Configuring Parameters](#)" was edited by mistake, WebAdmin may not run correctly.

- postgresql.conf
- pg_hba.conf

Changing the character set

1. Start WebAdmin and log in to the database server.
2. Select [Character Code] from the [Setting] menu.
3. Click [Edit].



The screenshot displays the Enterprise Postgres WebAdmin interface. The top navigation bar includes the title "Enterprise Postgres", the user "User: fsepuser", a "Logout" link, and the Fujitsu logo. Below the navigation bar, there are tabs for "Monitor" and "Setting", along with "Update" and "Manual" links. The main content area is divided into two panels. The left panel, titled "Instance List", contains a "Create" button and a table with one entry: "inst1" with a "Delete" button next to it. The right panel, titled "Instance Name: inst1", shows a tabbed interface with "Character Code" selected. Underneath, the "Character Set and Locale Settings" are displayed in a table:

Character Set	Not Applied
Message Locale	C

An "Edit" button is located at the bottom right of the settings table.

4. Edit Character Set, and then click [Modify].

Character Set And Locale Settings

Character Set: Apply

Client Character Set: BIG5

Refer to a manual about the details of a Client Character Set. And please choose the Client Character Set suitable for the Server Encoding.

Message Locale: C

Modify Cancel

See

.....

Select a client-side encoding system that can be converted to/from the database encoding system. Refer to "PostgreSQL Documentation" - "Server Administration" - "Automatic Character Set Conversion Between Server and Client" for information on the encoding system combinations that can be converted.

.....

Changing client authentication

1. Start WebAdmin and log in to the database server.
2. Select [Client Authentication] from the [Setting] menu.
Click [Add] to register new authentication information.
To change authentication information, select the information, and then click [Edit].
To delete authentication information, select the information, and then click [Delete].

Note

When creating the instance, do not delete the entry below, because it is a connection required for WebAdmin to monitor the operational status of the database:

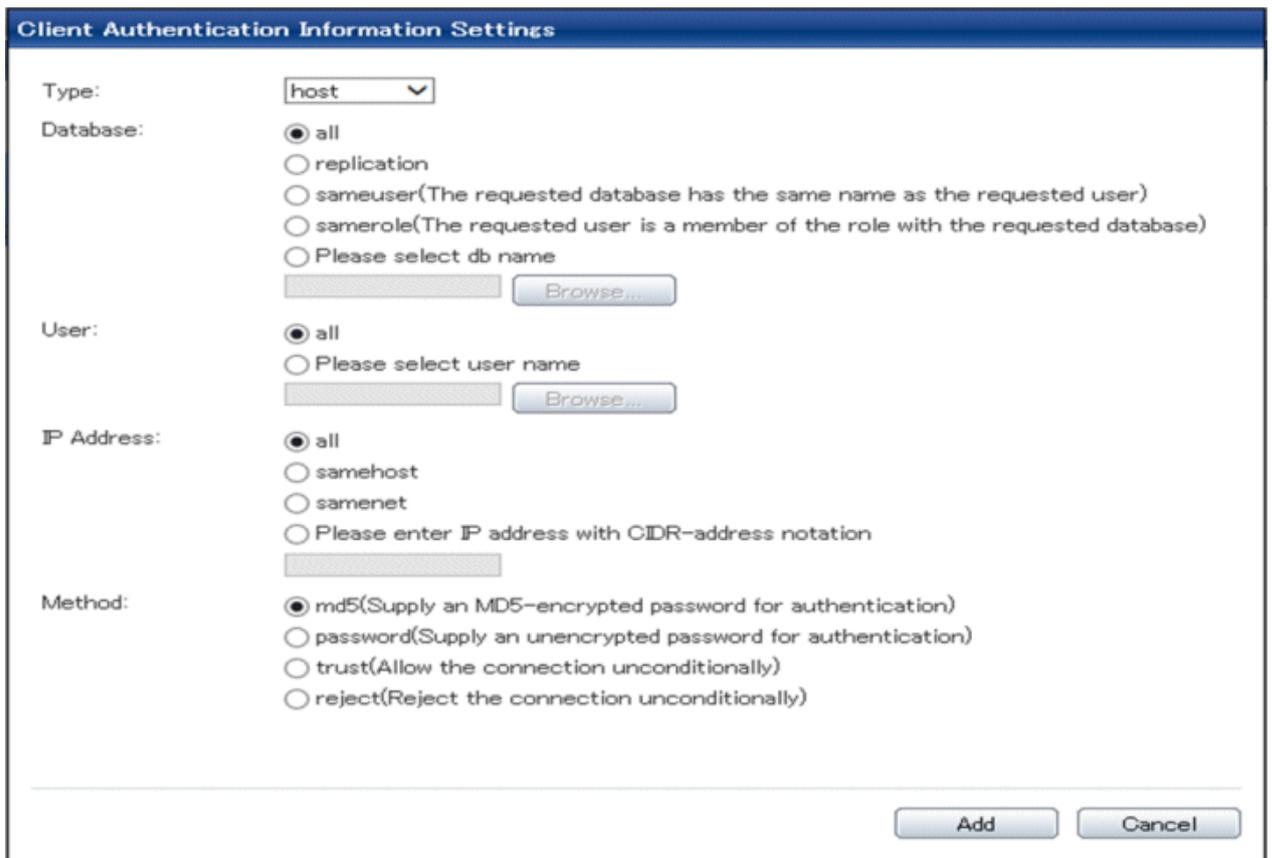
Type= host, Database=all, User=all, and Method=md5

The screenshot shows the Enterprise Postgres WebAdmin interface. The top navigation bar includes "Enterprise Postgres", "User: fsepuser", "Logout", and the Fujitsu logo. Below the navigation bar, there are tabs for "Monitor" and "Setting". The main content area is divided into two panels. The left panel, titled "Instance List", contains a "Create" button and a list of instances, with "inst1" selected and a "Delete" button next to it. The right panel, titled "Instance Name: inst1", shows the "Client Authentication" settings. It includes tabs for "Character Code", "Client Authentication", "Communication", "SQL Options", and "Memory". The "Client Authentication Information Settings" section contains a table with the following data:

	Type	Database	User	IP Address	Method
<input type="radio"/>	local	all	all		md5
<input type="radio"/>	host	all	all	127.0.0.1/32	md5
<input type="radio"/>	host	all	all	:::1/128	md5

Below the table are buttons for "Add", "Edit", and "Delete".

Click [Add] - the following window will be displayed.



The image shows a dialog box titled "Client Authentication Information Settings". It contains several configuration options:

- Type:** A dropdown menu with "host" selected.
- Database:** Radio buttons for "all" (selected), "replication", "sameuser(The requested database has the same name as the requested user)", "samerole(The requested user is a member of the role with the requested database)", and "Please select db name". Below the last option is a text input field and a "Browse..." button.
- User:** Radio buttons for "all" (selected) and "Please select user name". Below the second option is a text input field and a "Browse..." button.
- IP Address:** Radio buttons for "all" (selected), "samehost", "samenet", and "Please enter IP address with CIDR-address notation". Below the last option is a text input field.
- Method:** Radio buttons for "md5(Supply an MD5-encrypted password for authentication)" (selected), "password(Supply an unencrypted password for authentication)", "trust(Allow the connection unconditionally)", and "reject(Reject the connection unconditionally)".

At the bottom right of the dialog box are two buttons: "Add" and "Cancel".

Changing communication

1. Start WebAdmin and log into the database server.
2. Stop the instance.
3. Select [Communication] from the [Setting] menu.

4. Click [Edit].

The screenshot shows the Enterprise Postgres Administration console. The top navigation bar includes 'Enterprise Postgres', 'User: fsepuser', 'Logout', and the Fujitsu logo. Below the navigation bar are tabs for 'Monitor' and 'Setting'. The main content area is split into two panes. The left pane, titled 'Instance List', contains a 'Create' button and a table with one entry: 'inst1' with a 'Delete' button next to it. The right pane, titled 'Instance Name: inst1', shows configuration tabs: 'Character Code', 'Client Authentication', 'Communication', 'SQL Options', and 'Memory'. The 'Communication' tab is active, displaying 'Communication Settings' with a table for 'Port Number' (27500) and 'Max Connection' (100). An 'Edit' button is located at the bottom right of this pane.

5. Edit [Port Number] and [Max Connection], and then click [Modify].

The screenshot shows a 'Communication Settings' dialog box. It has two input fields: 'Port Number' with the value '27500' and 'Max Connection' with the value '100'. Below the fields is a note: 'The value of max_connections should include the number of connections(3) used by WebAdmin.' At the bottom of the dialog are 'Modify' and 'Cancel' buttons.

Note

- For the port number, specify an unused port number in the following range:
 - Windows Server(R) 2008, Windows Server(R) 2008 R2, Windows Server(R) 2012 or Windows Server(R) 2012 R2: 1024 to 49151
- Refer to "[4.4.3 Windows Firewall Settings](#)" for information on changing the port number.
- Calculate the maximum number of connections using the formula below:

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

*1: 3 is the default number of connections required by the system

Calculate the maximum number of connections using the following formula when changing either max_wal_senders (using streaming replication) or superuser_reserved_connections (connections reserved for use by the superuser) in postgresql.conf.

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

Refer to "[Appendix F Configuring Parameters](#)" for more information on postgresql.conf.

- Also check if the memory used exceeds the memory installed (refer to "[Parameters automatically set by WebAdmin according to the amount of memory](#)")

Changing the SQL option

1. Start WebAdmin and log in to the database server.
2. Select [SQL Options] from the [Setting] menu.
3. Click [Edit].

The screenshot shows the Enterprise Postgres WebAdmin interface. The top navigation bar includes "Enterprise Postgres", "User: fsepuser", "Logout", and the Fujitsu logo. Below the navigation bar, there are tabs for "Monitor" and "Setting". The "Setting" tab is active, and the "SQL Options" sub-tab is selected. The main content area is titled "Instance Name: inst1" and contains a table of SQL Options Settings. The table has two columns: the parameter name and its current value. An "Edit" button is located at the bottom right of the table.

SQL Options Settings	
Interpretation of Expression = NULL	Standard Format
Date Output Format	iso, mdy
Interval Output Format	postgres
The number of digits for floating values	0
Transaction Isolation Levels	read committed
Currency Format	C
Date and Time Format	C
Numerical Value Format	C

4. Edit [Interpretation of Expression = NULL], [Date Output Format], [Interval Output Format], [The number of digits for floating values], [Transaction Isolation Levels], [Currency Format], [Date and Time Format], and [Numerical Value Format], and then click [Modify].

SQL Options Settings

Interpretation of Expression = NULL: Standard Format Microsoft Access Format

Date Output Format: ISO, MDY(US,NonEuro,NonEuropean) ▼

Interval Output Format: postgres ▼

The number of digits for floating values : 0

Transaction Isolation Levels: read committed ▼

Currency Format: "C" ▼

Date and Time Format: "C" ▼

Numerical Value Format: "C" ▼

Modify Cancel

Note

If you select **postgres** in Date Output Format, dates will be output in the "12-17-1997" format, not the "Wed Dec 17 1997" format used in the PostgreSQL Documentation.

Changing allocated memory

1. Start WebAdmin and log in to the database server.
2. Stop the instance.
3. Select [Memory] from the [Setting] menu.

4. Click [Edit]

The screenshot shows the Enterprise Postgres WebAdmin interface. At the top, it displays 'Enterprise Postgres' and 'User: fsepuser'. Below the navigation bar, there are tabs for 'Monitor' and 'Setting'. The 'Instance List' on the left shows an instance named 'inst1' with a 'Delete' button. The main area is titled 'Instance Name: inst1' and has tabs for 'Character Code', 'Client Authentication', 'Communication', 'SQL Options', and 'Memory'. The 'Memory' tab is active, showing 'Memory Settings' with a table:

Sort Memory(Unit:KB)	24489KB
Share Buffer(Unit:KB)	4897792KB

An 'Edit' button is located at the bottom right of the table.

5. Edit [Sort Memory(Unit:KB)] and [Share Buffer(Unit:KB)], and then click [Modify].

The 'Memory Setting' dialog box is shown. It contains two input fields: 'Sort Memory(Unit:KB)' with the value '1024' and 'Share Buffer(Unit:KB)' with the value '32768'. At the bottom, there are 'Modify' and 'Cancel' buttons.

Note

Also check if the memory used exceeds the memory installed (refer to "[Parameters automatically set by WebAdmin according to the amount of memory](#)" for information on how to check).

4.3.2 Using the initdb Command

This section describes the procedure to create an instance using the initdb command.

Note

If a port is blocked (access permissions have not been granted) by a firewall, enable use of the port by granting access. Refer to the vendor document for information on how to grant port access permissions.
Consider the security risks carefully when opening ports.

4.3.2.1 Creating an Instance

Create an instance, with the database cluster storage destination specified in the PGDATA environment variable or in the -D option. Furthermore, the user that executed the initdb command becomes the instance administrator.

Note

- Instances created using the initdb command cannot be managed using WebAdmin.
- If creating multiple instances, ensure that there is no duplication of port numbers or the directories that store database clusters.

See

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

The procedure to create an instance is described below.

1. Use the OS user account that you want as the instance administrator.

Connect with the server using the OS user account that you want as the instance administrator.

2. Configure the environment variables

Configure the environment variables in the server with the newly created instance.

Set the following environment variables:

- PATH environment variables

Add installDir\bin and installDir\lib.

Example

The following is a setting example for environment variables in which "C:\Program Files\Fujitsu\fseserver64" is used as the installation folder:

```
> SET PATH=C:\Program Files\Fujitsu\fseserver64\bin;C:\Program Files\Fujitsu
\fseserver64\lib;%PATH%
```

3. Create a database cluster

Create the database cluster with the initdb command, specifying the storage destination directory.

Specify the transaction log storage destination and the locale setting option as required.

Example

```
> initdb -D D:\database\inst1 --xlogdir=E:\transaction\inst1 --lc-collate="C" --lc-ctype="C"
--encoding=UTF8
```

Point

In some features, instance names are requested, and those names are required to uniquely identify the instance within the system. These features allow names that conform to WebAdmin naming conventions, so refer to the following points when determining the names:

- Maximum of 16 characters
 - The first character must be ASCII alphabetic character
 - The other characters must be ASCII alphanumeric characters
-

Note

- To balance I/O load, consider deploying the transaction log storage destination to a disk device other than the database cluster storage destination and the backup data storage destination.
- Messages may not display correctly if a value other than "C" is specified as the display language for messages.
- Specify "C" for collation and character category. Performance deteriorates if you specify a value other than "C" , although the behavior will follow the rules for particular languages, countries and regions. Furthermore, this may need to be revised when running applications on systems with different locales.
For example, specify as follows:

```
initdb --locale="C" --lc-messages="C"  
initdb --lc-collate="C" --lc-ctype="C"
```

- Specify an encoding system other than SQL_ASCII for the database. If SQL_ASCII is used, there is no guarantee that the encryption system for data in the database will be consistent, depending on the application used to insert the data.
-

See

Refer to "Locale Support" in "Localization" in "Server Administration" in the PostgreSQL Documentation for information on locales.

4. Set port number.

Specify a port number in the port parameter of postgresql.conf. Ensure that the specified port number is not already used for other software. If a port number is not specified, "27500" is selected.

Register the specified port numbers in the /etc/services file if WebAdmin is used to create other instances. WebAdmin uses the services file to check if port numbers specified as available candidates have been duplicated.

Register any name as the service name.

Note

Make a note of the port number for use in the Windows firewall settings.

5. Set the corefile output destination.

Specify the output destination of the corefile, which can later be used to collect information for investigation, by setting the core_directory and core_contents parameters of postgresql.conf.

See

Refer to "Parameters" in the Operation Guide for information on the settings for these parameters.

6. Set the backup storage destination.

Specify the backup data storage destination and other backup settings when backup is to be performed as a provision against database errors.



See

Refer to "Backup Methods" in the Operation Guide for information on specifying backup settings.

7. Registering an instance in the Windows service

Use the register mode of the `pg_ctl` command to register an instance in the Windows service.

Specify the service name, user name, password and path to the instance in the `pg_ctl` command, and register the instance in the Windows service.

Example

The following is a setting example, in which the service name to register is "inst1", the user name is "fepuser", and the storage destination directory of the database cluster is "D:\database\inst1":

```
> pg_ctl register -N "inst1" -U fepuser -P ***** -D D:\database\inst1 -w
```



Note

- Registration should be performed by an instance administrator with "Administrator" privileges.
- For the following reasons, a user name and password must always be specified:

Because the Windows service is started up by the Network Service account, all user resources are created as resources of that account. This can result in error events such as failing to access database resources and not being able to perform backups/recovery.

Note that if not specifying a user name and password for security reasons, you should specify the account from the Windows services list immediately after registering the instance in Windows services.

- When entering the password that is specified in the `pg_ctl` command, for security reasons, you should be careful not to allow other users to access it.

This completes registration of an instance in the Windows service.

Commands such as `sc query` can be used to check the registration status.

8. Starting an instance

Use the following procedure to start the service:

- a. Display the [Services] window.

In Windows, select [Administrative Tools], and then click [Services].

- b. Start the service

From the services list, select the instance name that you wish to start, and click [Start Service].

If using commands to start the service, specify the service name using either the `net start` command or `sc start` command from the command prompt.

4.4 Configuring Remote Connections

This section describes the settings required when connecting remotely to Enterprise Postgres from a database application or a client command.

4.4.1 When an Instance was Created with WebAdmin

Settings related to connection

The default is to accept connections from remote computers to the database.

Change "listen_addresses" in postgresql.conf to change this.

Refer to "[Appendix F Configuring Parameters](#)" for more information on postgresql.conf.

Client Authentication Information settings

The following content is set by default when WebAdmin is used to create an instance.

- Authentication of remote connections from local machines is performed.

When changing Client Authentication Information, select [Client Authentication] from [Setting], and then change the settings.

4.4.2 When an Instance was Created with the initdb Command

Connection settings

The default setting only permits local connections from the client to the database. Remote connections are not accepted.

Change "listen_addresses" in postgresql.conf to perform remote connection.

All remote connections will be allowed when changed as shown below.

Example

```
listen_addresses = '*'
```

Also, configure the parameters shown below in accordance with the applications and number of client command connections.

Parameter name	Parameter description
superuser_reserved_connections	Number of connections reserved for database maintenance, for example backup or index rebuilding. If you need to simultaneously perform a large number of processes that exceed the default value, change this value accordingly.
max_connections	Set the value as: <i>numberOfSimultaneousConnectionsToInstance</i> + superuser_reserved_connections

Client authentication information settings

When trying to connect from a client to a database, settings are required to determine whether the instance permits connections from the client - if it does, then it is possible to make settings to determine if authentication is required.



See

.....
Refer to "The pg_hba.conf File" in "Server Administration" in the PostgreSQL Documentation for details.
.....

4.4.3 Windows Firewall Settings

If the Windows firewall feature is to be enabled, you should enable a port number on the database server. The following explains how to enable a port number:

Windows Server(R) 2008:

1. In the [Start] menu, click [Control Panel].
2. Click [Security] and then click [Windows Firewall].
3. In the [Windows Firewall] window, click [Change settings].
4. On the [Exceptions] tab, click [Add port].

5. In the [Add a Port] window, set the following:
 - a. In [Name], specify the desired name.
 - b. In [Port number], specify the Web server port number that was specified when creating the instance.
 - c. Select [TCP].
6. Click [OK].
7. On the [Exceptions] tab, in the [Program or port] list, check if the added port is enabled.
8. In the [Windows Firewall] window, click [OK].

Windows Server(R) 2008 R2:

1. In the [Start] menu, click [Control Panel].
2. Click [System and Security], and then click [Windows Firewall].
3. In the [Windows Firewall] window, click [Advanced settings] on the left side of the window.
4. In the [Windows Firewall with Advanced Security] window, click [Inbound Rules] on the left side of the window.
5. Click [New Rule] on the right side of the window.
6. In the [New Inbound Rule Wizard] window, select [Port], and then click [Next].
7. Select [TCP] and [Specific local ports], then specify the Web server port number specified when creating the instance, and then click [Next].
8. Select [Allow the connection], and then click [Next].
9. Select the profiles for which this rule applies, and then click [Next].
10. In [Name], specify the desired name, and then click [Finish].
11. In the [Windows Firewall with Advanced Security] window, check if the added rule is enabled under [Inbound Rules] in the center of the window.

Windows Server(R) 2012 or Windows Server(R) 2012 R2:

1. Right-click the [Start] screen, and then click [All apps] that is displayed in the bottom-right of the screen. In the [Apps] window, click [Control Panel].
2. Click [Windows Firewall].
3. In the [Windows Firewall] window, click [Advanced settings] on the left side of the window.
4. In the [Windows Firewall with Advanced Security] window, click [Inbound Rules] on the left side of the window.
5. Click [New Rule] on the right side of the window.
6. In the [New Inbound Rule Wizard] window, select [Port], and then click [Next].
7. Select [TCP] and [Specific local ports], then specify the Web server port number specified when creating the instance, and then click [Next].
8. Select [Allow the connection], and then click [Next].
9. Select the profiles for which this rule applies, and then click [Next].
10. In [Name], specify the desired name, and then click [Finish].
11. In the [Windows Firewall with Advanced Security] window, check if the added rule is enabled under [Inbound Rules] in the center of the window.

4.5 Other Settings

This section describes settings that are useful for operations.

4.5.1 Error Log Settings

This section explains the settings necessary to monitor errors in applications and operations, and to make discovering the causes easier.

Make error log settings only when instances are created with the `initdb` command.

When creating instances with WebAdmin, these settings are already made and hence do not need to be set. Furthermore, some parameters are used by WebAdmin, and if changed, may cause WebAdmin to no longer work properly. Refer to "[Appendix F Configuring Parameters](#)" for details.

Edit the following parameters in `postgresql.conf`:

Parameter name	Parameter description	How to enable the settings
<code>syslog_ident</code>	If using multiple instances, specify an event source name to be appended to messages in order to identify messages output to the event log.	<ul style="list-style-type: none"> - Restart services from the Windows services window. - Use the <code>net</code> command or <code>sc</code> command to stop and start services.
<code>event_source</code>	Specify the event source name to be attached to messages, for identifying messages output to the event log when using multiple instances.	<ul style="list-style-type: none"> - Restart services from the Windows services window. - Use the <code>net</code> command or <code>sc</code> command to stop and start services.
<code>logging_collector</code>	Specify "on" to ensure that messages are output by Enterprise Postgres to the server log file. The server log file is created in the <code>pg_log</code> directory in the database cluster.	<ul style="list-style-type: none"> - Restart services from the Windows services window. - Use the <code>net</code> command or <code>sc</code> command to stop and start services.
<code>log_destination</code>	Specify "stderr, eventlog" to output messages from Enterprise Postgres to the screen and either the system log or the event log.	reload option of the <code>pg_ctl</code> mode
<code>log_line_prefix</code>	<p>Specify information to be added at the start of messages output by an instance. This information is useful for automatic monitoring of messages.</p> <p>You can output the SQLSTATE value, output time, executing host, application name, and user ID.</p> <p>Refer to "What To Log" in the PostgreSQL Documentation for details.</p> <p>Example: <code>log_line_prefix = '%e: %t [%p]: [%l-1] user = %u,db = %d,remote = %r app = %a '</code></p>	reload option of the <code>pg_ctl</code> mode

Point

- If you want fewer application errors being output to the eventlog, refer to "When To Log" and "What To Log" in the PostgreSQL Documentation for information on how to reduce the output messages.
- If you want to separate errors output from other software, refer to "Where To Log" in the PostgreSQL Documentation to change the output destination to the server log file rather than the system log.

4.5.2 Configuring Automatic Start and Stop of an Instance

You can automatically start or stop an instance when the operating system on the database server is started or stopped.

Use the following procedure to configure automatic start and stop of an instance.

Note that, if an instance is activated in a failover operation, the cluster system will control the start or stop, therefore this feature should not be used. Also, when performing database multiplexing, refer to "Enabling Automatic Start and Stop of Mirroring Controller and Multiplexed Instances" in the Cluster Operation Guide.

When an instance was created with WebAdmin

If an instance is created with WebAdmin, the instance is registered in the Windows service and automatic start and stop is set for the instance.

To change the automatic start and stop setting for an instance, select the service for the applicable instance in the Windows services window, and in [Startup Type], select [Automatic] or [Manual].

When an instance was created with the initdb command

If the startup type of the service is set to [Manual], change it to [Automatic]. By setting the startup type to [Automatic], the service will start up automatically when the Windows(R) system is started up, and will stop automatically when the Windows(R) system is shut down.



Note

The settings should be performed by an instance administrator with "Administrator" privileges.

Use the following procedure to switch the service:

1. Display the [Services] window.
In Windows, select [Administrative Tools], and then click [Services].
2. Switch the startup type
Select the Enterprise Postgres service name, display the [Properties] dialog box, and then switch the startup type from [Manual] to [Automatic].

The above setting can also be changed using the sc config command.

4.5.3 Settings when Using the features compatible with Oracle databases

The compatibility feature for Oracle databases enables Enterprise Postgres to be used without any special settings by creating a database instance. Note that settings are required only when using the SUBSTR function.



See

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

4.6 Integration with Message-Monitoring Software

To monitor messages output by Enterprise Postgres using software, configure the product to monitor SQLSTATE, instead of the message text - this is because the latter may change when Enterprise Postgres is upgraded.

Configure Enterprise Postgres to output messages in a format that can be read by the message-monitoring software by specifying "%e" in the log_line_prefix parameter of postgresql.conf to output the SQLSTATE value.

A setting example is shown below - it outputs the output time, executing host, application name, and user ID, in addition to the SQLSTATE value.

Example

```
log_line_prefix = '%e: %t [%p]: [%l-1] user = %u,db = %d,remote = %r app = %a '
```



See

Refer to "What To Log" in the PostgreSQL Documentation for information on how to configure the settings.

4.7 Deleting Instances

This section explains how to delete instances.

- [4.7.1 Using WebAdmin](#)
- [4.7.2 Using Server Commands](#)



Note

- Always use WebAdmin to delete instances that were created using it. Because WebAdmin management information cannot be deleted, WebAdmin will determine that the instance is abnormal.

4.7.1 Using WebAdmin

This section explains how to delete instances using WebAdmin.

Use the following procedure to delete instances.

1. Stop the instance

In the [Monitor] window, click [Stop].

2. Back up files.

Before deleting the instance, back up any required files under the data storage destination, the backup data storage destination, and the transaction log storage destination.

3. Delete the instance

In the [Instance List] window, select the desired instance, and then click [Delete] - in the confirmation window, click [Delete].



Note

Deleting an instance deletes only the following lowest-level directories. If they are not required, delete them manually.

- Data storage destination
- Backup data storage destination
- Transaction log storage destination (if different from the data storage destination)

4.7.2 Using Server Commands

This section explains how to delete instances using server commands.

Use the following procedure to delete instances.

1. Stop the instance

Stop the instance in Windows services, or use the `pg_ctl` command `stop mode`.

Use the following procedure to stop a service in Windows services:

- a. Display the [Services] window.

In Windows, select [Administrative Tools], and then click [Services].

b. Stop the service

Select the instance name that you wish to stop from the services list, and click [Stop Service]. If you stop a service while applications and commands are running, Enterprise Postgres will force those applications and commands to close and will stop normally.

You can also stop a service by specifying the service name in the net stop command or sc stop command.

2. Back up files

Before deleting the instance, you should back up all necessary files contained in the data storage destination, backup data storage destination, and transaction log storage destination.

3. Delete the instance

Use a command such as rmdir to delete the following directories:

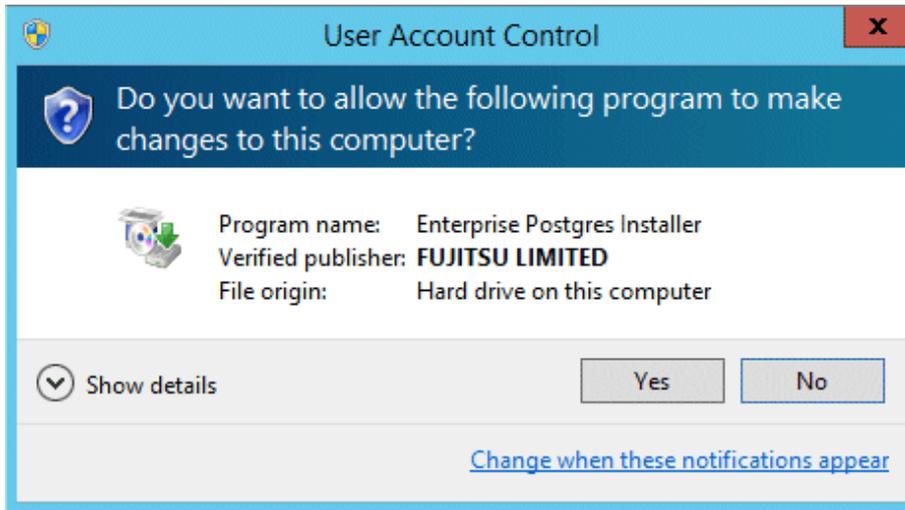
- Data storage destination directory
- Backup data storage destination directory
- Transaction log storage destination directory (if a different directory to the data storage destination directory was specified)

Chapter 5 Uninstallation

This chapter explains the uninstallation of Enterprise Postgres.

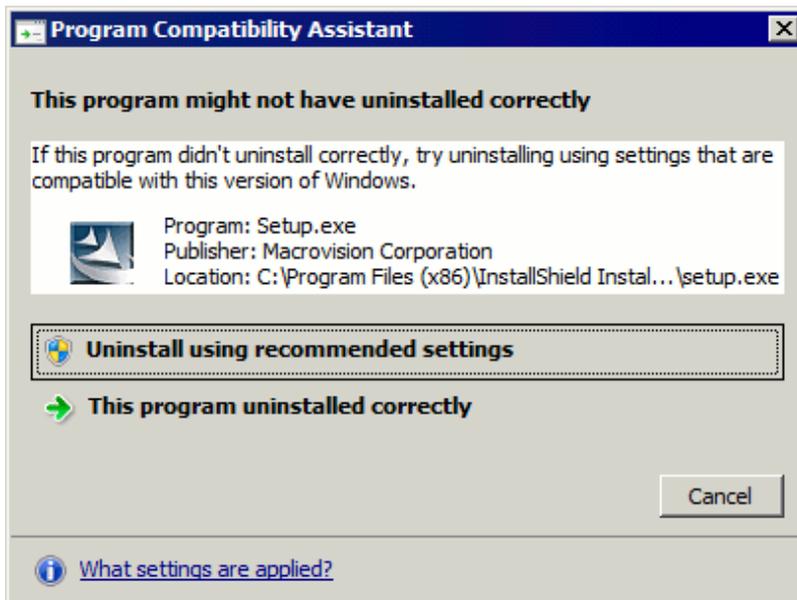
Information

- If a [User Account Control] dialog box such as the following is displayed at the start of the uninstallation, click [Yes] to continue processing:



If [No] is clicked, permission to continue is denied and an [Error] dialog box will be displayed. To continue the uninstallation, click [Retry] in the [Error] dialog box. To end the operation, click [Cancel].

- If uninstallation is suspended or processing terminates abnormally, a dialog box of the Program Compatibility Assistant similar to the one shown below may be displayed. Click [This program uninstalled correctly] and continue operation.



5.1 Uninstallation in Interactive Mode

The example described here shows the procedure for uninstalling Enterprise Postgres Standard Edition 64-bit from your system.

The uninstallation procedure is described below.

Note

- If performing operation with WebAdmin, back up the following file before uninstallation. Instances will not be recognized by WebAdmin even if Enterprise Postgres is reinstalled after uninstallation. If performing operation with WebAdmin after reinstalling Enterprise Postgres, replace the following backed up file after installation.

`install\Folder\gui\data\gui-instance`

By replacing the above file in the installation folder after installation, the instance will be recognized by WebAdmin, and the recognized instance will be set to automatically start and stop.

To disable the automatic start and stop setting for an instance, select the service for the applicable instance in the Windows services window, and in [Startup Type], select [Manual].

- If a PL/extJava environment has been built, delete the domain. Refer to "Setting up and Operating PL/extJava" in the Operation Guide for details.
- If using database multiplexing mode, refer to "Uninstalling in Database Multiplexing Mode" in the Enterprise Postgres Cluster Operation Guide before performing the uninstallation.

See

Refer to the Installation and Setup Guide for Client when uninstalling the Enterprise Postgres client feature.

Information

If an error occurs while the product is being uninstalled, refer to "[Appendix H Uninstall \(middleware\) Messages](#)" and take the required action.

1. Stop applications and the client program

Before starting the uninstallation, stop the following:

- Applications that use the product
- pgAdmin

2. Stop instances

Stop all instances that are using the product to be uninstalled.

Stopping of instances should be performed by the appropriate instance administrator.

When an instance was created with WebAdmin

In the [Monitor] window, click [Stop].

When an instance was created with the initdb command

Use the following procedure to stop a service:

- a. Display the [Services] window
In Windows, select [Administrative Tools], and then click [Services].

- b. Stop the service

Select the instance name that you wish to stop from the services list, and click [Stop Service]. If you stop a service while applications and commands are running, Enterprise Postgres will force those applications and commands to close and will stop normally.

You can also stop a service by specifying the service name in the net stop command or sc stop command.

3. Stop WebAdmin

If you are using WebAdmin, stop WebAdmin.

Refer to "[E.1.3 Stopping the Web Server Feature of WebAdmin](#)" for details.

4. Unregister Windows services

Perform this step if the instance was created with the initdb command.

Unregister the instance registered in Windows services.

Use the unregister mode of the pg_ctl command to specify the registered service name and unregister the instance from Windows services.

Example

The following is an example showing execution of this command on the registered service name "inst1".

```
> pg_ctl unregister -N "inst1"
```



You should unregister services before uninstalling Enterprise Postgres. If you uninstall Enterprise Postgres while services are running, several files will remain after the uninstallation.

If you have carried out the uninstallation without unregistering services beforehand, use the server command sc delete to unregister the services.

5. Delete registrations related to the event log

If you are outputting to the event log, a DLL registration mentioned in "[4.2.5 Preparing for Output to the Event Log](#)" has been performed.

To prevent unnecessary issues from occurring, you should delete this registration. Refer to "Server Setup and Operation", "Registering Event Log on Windows" in the PostgreSQL Documentation for details.

The following is an example showing deletion of the DLL registration for a 64-bit product under the default event source name.

```
regsvr32 /u "c:\Program Files\Fujitsu\fsepsvr64\lib\pgevent.dll"
```

If using multiple instances

DLL registration is performed so that you can output messages corresponding to the event source name assigned by the user, allowing you to identify messages output to the event log by instance.

Since it is necessary to delete the DLL registration for each instance, delete the DLL registration by event source name.

The following is an example showing deletion of the DLL of a 64-bit product registered under the event source name "Enterprise Postgres inst1".

```
regsvr32 /u /i:"Enterprise Postgres inst1" "C:\Program Files\Fujitsu\fsepsvr64\lib\pgevent.dll"
```

Note that this step is not required if the instance was created with WebAdmin.

If installing multiple versions

If the instances you created using this package have been set to output error logs to the event log, use the DLL path name that you took note of previously as explained in "[4.2.5 Preparing for Output to the Event Log](#)" to reregister the default event source name.



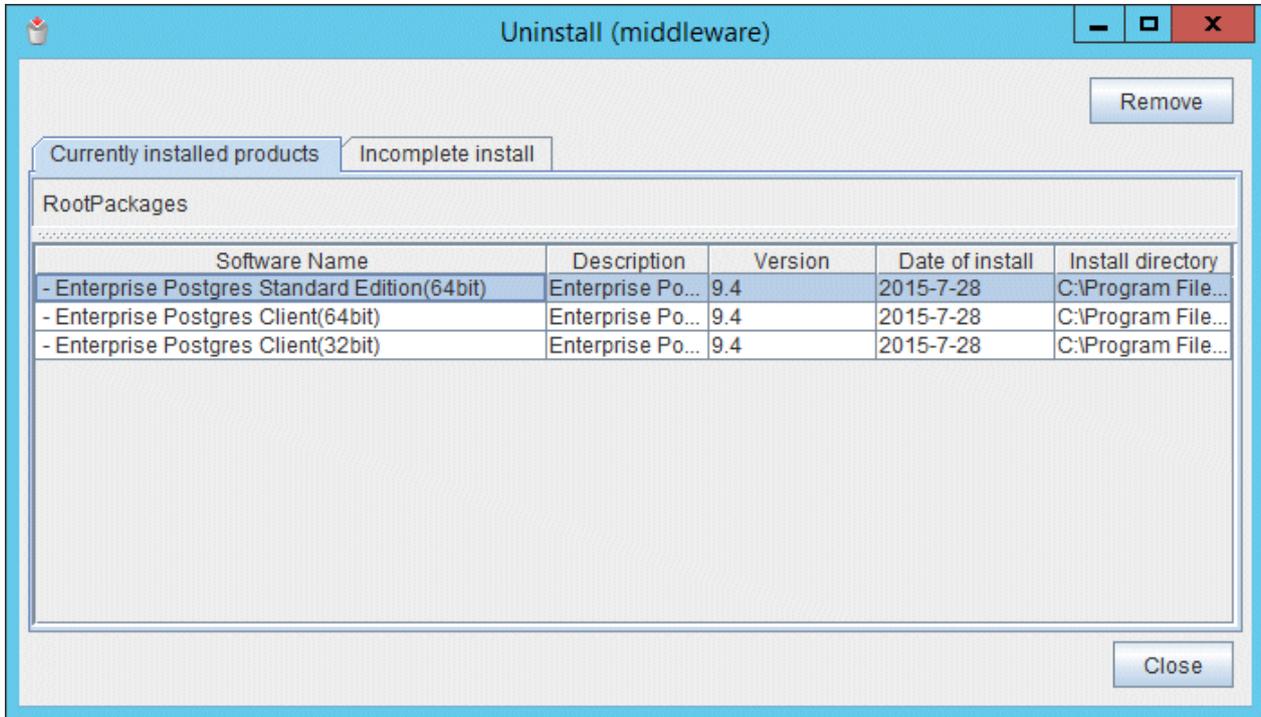
Ensure to delete DLLs before the uninstallation. If you perform the uninstallation without doing so, you may not be able to delete the DLLs at a later time.

6. Start the Uninstall (middleware) tool

In Windows, click [All Programs] or [All apps], then [Fujitsu], and then [Uninstall (middleware)].

7. Select the software

Select the product to be uninstalled from [Software Name], and then click [Remove].



8. Start the uninstallation

Click [Uninstall].



9. Finish the uninstallation

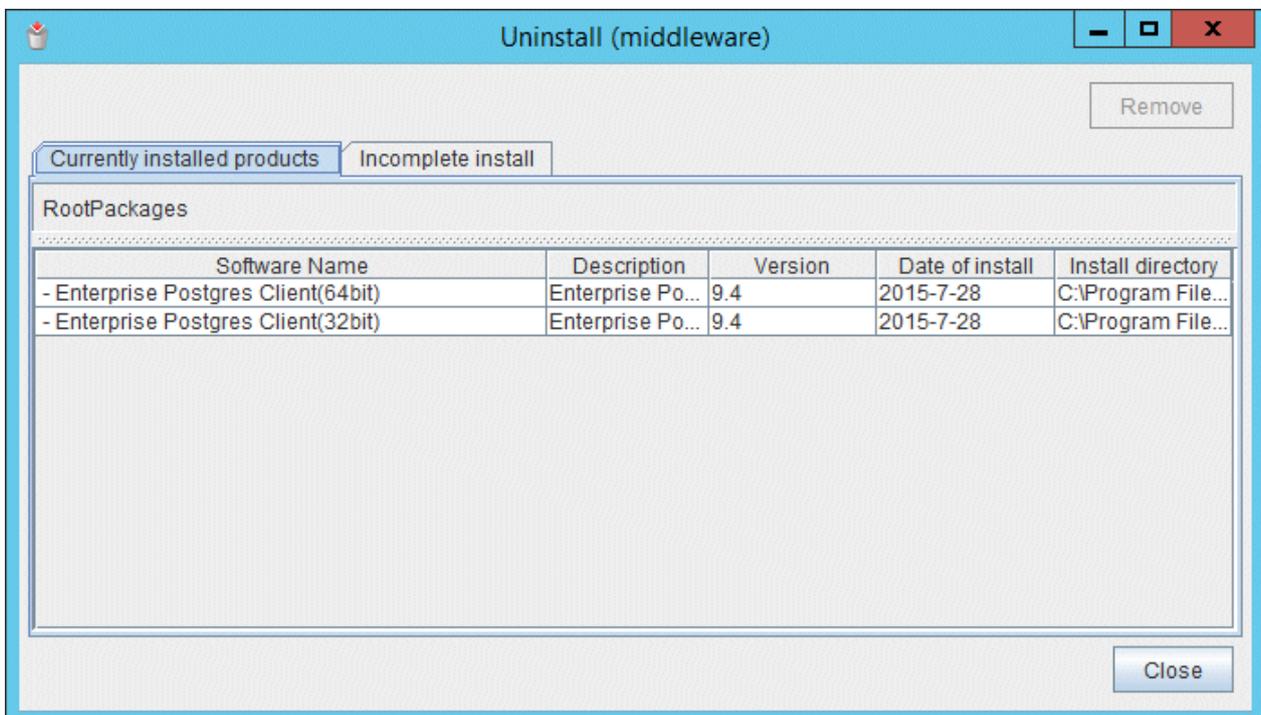
The uninstallation completion window will be displayed. Click [Finish].

The installation folder may remain after uninstallation. If it is not required, delete it.



10. Stop the Uninstall (middleware) tool

Click [Close].



5.2 Uninstallation in Silent Mode

The uninstallation procedure is described below.

Note

- If performing operation with WebAdmin, back up the following file before uninstallation. Instances will not be recognized by WebAdmin even if Enterprise Postgres is reinstalled after uninstallation. If performing operation with WebAdmin after reinstalling Enterprise Postgres, replace the following backed up file after installation.

install\Folder\gui\data\gui-instance

By replacing the above file in the installation folder after installation, the instance will be recognized by WebAdmin, and the recognized instance will be set to automatically start and stop.

To disable the automatic start and stop setting for an instance, select the service for the applicable instance in the Windows services window, and in [Startup Type], select [Manual].

- If a PL/extJava environment has been built, delete the domain. Refer to "Setting up and Operating PL/extJava" in the Operation Guide for details.
- If using database multiplexing mode, refer to "Uninstalling in Database Multiplexing Mode" in the Enterprise Postgres Cluster Operation Guide before performing the uninstallation.

See

Refer to the Installation and Setup Guide for Client when uninstalling the Enterprise Postgres client feature.

1. Stop applications and the client program

Before starting the uninstallation, stop the following:

- Applications that use the product
- pgAdmin

2. Stop all instances

Stop all instances that are using the product to be uninstalled.

When an instance was created with WebAdmin

In the [Monitor] window, click [Stop].

When an instance was created with the initdb command

Use the following procedure to stop a service:

- a. Display the [Services] window.
In Windows, select [Administrative Tools], and then click [Services].
- b. Stop the service

Select the instance name that you wish to stop from the services list, and click [Stop Service]. If you stop a service while applications and commands are running, Enterprise Postgres will force those applications and commands to close and will stop normally.

You can also stop a service by specifying the service name in the net stop command or sc stop command.

3. Stop WebAdmin

If you are using WebAdmin, stop WebAdmin.

Refer to "[E.1.3 Stopping the Web Server Feature of WebAdmin](#)" for details.

4. Unregister Windows services

Perform this step if the instance was created with the initdb command.

Unregister the instance registered in Windows services.

Use the unregister mode of the pg_ctl command to specify the registered service name and unregister the instance from Windows services.

Example

The following is an example showing execution of this command for the registered service name "inst1".

```
> pg_ctl unregister -N "inst1"
```



You should unregister services before uninstalling Enterprise Postgres. If you uninstall Enterprise Postgres while services are running, several files will remain after the uninstallation.

If you have carried out the uninstallation without unregistering services beforehand, use the server command sc delete to unregister the services.

5. Delete registrations related to the event log

If you are outputting to the event log, a DLL registration mentioned in ["4.2.5 Preparing for Output to the Event Log"](#) has been performed.

To prevent unnecessary issues from occurring, you should delete this registration. Refer to "Server Setup and Operation", "Registering Event Log on Windows" in the PostgreSQL Documentation for details.

The following is an example showing deletion of the DLL registration for a 64-bit product under the default event source name.

```
regsvr32 /u "c:\Program Files\Fujitsu\fsepserver64\lib\pgevent.dll"
```

If using multiple instances

DLL registration is performed so that you can output messages corresponding to the event source name assigned by the user, allowing you to identify messages output to the event log by instance.

Since it is necessary to delete the DLL registration for each instance, delete the DLL registration by event source name.

The following is an example showing deletion of the DLL of a 64-bit product registered under the event source name "Enterprise Postgres inst1".

```
regsvr32 /u /i:"Enterprise Postgres inst1" "c:\Program Files\Fujitsu\fsepserver64\lib\pgevent.dll"
```

Note that this step is not required if the instance was created with WebAdmin.

If installing multiple versions

If the instances you created using this package have been set to output error logs to the event log, use the DLL path name that you took note of previously as explained in ["4.2.5 Preparing for Output to the Event Log"](#) to reregister the default event source name.



Ensure to delete DLLs before the uninstallation. If you perform the uninstallation without doing so, you may not be able to delete the DLLs at a later time.

6. Start the command prompt

In Windows, right-click [Command Prompt] and then select [Run as administrator].

7. Run the uninstaller

Execute the command below.

The installation directory may remain after uninstallation. If it is not required, delete it.

Example

```
> installFolder\suninst.bat
```

X: Drive on which the product is installed

 Note

In Windows(R)7, Windows Server(R) 2008 or Windows Server(R) 2008 R2, an empty folder may remain in the [Start] menu after uninstallation. In Windows, select [All Programs], right-click the following, and then click [Delete]:

- Enterprise Postgres Client (32-bit) 9.4
- Enterprise Postgres Client (64-bit) 9.4

8. Check the results

Uninstaller results will be output to the log file.

Log file

xyz is the number part when the version is x.y.z.

```
pathName: %TEMP%\fsep_SERVER64_0xyz_uninstall.log
```

Return values

The following return values are output:

Return value	Description
0	Uninstallation was successful.
64	Failed to install because "Uninstall (middleware)" is running.
74	Failed to write product information (updating registry).
79	Failed to read product information (viewing registry).
83	Failed to delete the information from FJQSS.
129	Unsupported operating system.
144	Not executed with administrator privileges.
251	An error occurred when running a system function.
255	An internal inconsistency occurred.

Appendix A Installation in Silent Mode

This appendix provides specifications for installation in silent mode.

A.1 Specification Format

The installation parameters CSV file, which is specified as the argument for the silent installer, has three columns per line in CSV format.

```
sectionName, parameterName, value
sectionName, parameterName, value
:
```

Enter the following settings in respective columns.

Item	Settings	Optional
<i>sectionName</i>	Specify the section name. There are two types of section names: "installinfo": Set the product information. "parameters": Set the parameter information for this product.	Mandatory
<i>parameterName</i>	Specify the parameter name. Each section has a valid parameter.	Mandatory
<i>value</i>	Specify the value.	Optional

Note

- Blank lines cannot be included.
- Section names and parameter names cannot be omitted.
- Undefined parameters cannot be set in lines where the section name is "installInfo". Also, note that the same parameter cannot be specified multiple times.
- Specify at least one line with the section name "parameters".
- Undefined parameters specified in lines within the "parameters" section will be ignored during execution. Note that when the same parameter is specified multiple times, the settings in the lowest line will be valid.
- The setting values for lines where the section name is "installInfo" may contain alphanumeric characters (at least one), and symbols, excluding double quotation marks (") and commas (,).
- The following characters can be used for the setting value in the line of the section name "parameters":
[a-zA-Z0-9], -, and \

Information

The installation parameter CSV template is "(driveInWhichDvdIsSet)\sample\sample.csv".

A.2 List of Parameters

This section describes the parameters that can be set for each section.

installInfo section

The parameters that can be set in the installInfo section are shown below.

No.	Type	Parameter		Description
1	Software ID	Parameter name	Name	Mandatory. Specify the software identifier. For this software, specify the following value: "Enterprise Postgres"
		Value/Range	String	
		Default value	None	

Information

- In addition to "Name", the following parameters can be used in the "installInfo" section.
 - softwareName
 - OS
 - Version
 - Edition
- Note that specifying a parameter name other than "Name" and the parameters listed above will result in an error.
- The values set in these parameters do not affect silent installation.

Example

```
installInfo,Name,Enterprise Postgres
```

parameters section

The parameters that can be set in the parameters section are shown below.

No.	Type	Parameter		Description
1	Whether to install the server	Parameter name	ServerInstallExecute	Optional. Specify whether to run the installation of the server. Y: Install No: Do not install
		Value/Range	Y or N	
		Default value	Y	
2	Installation destination of the server	Parameter name	ServerInstallPath	Optional. Specify the installation destination of the software. The root folder (such as C:\) cannot be specified. PL/extJava, specify the directory to be defined in the path name using up to 50 bytes.
		Value/Range	Path name	
		Default value	C:\Program Files\Fujitsu \fsepsserver64	
3	Whether to install the client (32-bit)	Parameter name	Client32InstallExecute	Optional. Specify whether to run the installation of the client (32-bit). Y: Install N: Do not install
		Value/Range	Y or N	
		Default value	Y	
4	Installation destination of the client (32-bit)	Parameter name	Client32InstallPath	Optional. Specify the installation destination of the client (32-bit).
		Value/Range	Path name	

No.	Type	Parameter		Description
		Default value	C:\Program Files\Fujitsu \fsepclient32 or C:\Program Files (x86)\Fujitsu \fsepclient32	The root folder (such as C:\) cannot be specified.
5	Whether to install the client (64-bit)	Parameter name	Client64InstallExecute	Optional. Specify whether to run the installation of the client (64-bit). Y: Install N: Do not install This parameter is ignored on 32-bit operating systems.
		Value/Range	Y or N	
		Default value	Y	
6	Installation destination of the client (64-bit)	Parameter name	Client64InstallPath	Optional. Specify the installation destination of the client (64-bit). This parameter is ignored on 32-bit operating systems. The root folder (such as C:\) cannot be specified.
		Value/Range	Path name	
		Default value	C:\Program Files\Fujitsu \fsepclient64	
7	Whether to install Mirroring Controller	Parameter name	MirroringControllerInstallExecute	Optional. Specify whether to run the installation of Mirroring Controller. Y: Install No: Do not install This will be ignored in all cases
		Value/Range	Y or N	
		Default value	N	
8	Whether to set up WebAdmin	Parameter name	WebSetupExecute	Optional. Specify whether to set up WebAdmin. Y: Install N: Do not install
		Value/Range	Y or N	
		Default value	Y	
9	Web server port number	Parameter name	WebPortNumber1	Optional. Specify the port number of the Web server.
		Value/Range	1024 to 32767	
		Default value	27515	
10	WebAdmin internal port number	Parameter name	WebPortNumber2	Optional. Specify the WebAdmin internal port number.
		Value/Range	1024 to 32767	
		Default value	27516	

Example

```
parameters,ServerInstallExecute,Y
parameters,ServerInstallPath,C:\Program Files\Fujitsu\fsepserver64
parameters,Client32InstallExecute,Y
parameters,Client32InstallPath,C:\Program Files (x86)\Fujitsu\fsepclient32
parameters,Client64InstallExecute,Y
parameters,Client64InstallPath,C:\Program Files\Fujitsu\fsepclient64
parameters,WebSetupExecute,Y
```

```
parameters,WebPortNumber1,27515
parameters,WebPortNumber2,27516
```

A.3 Messages and Return Values

Messages are output when errors are detected during parametric analysis.

If an error is detected during installation of the product, a message is output to the log:

Log file

xyz is the number part when the version is x.y.z.

```
%TEMP%\fsep_SERVER64_media_0xyz_install.log
```

Messages and return values

CSV file errors

The following messages are output if errors are detected while parsing CSV files.

Return value	Message	Explanation and actions
4	CSV file error:code = 1, Invalid CSV error.	There is an error in the specification format of the CSV file.
4	CSV file error:code = 2, installInfo/Name is required.	Either installInfo or the Name parameter has not been specified.
4	CSV file error:code = 3, Invalid installInfo key.	There is an error in the installInfo specification. Or the section name is invalid.
4	CSV file error:code = 4, Duplicated installInfo key.	The same parameter has been defined more than once in installInfo.
4	CSV file error:code = 5, Invalid character length.	No setting value is specified, or the specified string is too long.
4	CSV file error:code = 6, Invalid character format or encoding.	An invalid format or character set is being used for the character string.
4	CSV file error:code = 8, Parameter is required.	There is no line in the "parameters" section.
20	The input file does not exist.	The input file does not exist.
21	The value of @1@ is incorrect,the value is @2@.	The value is incorrect. Specify the correct value. The parameter name is displayed in @1@. The specified parameter is displayed in @2@.
22	The value same at @1@ and @2@ is specified.	The same value is specified in different parameters. Specify different values. The parameter name is displayed in @1@ and @2@.
23	The @2@ of @1@ already exists.	The path already exists. Specify a different path. The parameter name is displayed in @1@. The specified value is displayed in @2@.
24	Drive @2@ of @1@ does not exist.	The drive does not exist. Specify an existing drive. The parameter name is displayed in @1@. The specified value is displayed in @2@.

Return value	Message	Explanation and actions
26	Port number @1@ is already used in the service file(/etc/services).	The port number is already being used. Specify an unused port number. The port number is displayed in @1@
28	There is no product for installation.	No product was specified for installation. Specify "Y" for at least one of the following parameters: - ServerInstallExecute - Client32InstallExecute - Client64InstallExecute
29	USAGE : silent.sh inputfile	The argument specified in the command is incorrect. Specify the correct argument.

Product installer errors

The following return values are returned when errors occur while the product installer is running.

Return value	Explanation
0	Completed successfully.
3	Failed to set up WebAdmin.
14	The resources required for installation do not exist on the media.
32	The same version of the same product is already installed.
60	An error occurred in Uninstall (middleware).
64	Installation failed because Uninstall (middleware) is running.
65	Failed to install Uninstall (middleware).
74	Failed to write to the registry.
79	Failed to read the registry.
85	Different editions of the same generation are installed.
86	A new StopChange version of the same generation is installed.
88	A conflicting product is installed.
129	Unsupported operating system.
142	The installation destination has insufficient disk capacity.
144	The command was not executed with administrator privileges.
220	Failed to create a directory in the installation directory.
221	Failed to create files in the installation directory.
251	An error occurred when running a system function.
255	An internal inconsistency occurred.

A.4 CSV File Format

The format of CSV files is based on RFC4180, with the following specifications.

Records

- Separate each record with a "CRLF" newline (operation is not guaranteed with only a "CR" or "LF" newline).
- Specify a newline at the end of a file.

- Separate each field within a record with a comma ",".

Format	Record		
aaa,bbb,ccc	aaa	bbb	ccc

- If several commas are entered in succession, or if a comma precedes a newline, the data following the comma is regarded as empty.

Format	Record		
aaa,,ccc	aaa		ccc
aaa,bbb,	aaa	bbb	

- Headers cannot be specified

Format	Record		
	field1	field2	field 3
aaa,bbb,ccc	aaa	bbb	ccc

Appendix B Enterprise Postgres Migration Procedures

The tasks shown in the following table are required to migrate from an older version/level:

Table B.1 Reference destination for migration tasks by product

Product before migration	Product after migration	Operation mode	Refer to:
Enterprise Postgres Standard Edition 9.2	Enterprise Postgres Standard Edition 9.4	Basic operations	B.1

B.1 Migration Tasks for Basic Operations

If the Enterprise Postgres version that was used to create the instance is 9.2, execute "CREATE EXTENSION" in order to reference Information added by Enterprise Postgres.



See

Refer to "Information that can be Collected" in the Operation Guide for information on Information added by Enterprise Postgres.

Follow the procedure below to install any missing extensions:

1. Check if CREATE EXTENSION is required for the target database.

Use the psql command to connect to the target database (in the example below, "postgres"), and then execute the following SELECT statements:

```
postgres=# SELECT * FROM pgx_stat_sql;
ERROR: Relation "pgx_stat_sql" does not exist (10292)
postgres=# SELECT * FROM pgx_stat_lwlock;
ERROR: Relation "pgx_stat_lwlock" does not exist (10292)
postgres=# SELECT * FROM pgx_stat_latch;
ERROR: Relation "pgx_stat_latch" does not exist (10292)
postgres=# SELECT * FROM pgx_stat_walwriter;
ERROR: Relation "pgx_stat_walwriter" does not exist (10292)
```

CREATE EXTENSION is required if the view specified in the SELECT statement does not exist.

2. Execute the statements below to install the statistics views independently added by Enterprise Postgres and that were not displayed in step 1. If there are multiple databases, execute the statements for all of them.

If the connected database is not a registration target, these views cannot be browsed. If the database has been registered for template1, there is no need to consider this point for databases added thereafter.

The statistics are reset when the view is installed.

```
postgres=# CREATE EXTENSION pgx_stat_sql;
CREATE EXTENSION
postgres=# CREATE EXTENSION pgx_stat_lwlock;
CREATE EXTENSION
postgres=# CREATE EXTENSION pgx_stat_latch;
CREATE EXTENSION
postgres=# CREATE EXTENSION pgx_stat_walwriter;
CREATE EXTENSION
postgres=# \q
```

Appendix C Using WebAdmin After an Upgrade

The following steps are required if performing an upgrade installation in an environment in which WebAdmin is used.

Check the service name

When starting WebAdmin for the first time after an upgrade, the service name of the instance created in WebAdmin is changed to the post-upgrade version. Therefore, if you have a job that is using the service name, you need to change the service name to the post-upgrade version.

Example)

If the instance name "inst1" with the user name "fepuser" was created in the Enterprise Postgres V9.2 Standard Edition, and the product is upgraded to 9.4:

[Service name of the instance created in WebAdmin]

Before upgrade: fepuser_0920_SE_64_fepuser_inst1

After upgrade: fsep_94_SE_64_fepuser_inst1

Change the settings information

1. Change the parameter file "postgresql.conf"

Change the following parameters in the "postgresql.conf" parameter file.

Change the "version" portion in the same manner as for the service name.

- log_directory
- core_directory
- event_source

Example)

If the instance name "inst1" with the user name "fepuser" was created in the Enterprise Postgres V9.2 Standard Edition, and the product is upgraded to 9.4:

[log_directory] [core_directory]

Before modification: C:\\Users\\fepuser\\AppData\\Local\\Fujitsu\\fep_0920_SE_64\\inst1\\log

After modification: C:\\Users\\fepuser\\AppData\\Local\\Fujitsu\\fsep_94_SE_64\\inst1\\log

[event_source]

Before modification: fep_0920_SE_64_fep_inst1

After modification: fsep_94_SE_64_fep_inst1

2. Change the services file instance name

Change the instance name of the services file.

Example)

If the instance name "inst1" with the user name "fepuser" was created using the port number "26500" in the Enterprise Postgres V9.2 Standard Edition version, and the product is upgraded to 9.4:

Directory for storing the services file: %SystemRoot%\system32\drivers\etc\services

Before modification: fep_0920_SE_64_fep_inst1 26500/tcp

After modification: fsep_94_SE_64_fep_inst1 27500/tcp

Creating directories for resource deployment

If the directory specified in "log_directory" and "core_directory" does not exist, create it.

Additionally, the directory access privileges must satisfy the following requirements:

- The directory owner is the operating system user account that will become the instance administrator
- There are permissions to write to the directory
- The directory is empty

Note

.....

If operations management middleware (for example, Systemwalker Centric Manager) is used to monitor instance ID information for the service name or server log, modify the operations management middleware settings together with these changes. If multiple methods were specified for the same table combination, the method specified last will be effective.

.....

Appendix D Recommended WebAdmin Environments

This appendix describes the recommended WebAdmin environment.

D.1 Recommended Browser Settings

- Use a display resolution to be 1024 x 768 or higher, and 256 colors or more.
- Select [View] >> [Text size] >> [Medium].
- Click [Tools] >> [Internet options] >> [General] >> [Fonts], and then:
 - Set [Webpage font] to [Times New Roman]
 - Set [Plain text font] to be [Courier New]

D.2 How to Set Up the Pop-up Blocker

If the Pop-up Blocker is enabled, use the procedure below to configure settings to allow pop-ups from the server where Enterprise Postgres is installed.

1. Click [Tools] >> [Internet options], and then select the [Privacy] tab.
If [Turn on Pop-up Blocker] is not selected, the Pop-up Blocker feature will not operate, and therefore steps below are not required.
2. Click [Settings].
3. In the [Pop-up Blocker Settings] window, enter in the [Address of website to allow] the URL of the server where Enterprise Postgres is installed, and then click [Add].
4. Click [Close].
5. In the [Internet Options] window, click [OK].

Appendix E Setting Up and Removing WebAdmin

This appendix describes how to set up and remove WebAdmin.

E.1 Setting Up WebAdmin

This section explains how to set up WebAdmin.

E.1.1 Setting Up WebAdmin

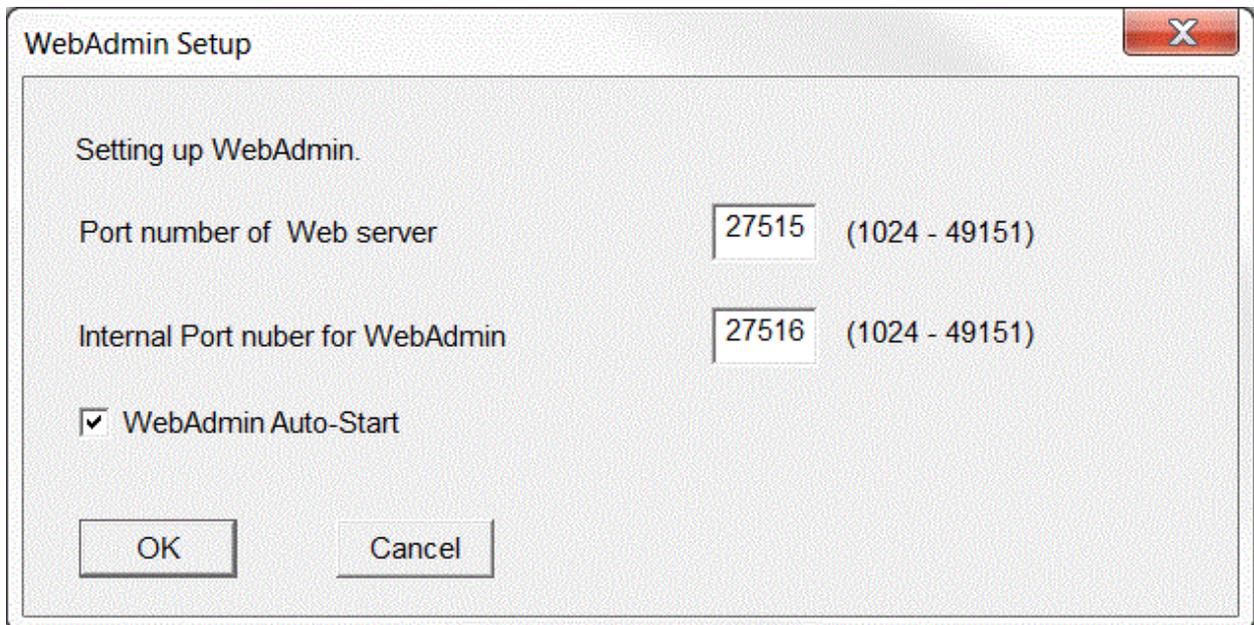
Follow the procedure below to set up WebAdmin.

1. Log in

Log in as a user that belongs to the Administrators group.

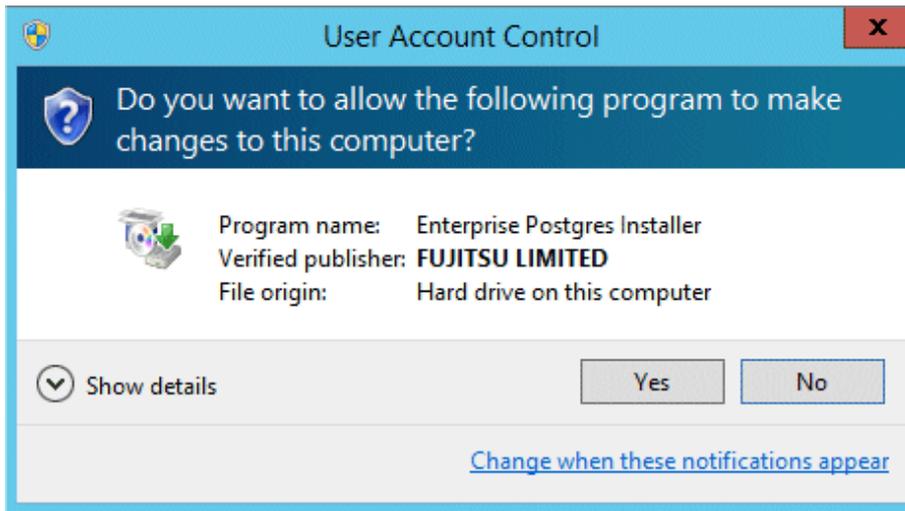
2. Display the setup window

In Windows, click [All Programs] or [All apps], then [Product name], and then [WebAdmin Setup].



Note

If the same [User Account Control] dialog box as that shown below is displayed, click [Yes] to continue processing.



3. Specify the port number

Specify the following port numbers to be used in WebAdmin.

Refer to the services file. Only change to a different port number if the same port number is being used by another service.

Make a note of the port number for the Web server, because it will be required for activating the WebAdmin window.

Item	Value (recommended value)
Web server port number	27515
WebAdmin internal port number	27516
WebAdmin automatic start	Checked

Web server port number

Specify an unused port number in the following range for the port number used for communication between the Web browser and Web server:

- Windows Server(R) 2008, Windows Server(R) 2008 R2, Windows Server(R) 2012 or Windows Server(R) 2012 R2: 1024 to 49151

The Web server port number is registered as a port number of the following service name in the services file:

64-bit product:

fsep_94_64_WebAdmin_Port1

32-bit product:

fsep_94_32_WebAdmin_Port1

WebAdmin internal port number

Specify an unused port number in the following range for the port number used for communication between the Web server and WebAdmin runtime environment:

- Windows Server(R) 2008, Windows Server(R) 2008 R2, Windows Server(R) 2012 or Windows Server(R) 2012 R2: 1024 to 49151

The WebAdmin internal port number is registered as a port number of the following service name in the services file:

64-bit product:

fsep_94_64_WebAdmin_Port2

32-bit product:

fsep_94_32_WebAdmin_Port2

WebAdmin automatic start

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



- Make a note of the Web server port number for use in the Windows firewall settings.
- Unused port numbers
In the operating system and other products, regardless of the information in the service file, unused port numbers may be automatically numbered and then used, or port numbers specified in environment files within products may also be used. Check the port numbers used by the OS and other products, and ensure that these are not duplicated.
- Access restrictions
Prevent unauthorized access and maintain security by using a firewall product, or the packet filtering feature of a router device, to restrict access to the server IP address and the various specified port numbers.
- Port access permissions
If a port is blocked (access permissions have not been granted) by a firewall, enable use of the port by granting access. Refer to the vendor document for information on how to grant port access permissions.
Consider the security risks carefully when opening ports.

4. Prepare for setup

Click [OK] in the setup window, and after completing the WebAdmin setup, refer to "[4.2 Preparations for Setup](#)" and perform the required preparations for setting up Enterprise Postgres if using WebAdmin for operation.

E.1.2 Activating the Web Server Feature of WebAdmin

Follow the procedure below to activate the Web server feature of WebAdmin.

1. Display the Services window

In Windows, select [Administrative Tools], and then click [Services].

2. Start the service

Select the displayed name "FUJITSU Enterprise Postgres WebAdmin *version*", and then click [Start Service].

You can also start the service by specifying the service name of the Web server feature of WebAdmin in the net start command or sc start command.

E.1.3 Stopping the Web Server Feature of WebAdmin

Follow the procedure below to stop the Web server feature of WebAdmin.

1. Display the Services window

In Windows, select [Administrative Tools], and then click [Services].

2. Stop the service

Select the displayed name "FUJITSU Enterprise Postgres WebAdmin *version*", and then click [Stop Service].

You can also stop the service by specifying the service name of the Web server feature of WebAdmin in the net stop command or sc stop command.

E.2 Removing WebAdmin

This section explains how to remove WebAdmin.

This removal procedure stops WebAdmin and ensures that it no longer starts automatically when the machine is restarted.

To remove the setup, execute the command shown below.

Example

When Enterprise Postgres is installed in "C:\Program Files\Fujitsu\fsepserver64":

```
> C:  
> cd C:\Program Files\Fujitsu\fsepserver64\gui\sbin  
> WebAdminSetup --delete
```

Note

- The removal of the WebAdmin setup must be performed by a user with administrator privileges (a user ID that belongs to the Administrators group).
- In Windows Server(R) 2008, Windows Server(R) 2008 R2, Windows Server(R) 2012 or , Windows Server(R) 2012 R2 commands that require administrator privileges must be executed from the [Administrator: Command Prompt] window. Right-click [Command Prompt], and then select [Run as administrator] from the context menu to display the [Administrator: Command Prompt] window.

Appendix F Configuring Parameters

WebAdmin operates and manages databases according to the contents of the following configuration files:

- [postgresql.conf](#)

Contains various items of information that define the operating environment of Enterprise Postgres.

- [pg_hba.conf](#)

Contains various items of information related to client authentication.

These configuration files are deployed to a data storage destination. Data is written to them when the instance is created by WebAdmin and when settings are changed, and data is read from them when the instance is started and when information from the [Setting] menu is displayed.

Direct editing of each configuration file is possible with a text editor.



See

Refer to "Server Configuration" and "Client Authentication" in "Server Administration" in the PostgreSQL Documentation for information on the parameters.

postgresql.conf

Parameters that can be changed in WebAdmin

The postgresql.conf parameters that can be changed in WebAdmin are shown below:

Tab	WebAdmin item	Parameter
Character Code	Character Set	client_encoding
	Message Locale	lc_messages (*1)
Communication	Port Number	port
	Max Connection	max_connections
SQL Options	Interpretation of Expression = NULL	transform_null_equals
	Date Output Format	DateStyle (*2)
	Interval Oputput Format	IntervalStyle
	The number of digits for floating values	extra_float_digits
	Transaction Isolation Levels	default_transaction_isolation
	Currency Format	lc_monetary
	Date and Time Format	lc_time
Memory	Sort Memory(Unit:KB)	work_mem
	Share Buffer(Unit:KB)	shared_buffers

*1: In the current version, you can only specify "C".

*2: If you specify "Postgres" as the output format, dates will be output in the "12-17-1997" format, not the "Wed Dec 17 1997" format used in the PostgreSQL Documentation.

Parameters set by WebAdmin

Parameters set by WebAdmin during instance startup are shown below (they will be ignored even if specified in postgresql.conf):

Parameter	Value
event_source	'fsep_version_userName_instanceName'
listen_addresses	*
log_destination	'stderr,eventlog'
logging_collector	on
log_line_prefix	'%e: %t [%p]: [%l-1] user = %u,db = %d,remote = %r app = %a '
log_directory	'userProfileFolder\localSettingsFolder\Fujitsu\fsep_version\instanceName\log' Example: If using Windows Server(R) 2008 R2, userProfileFolder\localSettingsFolder will be "C:\Users\userName\AppData\Local".
log_filename (*1)(*2)	'logfile-%a.log'
log_truncate_on_rotation	on
log_rotation_age	1d

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

*2) If the date changes while the instance is stopped, old logs are not deleted and continue to exist. Manually delete old logs that are no longer required to release disk space.

Parameters automatically set by WebAdmin according to the amount of memory

The postgresql.conf parameters automatically set according to the amount of installed memory, during the creation of instances by WebAdmin, are shown below:

Parameter	Value
shared_buffers	30% of the machine's installed memory
work_mem	30% of the machine's installed memory / max_connections / 2
effective_cache_size	75% of the machine's installed memory
maintenance_work_mem	10% of the machine's installed memory / (1 + autovacuum_max_workers)

When determining the values to be configured in the above parameters, you must take into account any anticipated increases in access volume or effects on performance during business operations, such as the number of applications and commands that will access the instance, and the content of processes. Also, note that in addition to Enterprise Postgres, other software may be running on the actual database server. You will need to determine the degree of priority for the database and other software, as well as the memory allocation size.

WebAdmin automatically configures complex parameter settings such as those mentioned above, based on the size of the internal memory of the machine. This enables maximum leverage of the machine memory to facilitate resistance against fluctuations during business operations.

Accordingly, the effects of the above-mentioned factors must be estimated and taken into account when determining and configuring parameter values, so that memory resources can be effectively allocated among other software or instances, and so that adverse effects can be mutually avoided. Refer to "Memory" in "Resource Consumption", and "Planner Cost Constants" in "Query Planning", under "Server Administration" in the PostgreSQL Documentation for information on parameter values and required considerations.

Parameter values can be modified using the WebAdmin [Setting] menu, or edited directly using a text editor.

If adding an instance, determine the parameter values, including for existing instances, and make changes accordingly.

Note

- Do not directly edit the following postgresql.conf parameters with a text editor, otherwise WebAdmin may not work properly if you make a mistake):
 - port
 - archive_mode
 - archive_command
 - wal_level
 - wal_sync_method
 - log_line_prefix
 - log_destination
 - logging_collector
 - log_directory
 - log_file_mode
 - log_filename
 - log_truncate_on_rotation
 - log_rotation_age
 - event_source
 - backup_destination
 - You must take care with the following parameter:
 - superuser_reserved_connections

Set it to a number that includes the 3 connections required in WebAdmin (the default is 3).
-

pg_hba.conf

Refer to "Client Authentication" in "Server Administration" in the PostgreSQL Documentation for information on content that can be configured in pg_hba.conf.

Note

- Use the following client authentication settings to allow the instance administrator to connect to the database using WebAdmin:
 - The connection type: "host"
 - The IP address is a loopback address ("127.0.0.1/32")
 - If you specify an item or value that cannot be set by WebAdmin when editing the pg_hba.conf file with a text editor, it will not be possible to reference that line from WebAdmin.
-

Appendix G Uninstall (middleware)

G.1 Features that are Installed

[Windows]

Feature	Package name	Component name	Remarks	Selectively installed?
Common tool	FJSVcir	CIRuntime Application	Controls the installation and uninstallation of Fujitsu middleware products. Manages the installation information of Fujitsu middleware products and includes a management function for viewing installation information and launching the uninstaller of each product. ("Uninstall (middleware)")	No

[Linux]

Feature	Package name	Component name	Remarks	Selectively installed?
Common tool	FJSVcir	CIRuntime Application	Controls the installation and uninstallation of Fujitsu middleware products. Manages the installation information of Fujitsu middleware products and includes a management function for viewing installation information and launching the uninstaller of each product. ("Uninstall (middleware)")	No

G.2 Notes

G.2.1 Notes on the Uninstall (middleware) Tool

"Uninstall (middleware)" is a tool used by Fujitsu middleware products. It performs information management for middleware products, and also, launching the uninstaller for removing these products from the installed systems.



- To uninstall Enterprise Postgres, uninstall from "Uninstall (middleware)".

This tool manages information about other Fujitsu middleware products even if Enterprise Postgres. For this reason, do not uninstall it unless absolutely necessary. If you have accidentally uninstalled this tool, re-install it as described below.

[Windows]

1. Log on to the machine to be installed using a user name belonging to the Administrators group, or switch to an account with administrator privileges.
2. Insert the server program DVD in the drive device.
3. Execute the installation command.

```
z:\CIR\Windows\cirinst.exe
```

In the example above, for z, specify the drive name of the drive where the DVD has been inserted.

[Linux]

1. Log in as a superuser to the installation target machine, or switch to a user with administrative privileges.
2. Insert the product media in the drive.
3. Execute the installation command.

```
# /media/dvd/CIR/Linux/cirinst.sh
```

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

- If you uninstall this tool as described below.

1. Check that there are no Fujitsu middleware products installed in the system.
You can check this by using "Uninstall (middleware)", which is launched by the command below.

[Windows]

Select [All Programs] > [Fujitsu] > [Uninstall (middleware)]

[Linux]

```
/opt/FJSVcir/cir/bin/cimanager.sh -c
```

-c : CUI

2. If no Fujitsu middleware products have been installed, then issue the command below.

[Windows]

```
%SystemDrive%\FujitsuF4CR\bin\cirremove.exe
```

[Linux]

```
# /opt/FJSVcir/bin/cirremove.sh
```

3. At the uninstallation confirmation prompt, enter "y".

```
This software is a common tool of Fujitsu products.Are you sure you want to remove it?[y/n]:
```

Uninstallation will start.

Appendix H Uninstall (middleware) Messages

H.1 Messages output by FJSVcir

FSP_FJSVCIR_CIRINST: ERROR: 101: CIRINST101: Administrator privilege is required.

Description

The user does not have administrator privileges.

Action method

Login with administrator privileges and run the command again.

FSP_FJSVCIR_CIRINST: ERROR: 102: CIRINST102: Failed to create %s

Parameters

%s: Directory name

Description

Failed to create a directory during the installation.

Action method

Creating the directory might have failed for some reason. Run the installer again, or run cirinst command again. If the problem persists, check the system log.

FSP_FJSVCIR_CIRINST: ERROR: 103: CIRINST103: It is required at least %dMB

Parameters

%d: Disk spaces

Description

Insufficient disk space.

Action method

Confirm that the system has sufficient disk spaces.

[Windows] FSP_FJSVCIR_CIRINST: ERROR: 104: CIRINST104: Failed to copy CIR files.

[Linux] FSP_FJSVCIR_CIRINST: ERROR: 104: CIRINST104: Failed to copy CIR installer files.

Description

Failed to copy files required for the installation.

Action method

Copying the files failed for some reason. Run the installer again, or run cirinst command again. If the problem persists, check the system log.

FSP_FJSVCIR_CIRINST: ERROR: 105: CIRINST105: Failed to copy JRE.

Description

Failed to copy JRE required for installation.

Action method

Copying the files failed for some reason. Run the installer again, or run cirinst command again. If the problem persists, check the system log.

FSP_FJSVCIR_CIRINST: ERROR: 106: CIRINST106: Failed to copy CIR**Description**

Failed to copy files during the installation.

Action method

Copying the files failed for some reason. Run the installer again, or run cirinst command again. If the problem persists, check the system log.

FSP_FJSVCIR_CIRINST: ERROR: 107: CIRINST107: invalid option.**Description**

An invalid option was specified.

Action method

Specify a valid option. [-c]: CUI

FSP_FJSVCIR_CIRINST: ERROR: 108: CIRINST108: JRE which should be deployed is not bundled.**Description**

An appropriate JRE for this system is not bundled.

Action method

Confirm that the product supports the platform.

FSP_FJSVCIR_CIRINST: ERROR: 109: CIRINST109: Unsupported OS.**Description**

Unsupported operating system.

Action method

Confirm that the product supports the operating system.

FSP_FJSVCIR_CIRINST: ERROR: 185: CIRINST185: lsb_release was not found. Please install the package below.**- redhat-lsb****Description**

lsb_release command was not found.

Action method

Add redhat-lsb package and then execute install again.

FSP_FJSVCIR_CIRINST: ERROR: 186: CIRINST186: lsb_release was not found.**Description**

lsb_release command was not found.

Action method

Add the relevant package and then execute install again.

FSP_FJSVCIR_CIRINST: ERROR: 201: CIRINST201: Administrator privilege is required.**Description**

The user does not have administrator privileges.

Action method

Login with administrator privileges and run the command again.

FSP_FJSVCIR_CIRINST: ERROR: 203: CIRINST203: Failed to delete files.

Description

Failed to delete FJSVcir files.

Action method

Deleting the files failed for some reason. Run cirremove command again. If the problem persists, check the system log.

FSP_FJSVCIR_CIRINST: ERROR: 205: CIRINST205: invalid option.

Description

Invalid option is specified.

Action method

No option is supported. Retry without any options.

Specified installation parameters file was not found.

Description

Specified installation parameters file was not found. (silent mode)

Action method

Specify the installation parameters file.

It failed to install. See log for details

Description

An error occurred during install. (silent mode)

Action method

Collect the files stored in the following directories.

[UNIX]

/var/opt/FJSVcir/cir/logs/cirlog0.log.0

[Windows]

%ProgramData%\Fujitsu\FujitsuF4CR\cir\logs\cirlog0.log.0

Parameters are invalid. Please specify valid parameters.

Description

Invalid option is specified. (silent mode)

Action method

Specify a valid option.

Internal data is broken.

Description

Internal data used by this product is corrupt.

Action method

Collect the files stored in the following directories and contact Fujitsu technical support:

[UNIX]

/var/opt/FJ SVCIR/cir/

/etc/opt/FJ SVCIR/cir/CIR.properties

[Windows]

%ProgramData%\Fujitsu\FujitsuF4CR\cir\

H.2 Messages output by Uninstall (middleware)

It has already been running in another process.

Description

An attempt was made to start Uninstall (middleware) more than once. Alternatively, an attempt was made to start Uninstall (middleware) while the installer was starting.

Action method

End any Uninstall (middleware) that have already started. Alternatively, start Uninstall (middleware) after the installer has ended.

Specified option is invalid.

Description

An invalid option was specified.

Action method

Specify valid options and run the command again.

Please enter a valid option.

Description

An invalid character was entered. Valid characters are y, n, or item numbers.

Action method

Enter a valid character.

Failed to uninstall of the following software packages:

Description

Failed to uninstall the software listed in this message.

Action method

Confirm the information shown after this message and take the appropriate action.

Unable to uninstall the software as it is referred from more than one software.

Description

Failed to uninstall the software because it is used by other software.

Action method

No specific action is required.

Unable to uninstall the software as the software is mandatory.

Description

Failed to uninstall the software because it is required by another program.

Action method

No specific action is required.

Internal data is broken.

Description

Internal data used by this product is corrupt.

Action method

Collect the files stored in the following directories and contact Fujitsu technical support:

[UNIX]

/var/opt/FJSVCIR/cir/

/etc/opt/FJSVCIR/cir/CIR.properties

[Windows]

%ProgramData%\Fujitsu\FujitsuF4CR\cir\

Unable to uninstall the software you were about to uninstall as the existence of the software is a precondition for the operation of another software.

Description

Failed to uninstall the software because it is required by another program.

Action method

No specific action is required.

The program terminated abnormally.

Description

The program has terminated abnormally.

Action method

Collect the files stored in the following directories and contact Fujitsu technical support:

[UNIX]

/var/opt/FJSVCIR/cir/

/etc/opt/FJSVCIR/cir/CIR.properties

[Windows]

%ProgramData%\Fujitsu\FujitsuF4CR\cir\

An unexpected error has occurred during uninstall.

Description

An error occurred during uninstall.

Action method

Collect the files stored in the following directories and contact Fujitsu technical support:

[UNIX]

/var/opt/FJSVCIR/cir/

/etc/opt/FJSVCIR/cir/CIR.properties

[Windows]

%ProgramData%\Fujitsu\FujitsuF4CR\cir\

It failed to uninstall. See log for details.

Description

An error occurred during uninstall. (silent mode)

Action method

Collect the files stored in the following directories.

[UNIX]

/var/opt/FJSVcir/cir/logs/cirlog0.log.0

[Windows]

%ProgramData%\Fujitsu\FujitsuF4CR\cir\logs\cirlog0.log.0

Failed to initialize the temp directory.

Description

Can not start Uninstall (middleware) because failed to initialize the temp directory.

Action method

Run Uninstall (middleware) again. If the problem persists, check whether other processes have accessed the files in the following directories.

[UNIX]

/var/opt/FJSVcir/cir/temp/meta_db

[Windows]

%ProgramData%\Fujitsu\FujitsuF4CR\cir\temp\meta_db

[Notice]Need to restart for uninstall completion.

Description

Uninstallation was completed. (silent mode)

Action method

Restart the system.

Appendix I Estimating Database Disk Space Requirements

This appendix describes how to estimate database disk space requirements.

I.1 Estimating Table Size Requirements

The following tables provide the formulas for estimating table size requirements.

Table I.1 Estimation formula when the record length is 2032 bytes or less

Item	Estimation formula (bytes)
(1) Record length	<p>$27(*1) + \text{NULL map} + \text{OID} + \text{column data}$</p> <p>NULL map: Number of columns / 8 (*2) OID: 4 for a table "WITH OID". The default is 0. Column data: Sum of column lengths</p> <p>*1: Record head section *2: Round the result up to the next integer.</p> <ul style="list-style-type: none"> - Because the column data is placed in boundaries of 8 bytes, you need to make an adjustment so that the sum of the record head section, NULL map and OID is a multiple of 8. For example, if the calculated length is $27 + 1 / 8$ (rounded up) + 0 = 28 bytes, add 4 to make the length 32 bytes. - Because the data of each column is placed in boundaries of the defined data type, take the boundary of each data type into account for the length of the column data. For example, the length of the column data in the table below will not be the sum of the data types, which is 37 bytes, but will instead be 64 bytes following boundary adjustment. Definition: create table tb1(c1 char(1), c2 long, c3 int, c4 box) Estimation: CHAR type 1 byte + boundary adjustment of 7 bytes for LONG type 8 bytes + LONG type 8 bytes + INT type 4 bytes + boundary adjustment of 12 bytes for BOX type 32 bytes + BOX type 32 bytes = 64 bytes - Because each record is placed in boundaries of 8 bytes, you need to make an adjustment so that the length of the column data is a multiple of 8. - If the calculated record length exceeds 2,032 bytes, the variable length data in the record might be compressed automatically. If so, use the estimation formulas in "Table I.2 Estimation formula when the record length exceeds 2032 bytes" to estimate the table size.
(2) Number of records per page	<p>$8168 (*1) / ((1) \text{ record length} + 4 (*2))$</p> <p>*1: Page length (8192) - page head (24) *2: Pointer length (4)</p> <ul style="list-style-type: none"> - The result will be rounded down to the next integer.
(3) Number of pages required for storing records	<p>Total number of records / (2) number of records per page</p> <ul style="list-style-type: none"> - The result will be rounded up to the next integer.
(4) Amount of space	<p>(3) Number of pages required for storing records x page length x safety factor (*1)</p> <p>*1: Specify 2.0 or higher.</p> <ul style="list-style-type: none"> - This is the safety factor assumed if vacuuming is performed for garbage collection in tables and indexes.

Table I.2 Estimation formula when the record length exceeds 2032 bytes

Item	Estimation formula (bytes)
(4) Amount of space	Total number of records x (1) record length x safety factor (*1) *1: Specify 2.0 or higher. - This is the safety factor assumed if vacuuming is performed for garbage collection in tables and indexes.

I.2 Estimating Index Size Requirements

This section provides the formulas for estimating index size requirements.

Enterprise Postgres provides five index types: B-tree, Hash, GiST, GIN, and SP-GiST (the same as PostgreSQL). If you do not specify the index type in the CREATE INDEX statement, a B-tree index is generated.

The following describes how to estimate a B-tree index.

A B-tree index is saved as a fixed-size page of 8 KB. The page types are meta, root, leaf, internal, deleted, and empty. Since leaf pages usually account for the highest proportion of space required, you need to calculate the requirements for these only.

Table I.3 Estimation formula when the key data length is 512 bytes or less

Item	Estimation formula (bytes)
(1) Entry length	8 (*1) + key data length (*2) *1: Entry head *2: The key data length depends on its data type (refer to "I.3 Sizes of Data Types" for details). Because each entry is placed in boundaries of 8 bytes, you need to make an adjustment so that the length of the key data is a multiple of 8. For example, if the calculated length is 28 bytes, add 4 to make the length 32 bytes. - If the key data length exceeds 512 bytes, key data may be automatically compressed. In this case, use the estimation formula given in "Table I.4 Estimation formula when the key data length exceeds 512 bytes" to estimate the key data length.
(2) Page size requirement	8192 (*1) *1: Page length (8192) - page header (24) - special data (16) = 8152
(3) Number of entries per page	(2) Page size requirement / ((1) entry length + 4 (*1)) *1: Pointer length - Result of (3) number of entries per page will be rounded down to the nearest integer.
(4) Number of pages required for storing indexes	Total number of records / (3) number of entries per page - Result of (4) number of pages required for storing indexes will be rounded up to the nearest integer.
(5) Space requirement	(4) Number of pages required for storing indexes x 8192 (*1) / usage rate (*2) *1: Page length *2: Specify 0.7 or lower.

Table I.4 Estimation formula when the key data length exceeds 512 bytes

Item	Estimation formula (bytes)
(5) Space requirement	Total number of records x key data length x compression ratio (*1) / usage rate (*2) *1: The compression ratio depends on the data value, so specify 1. *2: Specify 0.7 or lower as the usage rate.

I.3 Sizes of Data Types

This section lists the sizes of the data types.

I.3.1 Sizes of Fixed-Length Data Types

The following table lists the sizes of fixed-length data types.

Data type	Size (bytes)
SMALLINT (INT2)	2
INTEGER (INT4)	4
BIGINT (INT8)	8
REAL	4
DOUBLE PRECISION	8
SERIAL (SERIAL4)	4
BIGSERIAL (SERIAL8)	8
MONEY	8
FLOAT	8
FLOAT (1-24)	4
FLOAT (25-53)	8
TIMESTAMP WITHOUT TIME ZONE	8
TIMESTAMP WITH TIME ZONE	8
DATE	4
TIME WITHOUT TIME ZONE	8
TIME WITH TIME ZONE	12
INTERVAL	12
BOOLEAN	1
CIDR	IPv4: 7 IPv6: 19
INET	IPv4: 7 IPv6: 19
MACADDR	6
POINT	16
LINE	32
LSEG	32
BOX	32
CIRCLE	24

I.3.2 Sizes of Variable-Length Data Types

The following table lists the sizes of variable-length data types.

Data type	Size (bytes)	Remarks
path	Length of size portion + 12 + 16 x number of vertices	1) When carrying out division, round to the next integer.

Data type	Size (bytes)	Remarks
polygon	Length of size portion + 36 + 16 x number of vertices	2) If the real data length is less than 127, then the length of the size portion is 1 byte, otherwise it is 4 bytes. 3) The number of bytes per character depends on the character set (refer to " I.3.4 Number of Bytes per Character " for details).
decimal	Length of size portion + 2 + (integer precision / 4 + decimal precision / 4) x 2	
numeric		
bytea	Length of size portion + real data length	
character varying(n), varchar(n)	Length of size portion + number of characters x number of bytes per character	
character(n), char(n)	Length of size portion + n x number of bytes per character	
text	Length of size portion + number of characters x number of bytes per character	

I.3.3 Sizes of Array Data Types

The following table lists the sizes of array data types.

Data type	Size (bytes)	Remarks
Array	Length of size portion + 12 + 8 x number of dimensions + data size of each item	If the real data length is less than 127, then the length of the size portion is 1 byte, otherwise it is 4 bytes. - Example of estimation when array data is "ARRAY[[1,2,3],[1,2,3]]" Number of dimensions: 2 INTEGER data size: 4 Total size = 1+12+8x2+6x4 = 53

I.3.4 Number of Bytes per Character

The following table lists the number of bytes per character.

The given values relate to the common character sets EUC-JP and UTF8.

Character type	Character set	Number of bytes per character
ASCII	EUC_JP	1
Halfwidth katakana	EUC_JP	2
JIS X 0208 kanji characters	EUC_JP	2
JIS X 0212 kanji characters	EUC_JP	3
ASCII	UTF8	1
Halfwidth katakana	UTF8	3
JIS X 0208 kanji characters	UTF8	3
JIS X 0212 kanji characters	UTF8	3

I.4 Estimating Transaction Log Space Requirements

This section provides the formula for estimating transaction log space requirements.

```
Transaction log space requirements = ( checkpoint_segments x 3 + 1 ) x 16 MB
```

However, if the update volume is extremely high (for example, due to a large data load and batch processing), disk writing at a checkpoint may not be able to keep up with the load, and a higher number of transaction logs than indicated here may temporarily be accumulated.

I.5 Estimating Archive Log Space Requirements

This section explains how to estimate archive log space requirements.

The archive log is an archive of the transaction logs from the time of a previous backup to the present, so it fluctuates depending on the backup period and the content of update transactions.

The longer the backup period and the more update transactions, the greater the space required for the archive log.

Therefore, measure the actual archive log space by using a test environment to simulate backup scheduling and database update in a real operating environment.

I.6 Estimating Backup Disk Space Requirements

This section provides the formula for estimating backup disk space requirements.

```
Backup disk space requirements = size of the database cluster + transaction log space requirements  
+ archive log space requirements
```

Appendix J Estimating Memory Requirements

This appendix explains how to estimate the memory.

J.1 Enterprise Postgres Memory Requirements

This section describes the formulas for estimating Enterprise Postgres memory requirements.

Use the following formula to obtain a rough estimate of memory required for Enterprise Postgres:

$$\text{enterprisePostgresRequiredMemory} = \text{sharedMemoryAmount} + \text{LocalMemoryAmount}$$

Shared memory amount

Refer to "Shared Memory and Semaphores" under "Server Administration" in the PostgreSQL Documentation for information on shared memory.

However, note that if instances have been created using WebAdmin, the parameters below will be configured automatically when the instances are created. Take this into account when calculating the shared memory size.

Parameter name	Set value
shared_buffers	30 percent of the internal memory of the machine.
max_connections	100
max_prepared_transactions	100

Local memory amount

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

Process stack area

$$\begin{aligned} \text{processStackArea} \\ = & \text{max_stack_depth} \times (\text{max_connections} + \text{autovacuum_max_workers} + 9) \end{aligned}$$

This formula evaluates to the maximum value.

Actually it is used according to the growth of the stack.

In the formula above, 9 is the number of processes that perform roles specific to servers.

Memory used in database sessions that use temporary tables

$$\begin{aligned} \text{memoryUsedInDbSessionsThatUseTempTables} \\ = & \text{temp_buffers} \times \text{max_connections} \end{aligned}$$

This formula evaluates to the maximum value.

Memory is gradually used as temporary buffers are used, and is released when the session ends.

Memory used in database sessions that perform sort and hash table operations

$$\begin{aligned} \text{memoryUsedInDbSessionsThatPerformSortAndHashTableOperations} \\ = & \text{work_mem} \times \text{max_connections} \end{aligned}$$

This formula evaluates to the maximum value.

Memory is gradually used as operations such as sort are performed, and is released when the query ends.

Memory used in maintenance operations

$$\begin{aligned} & \text{memoryUsedInMaintenanceOperations} \\ & = \text{maintenance_work_mem} \times (\text{numOfSessionsPerformingMaintenance} + \text{autovacuum_max_workers}) \end{aligned}$$

Note that 'maintenance operations' are operations such as VACUUM, CREATE INDEX, and ALTER TABLE ADD FOREIGN KEY.

Base memory used in each process

$$\begin{aligned} & \text{baseMemoryUsedInEachProcess} \\ & = 3\text{MB} \times (\text{max_connections} + \text{autovacuum_max_workers} + 9) \end{aligned}$$

This formula evaluates to the memory used when server processes are running.

In the formula above, 9 is the number of processes that perform roles specific to servers.

Memory used preparing for data access

$$\begin{aligned} & \text{memoryUsedPreparingForDataAccess} \\ & = \text{variationAmount} \times (\text{max_connections} + \text{autovacuum_max_workers} + 4) \end{aligned}$$

where $\text{variationAmount} = \text{shared_buffers} / 8\text{KB} \times 4 \text{ bytes}$
(note that 8KB is the page length, and 4 bytes is the size of page management data)

This formula evaluates to the memory required to access the database cache in the shared memory.

In the formula above, among the processes that perform roles specific to servers, 4 is the number of processes that access the database.

J.2 Database Multiplexing Memory Requirements

This section describes the formula for estimating database multiplexing memory requirements.

Use the following formula to obtain a rough estimate of memory required for database multiplexing:

$$\begin{aligned} & \text{Memory usage of the database multiplexing feature} \\ & = \text{Peak memory usage of the Mirroring Controller processes} \\ & \quad + \text{Peak memory usage of the Mirroring Controller commands} \end{aligned}$$

Peak memory usage of the Mirroring Controller processes=150 MB

Peak memory usage of the Mirroring Controller commands=50 MB * Number of commands executed simultaneously

J.3 PL/extJava Memory Requirements

This section describes the formula for estimating PL/extJava memory requirements.

Use the following formula to obtain a rough estimate of memory required for PL/extJava:

$$\text{memoryUsedByPLExtJavaInMegabytes} = 600 + (\text{numOfServerInstances(JavaVMs)} \times 500)$$



Note

- Creating a container will create one server instance (Java VM).
- The number of server instances (Java VMs) may fluctuate based on the following actions:
 - Adding a container (one server instance (Java VM) will be added)
 - Adding a server instance (Java VM) to a container
 - Deleting a server instance (Java VM) from a container

Appendix K Quantitative Limits

This appendix lists the quantitative limits of Enterprise Postgres.

Table K.1 Length of identifier

Item	Limit
Database name	Up to 63 bytes (*1)(*2)(*3)
Schema name	Up to 63 bytes (*1)(*2)
Table name	Up to 63 bytes (*1)(*2)
View name	Up to 63 bytes (*1)(*2)
Index name	Up to 63 bytes (*1)(*2)
Table space name	Up to 63 bytes (*1)(*2)
Cursor name	Up to 63 bytes (*1)(*2)
Function name	Up to 63 bytes (*1)(*2)
Aggregate function name	Up to 63 bytes (*1)(*2)
Trigger name	Up to 63 bytes (*1)(*2)
Constraint name	Up to 63 bytes (*1)(*2)
Conversion name	Up to 63 bytes (*1)(*2)
Role name	Up to 63 bytes (*1)(*2)(*4)
Cast name	Up to 63 bytes (*1)(*2)
Collation sequence name	Up to 63 bytes (*1)(*2)
Encoding method conversion name	Up to 63 bytes (*1)(*2)
Domain name	Up to 63 bytes (*1)(*2)
Extension name	Up to 63 bytes (*1)(*2)
Operator name	Up to 63 bytes (*1)(*2)
Operator class name	Up to 63 bytes (*1)(*2)
Operator family name	Up to 63 bytes (*1)(*2)
Rewrite rule name	Up to 63 bytes (*1)(*2)
Sequence name	Up to 63 bytes (*1)(*2)
Text search settings name	Up to 63 bytes (*1)(*2)
Text search dictionary name	Up to 63 bytes (*1)(*2)
Text search parser name	Up to 63 bytes (*1)(*2)
Text search template name	Up to 63 bytes (*1)(*2)
Data type name	Up to 63 bytes (*1)(*2)
Enumerator type label	Up to 63 bytes (*1)(*2)

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

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

*3: Names of databases that use PL/extJava must be 28 bytes or less.

*4: Names of roles for connecting from the PL/extJava application server must be 8 bytes or less.

Table K.2 Database object

Item	Limit
Number of databases	Less than 4,294,967,296 (*1)
Number of schemas	Less than 4,294,967,296 (*1)
Number of tables	Less than 4,294,967,296 (*1)
Number of views	Less than 4,294,967,296 (*1)
Number of indexes	Less than 4,294,967,296 (*1)
Number of table spaces	Less than 4,294,967,296 (*1)
Number of functions	Less than 4,294,967,296 (*1)
Number of aggregate functions	Less than 4,294,967,296 (*1)
Number of triggers	Less than 4,294,967,296 (*1)
Number of constraints	Less than 4,294,967,296 (*1)
Number of conversion	Less than 4,294,967,296 (*1)
Number of roles	Less than 4,294,967,296 (*1)
Number of casts	Less than 4,294,967,296 (*1)
Number of collation sequences	Less than 4,294,967,296 (*1)
Number of encoding method conversions	Less than 4,294,967,296 (*1)
Number of domains	Less than 4,294,967,296 (*1)
Number of extensions	Less than 4,294,967,296 (*1)
Number of operators	Less than 4,294,967,296 (*1)
Number of operator classes	Less than 4,294,967,296 (*1)
Number of operator families	Less than 4,294,967,296 (*1)
Number of rewrite rules	Less than 4,294,967,296 (*1)
Number of sequences	Less than 4,294,967,296 (*1)
Number of text search settings	Less than 4,294,967,296 (*1)
Number of text search dictionaries	Less than 4,294,967,296 (*1)
Number of text search parsers	Less than 4,294,967,296 (*1)
Number of text search templates	Less than 4,294,967,296 (*1)
Number of data types	Less than 4,294,967,296 (*1)
Number of enumerator type labels	Less than 4,294,967,296 (*1)
Number of default access privileges defined in the ALTER DEFAULT PRIVILEGES statement	Less than 4,294,967,296 (*1)
Number of large objects	Less than 4,294,967,296 (*1)
Number of rows in tables defined by WITH OIDS	Less than 4,294,967,296 (*1)
Number of index access methods	Less than 4,294,967,296 (*1)

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

Table K.3 Schema element

Item	Limit
Number of columns that can be defined in one table	From 250 to 1600 (according to the data type)
Table row length	Up to 400 gigabytes

Item	Limit
Number of columns comprising a unique constraint	Up to 32 columns
Data length comprising a unique constraint	Less than 2,000 bytes (*1)(*2)
Table size	Up to 32 terabyte
Search condition character string length in a trigger definition statement	Up to 800 megabytes (*1)(*2)
Item size	Up to 1 gigabyte

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

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

Table K.4 Index

Item	Limit
Number of columns comprising a key (B-tree, GiST, GIN index)	Up to 32 columns
Key length	Less than 2,000 bytes (*1)

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

Table K.5 Data types and attributes that can be handled

Item		Limit	
Character	Data length	Data types and attributes that can be handled (*1)	
	Specification length (n)	Up to 10,485,760 characters (*1)	
Numeric	External decimal expression	Up to 131,072 digits before the decimal point, and up to 16,383 digits after the decimal point	
	Internal binary expression	2 bytes	From -32,768 to 32,767
		4 bytes	From -2,147,483,648 to 2,147,483,647
		8 bytes	From -9,223,372,036,854,775,808 to 9,223,372,036,854,775,807
	Internal decimal expression		Up to 13,1072 digits before the decimal point, and up to 16,383 digits after the decimal point
	Floating point expression	4 bytes	From -3.4E+38 to -7.1E-46, 0, or from 7.1E-46 to 3.4E+38
		8 bytes	From -1.7E+308 to -2.5E-324, 0, or from 2.5E-324 to 1.7E+308
bytea		Up to one gigabyte minus 53 bytes	
Large object		Up to two gigabytes	

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

Table K.6 Function definition

Item	Limit
Number of arguments that can be specified	Up to 100
Number of variable names that can be specified in the declarations section	No limit

Item	Limit
Number of SQL statements or control statements that can be specified in a function processing implementation	No limit

Table K.7 Data operation statement

Item	Limit
Maximum number of connections for one process in an application (remote access)	4,000 connections
Number of expressions that can be specified in a selection list	Up to 1,664
Number of tables that can be specified in a FROM clause	No limit
Number of unique expressions that can be specified in a selection list/DISTINCT clause/ORDER BY clause/GROUP BY clause within one SELECT statement	Up to 1,664
Number of expressions that can be specified in a GROUP BY clause	No limit
Number of expressions that can be specified in an ORDER BY clause	No limit
Number of SELECT statements that can be specified in a UNION clause/INTERSECT clause/EXCEPT clause	Up to 4,000 (*1)
Number of nestings in joined tables that can be specified in one view	Up to 4,000 (*1)
Number of functions or operator expressions that can be specified in one expression	Up to 4,000 (*1)
Number of expressions that can be specified in one row constructor	Up to 1,664
Number of expressions that can be specified in an UPDATE statement SET clause	Up to 1,664
Number of expressions that can be specified in one row of a VALUES list	Up to 1,664
Number of expressions that can be specified in a RETURNING clause	Up to 1,664
Total expression length that can be specified in the argument list of one function specification	Up to 800 megabytes (*2)
Number of cursors that can be processed simultaneously by one session	No limit
Character string length of one SQL statement	Up to 800 megabytes (*1) (*3)
Number of input parameter specifications that can be specified in one dynamic SQL statement	No limit
Number of tokens that can be specified in one SQL statement	Up to 10,000
Number of values that can be specified as a list in a WHERE clause IN syntax	No limit
Number of expressions that can be specified in a USING clause	No limit
Number of JOINS that can be specified in a joined table	Up to 4,000 (*1)
Number of expressions that can be specified in COALESCE	No limit
Number of WHEN clauses that can be specified for CASE in a simple format or a searched format	No limit
Data size per record that can be updated or inserted by one SQL statement	Up to one gigabyte minus 53 bytes
Number of objects that can share a lock simultaneously	Up to 256,000 (*1)

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

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

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

Table K.8 Data size

Item	Limit
Data size per record for input data files (COPY statement, psql command \copy meta command)	Up to 800 megabytes (*1)
Data size per record for output data files (COPY statement, psql command \copy meta command)	Up to 800 megabytes (*1)

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

Index

	[C]		[W]
Changing allocated memory.....	42	WebAdmin automatic start.....	73
Changing client authentication.....	37	Web server port number.....	72
Changing communication.....	39	When an Instance was Created with the initdb Command.....	47
Changing the character set.....	36	When an Instance was Created with WebAdmin.....	46
Changing the settings.....	35		
Changing the SQL option.....	41		
Check the disk space.....	9		
Check the installed product and determine the installation method	9		
Check the results.....	60		
Client Authentication Information settings.....	47		
Creating an Instance Administrator.....	20		
	[E]		
Excluded Software.....	5		
	[F]		
Firewall.....	26,47		
	[H]		
Hardware Environment.....	5		
How to Set Up the Pop-up Blocker	70		
	[I]		
Installation Procedure.....	1		
Installation Types.....	1		
	[N]		
New Installation.....	1		
	[O]		
Operating Environment.....	3		
Operating Method Types and Selection.....	19		
	[P]		
Port number to use when Tomcat is stopped.....	72		
Post-installation Tasks.....	18		
postgresql.conf.....	75		
Pre-installation Tasks.....	9		
Preparations for Setup.....	20		
	[R]		
Recommended Browser Settings.....	70		
Related Software.....	4		
Remove applied updates.....	10		
Required Operating System.....	3		
Required Patches.....	5		
	[S]		
Settings related to connection.....	46		
	[T]		
TCP/IP Protocol.....	5		
	[U]		
uninstaller.....	59		