

Fujitsu Enterprise Postgres 14 SP1 on IBM Power®

Installation and Setup Guide for Server

SP1

Linux



Preface

Purpose of this document

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

This document describes how to install and set up "FUJITSU Enterprise Postgres".

Intended readers

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

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

- PostgreSQL
- SQL
- 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 FUJITSU Enterprise Postgres

[Chapter 3 Installation](#)

Describes how to perform a new installation of FUJITSU Enterprise Postgres

[Chapter 4 Setup](#)

Describes the setup to be performed after installation

[Chapter 5 Uninstallation](#)

Describes how to uninstall FUJITSU Enterprise Postgres

[Appendix A Recommended WebAdmin Environments](#)

Describes the recommended WebAdmin environment.

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

Describes how to set up and remove WebAdmin

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

Describes characters that are not allowed in WebAdmin.

[Appendix D Configuring Parameters](#)

Describes FUJITSU Enterprise Postgres parameters.

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

Describes how to estimate database disk space requirements

[Appendix F Estimating Memory Requirements](#)

Describes the formulas for estimating memory requirements

[Appendix G Quantitative Limits](#)

Describes the quantity range

[Appendix H Configuring Kernel Parameters](#)

Describes the settings for kernel parameters

[Appendix I Determining the Preferred WebAdmin Configuration](#)

Describes the two different configurations in which WebAdmin can be used and how to select the most suitable configuration

[Appendix J System Configuration when using Pgpool-II](#)

Describes the system configuration when using Pgpool-II.

[Appendix K Supported contrib Modules and Extensions Provided by External Projects](#)

Lists the PostgreSQL contrib modules and the extensions provided by external projects supported by FUJITSU Enterprise Postgres.

[Appendix L Procedure when Modifying the JRE Installation](#)

Describes the procedures to follow when modifying the JRE installation.

Export restrictions

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

Issue date and version

Edition 2.0: September 2022 Edition 1.0: April 2022
--

Copyright

Copyright 2022 FUJITSU LIMITED

Contents

Chapter 1 Overview of Installation.....	1
1.1 Features that can be Installed.....	1
1.2 Installation Types.....	1
1.2.1 New Installation.....	1
1.2.2 Reinstallation.....	1
1.3 Uninstallation.....	1
Chapter 2 Operating Environment.....	2
2.1 Required Operating System.....	2
2.2 Related Software.....	5
2.3 Excluded Software.....	5
2.4 Required Patches.....	5
2.5 Hardware Environment.....	5
2.6 Disk Space Required for Installation.....	5
2.7 Supported System Environment.....	6
2.7.1 TCP/IP Protocol.....	6
2.7.2 File System.....	6
2.8 PostgreSQL Version Used for FUJITSU Enterprise Postgres.....	6
2.9 Notes on Using Streaming Replication.....	6
2.10 Key Management System Requirements.....	6
Chapter 3 Installation.....	8
3.1 Pre-installation Tasks.....	8
3.2 Run Installation	9
Chapter 4 Setup.....	12
4.1 Operating Method Types and Selection.....	12
4.2 Preparations for Setup.....	13
4.2.1 Creating an Instance Administrator.....	13
4.2.2 Preparing Directories for Resource Deployment.....	14
4.2.3 Estimating Resources.....	16
4.2.4 Configuring Corefile Names.....	17
4.3 Creating Instances.....	17
4.3.1 Using WebAdmin.....	17
4.3.1.1 Logging in to WebAdmin.....	18
4.3.1.2 Creating an Instance.....	19
4.3.1.3 Changing Instance Settings.....	21
4.3.1.4 Importing Instances.....	23
4.3.2 Using the initdb Command.....	24
4.3.2.1 Editing Kernel Parameters.....	24
4.3.2.2 Creating an Instance.....	25
4.4 Configuring Remote Connections.....	27
4.4.1 When an Instance was Created with WebAdmin.....	27
4.4.2 When an Instance was Created with the initdb Command.....	28
4.5 Other Settings.....	28
4.5.1 Error Log Settings.....	28
4.5.2 Configuring Automatic Start and Stop of an Instance.....	29
4.5.3 Settings when Using the Features Compatible with Oracle Databases.....	30
4.6 Setting Up and Removing OSS.....	31
4.6.1 oracle_fdw.....	31
4.6.1.1 Setting Up oracle_fdw.....	31
4.6.1.2 Removing oracle_fdw.....	32
4.6.2 pg_bigm.....	32
4.6.2.1 Setting Up pg_bigm.....	32
4.6.2.2 Removing pg_bigm.....	32
4.6.3 pg_hint_plan.....	33

4.6.3.1 Setting Up pg_hint_plan.....	33
4.6.3.2 Removing pg_hint_plan.....	33
4.6.4 pg_dbms_stats.....	34
4.6.4.1 Setting Up pg_dbms_stats.....	34
4.6.4.2 Removing pg_dbms_stats.....	34
4.6.5 pg_repack.....	35
4.6.5.1 Setting Up pg_repack.....	35
4.6.5.2 Removing pg_repack.....	35
4.6.6 pg_rman.....	35
4.6.6.1 Setting Up pg_rman.....	35
4.6.6.2 Removing pg_rman.....	36
4.6.7 pg_statsinfo.....	36
4.6.7.1 Setting Up pg_statsinfo.....	36
4.6.7.2 Removing pg_statsinfo.....	37
4.6.8 pgBadger.....	37
4.6.8.1 Setting Up pgBadger.....	37
4.6.8.2 Removing pgBadger.....	37
4.6.9 Pgpool-II.....	37
4.6.9.1 Setting Up Pgpool-II.....	37
4.6.9.2 Removing Pgpool-II.....	38
4.6.10 Build with PGXS.....	38
4.6.10.1 Using the Default Version of llvm.....	38
4.6.10.2 Using a Non-Default Version of llvm.....	38
4.6.10.3 Without llvm.....	39
4.7 Integration with Message-Monitoring Software.....	39
4.8 Deleting Instances.....	39
4.8.1 Using WebAdmin.....	40
4.8.2 Using Server Commands.....	40
Chapter 5 Uninstallation.....	41
5.1 Run Uninstallation.....	41
Appendix A Recommended WebAdmin Environments.....	43
A.1 Recommended Browser Settings	43
A.2 How to Set Up the Pop-up Blocker.....	43
Appendix B Setting Up and Removing WebAdmin.....	44
B.1 Setting Up WebAdmin.....	44
B.1.1 Setting Up WebAdmin.....	44
B.1.2 Starting the Web Server Feature of WebAdmin.....	45
B.1.3 Stopping the Web Server Feature of WebAdmin.....	45
B.2 Removing WebAdmin.....	46
B.3 Using an External Repository for WebAdmin.....	46
B.4 Using the WebAdmin Auto-Refresh Feature.....	48
Appendix C WebAdmin Disallow User Inputs Containing Hazardous Characters.....	49
Appendix D Configuring Parameters.....	50
Appendix E Estimating Database Disk Space Requirements.....	54
E.1 Estimating Table Size Requirements.....	54
E.2 Estimating Index Size Requirements.....	55
E.3 Sizes of Data Types.....	56
E.3.1 Sizes of Fixed-Length Data Types.....	56
E.3.2 Sizes of Variable-Length Data Types.....	57
E.3.3 Sizes of Array Data Types.....	57
E.3.4 Number of Bytes per Character.....	58
E.4 Estimating Transaction Log Space Requirements.....	58
E.5 Estimating Archive Log Space Requirements.....	58

E.6 Estimating Backup Disk Space Requirements.....	58
E.7 Estimating VCI Disk Space Requirements.....	58
Appendix F Estimating Memory Requirements.....	60
F.1 FUJITSU Enterprise Postgres Memory Requirements.....	60
F.2 Database Multiplexing Memory Requirements.....	61
F.3 VCI Memory Requirements.....	61
F.4 High-Speed Data Load Memory Requirements.....	63
F.5 Global Meta Cache Memory Requirements.....	63
Appendix G Quantitative Limits.....	64
Appendix H Configuring Kernel Parameters.....	69
Appendix I Determining the Preferred WebAdmin Configuration.....	70
I.1 WebAdmin Configurations.....	70
I.1.1 Single-Server Configuration.....	70
I.1.2 Multiserver Configuration.....	70
I.2 Installing WebAdmin in a Single-Server Configuration.....	71
I.3 Installing WebAdmin in a Multiserver Configuration.....	72
Appendix J System Configuration when using Pgpool-II.....	73
J.1 Pgpool-II Configuration.....	73
J.1.1 Single-Machine Configuration.....	73
J.1.2 Two-Machine Configuration.....	74
J.1.3 Three-Machine Configuration.....	74
J.2 Installing Pgpool-II.....	74
J.3 Pgpool-II Setup.....	75
J.3.1 Setting Environment Variables.....	75
J.3.2 Configuration file.....	75
J.3.2.1 Configuring pgpool.conf.....	75
J.3.2.2 Using Configuration Files.....	75
Appendix K Supported contrib Modules and Extensions Provided by External Projects.....	76
Appendix L Procedure when Modifying the JRE Installation.....	77
L.1 When Using WebAdmin.....	77
L.2 When Performing Database Multiplexing.....	77
Index.....	79

Chapter 1 Overview of Installation

This chapter provides an overview of FUJITSU Enterprise Postgres installation.

1.1 Features that can be Installed

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

The basic features of FUJITSU Enterprise Postgres (server feature, client feature) can be installed.

1.2 Installation Types

The following installation types are available for FUJITSU Enterprise Postgres:

- New installation
- Reinstallation

1.2.1 New Installation

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

1.2.2 Reinstallation

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

1.3 Uninstallation

Uninstallation removes the system files of the installed FUJITSU Enterprise Postgres.

Chapter 2 Operating Environment

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



See

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

2.1 Required Operating System

One of the operating systems shown below is required in order to use FUJITSU Enterprise Postgres. Check and use minor version, which is certified and currently supported by Red Hat or SUSE for IBM Power LE (POWER9 and POWER10).

- RHEL8.4 or later minor version
- SLES 15 SP3



Note

The sepgsql module, which is a PostgreSQL extension, can be used in RHEL8.



Information

- The following packages are required for operations on RHEL8.

Package name	Remarks
alsa-lib	-
audit-libs	-
cyrus-sasl-lib	-
pcp-system-tools	Required when using parallel scan.
gdb	Required when using pgx_fjqssinf command.
glibc	-
iputils	Required for Mirroring Controller.
libnsl2	-
libc	Provides collation support. Install 60.x.
libgcc	-
libmemcached	Required when using Pgpool-II.
libstdc++	-
libtool-ltdl	Required when using ODBC drivers.
llvm	Versions 12.0.x or 11.0.x of llvm is required to run SQL with runtime compilation (just-in-time compilation). Install the package that contains libLLVM-12.so, or libLLVM-11.so. For example, version 12.0.x of "llvm-libs" published with Application Streams includes libLLVM-12.so.

Package name	Remarks
	<p>By default, version 12.0.x is used.</p> <p>If you use a version other than 12.0.x, specify the version you want to use in the <code>jit_provider</code> parameter in <code>postgresql.conf</code>.</p> <p>For example, use <code>llvmjit-vsn11</code> when using version 11.0.x. FUJITSU Enterprise Postgres uses runtime compilation by default. If you do not want to use runtime compilation, turn off the <code>jit</code> parameter in <code>postgresql.conf</code>. You do not need to install <code>llvm</code> if you turn off the <code>jit</code> parameter.</p> <p>Failure to install <code>llvm</code> without turning off the <code>jit</code> parameter may result in an error when executing SQL. For more information about runtime compilation, see "Just-in-Time Compilation (JIT)" in the PostgreSQL Documentation.</p> <p>Note:</p> <p>If JIT compilation using LLVM11 is performed on RHEL8, a verbose server message is logged one or more times on first query execution, for each associated database connection.</p> <p><Message></p> <p>ORC error: No callback manager available for powerpc64le-unknown-linux-gnu</p> <p>This error message results from a known deficiency in one of the features of the LLVM library for POWER, but this feature is not used by the server and the error does not otherwise affect LLVM operation.</p> <p>To prevent this message from being logged, JIT compilation may be disabled, by setting the value "off" to the configuration variable "jit".</p>
lz4-libs	-
ncurses-libs	-
net-tools	-
nss-softoken-freebl	-
pam	Required when using PAM authentication.
perl-libs	Required when using PL/Perl. Install 5.26.
protobuf-c	Required if using the Transparent Data Encryption feature when using a key management system as a keystore. Install 1.3.0.
python3	Required when using PL/Python based on Python 3. Install 3.6.x.
rsync	Required when using Pgpool-II.
sysstat	Required when using <code>pgx_fjqssinf</code> command. Set up the <code>sar</code> command after installation.
redhat-lsb	-
libselinux	Required for <code>sepgsql</code> .
tcl	Required when using PL/Tcl. Install 8.6.
unzip	-

Package name	Remarks
xz-libs	-
zlib	-
java-1.8.0-openjdk	Required when using the database multiplexing and WebAdmin. Use build 1.8.0.312.b07 or later for ppc64le architecture.

- The following packages are required for operations on SLES 15.

Package name	Remarks
dstat	Required when using parallel scan.
gdb	Required when using pgx_fjqssinf command.
glibc	-
iputils	Required for Mirroring Controller.
libaudit1	-
libgcc_s1	-
libc-2.27-1	Provides collation support. Install 65.
liblz4-1	-
libmemcached	Required when using Pgpool-II.
libncurses6	-
libstdc++6	-
libz1	-
llvm	Install version 7.0.x of llvm to run SQL with runtime compilation (just-in-time compilation) and add the directory where the shared library libLLVM-*.so is located to the environment variable LD_LIBRARY_PATH. FUJITSU Enterprise Postgres 14 uses runtime compilation by default. If you do not want to use runtime compilation, turn off the jit parameter in postgresql.conf. You do not need to install llvm if you turn off the jit parameter. Failure to install llvm without turning off the jit parameter may result in an error when executing SQL. For more information about runtime compilation, see "Just-in-Time Compilation (JIT)" in the PostgreSQL Documentation.
libLLVM7	Install version 7.0.x.
net-tools	-
pam	Required when using PAM authentication.
perl	Required when using PL/Perl. Install 5.26.
protobuf-c	Required if using the Transparent Data Encryption feature when using a key management system as a keystore. Install 1.3.2.
python3	Required when using PL/Python based on Python 3. Install 3.6.x.
rsync	Required when using Pgpool-II.

Package name	Remarks
sysstat	Required when using pgx_fjqssinf command. Set up the sar command after installation.
tcl	Required when using PL/Tcl. Install 8.6.
java-1_8_0-openjdk	Required when using the database multiplexing and WebAdmin. Use build 1.8.0.312 or later for ppc64le architecture.

2.2 Related Software

There is no exclusive Software.

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

Table 2.1 Connectable client

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



Note

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

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

Table 2.2 Connectable server assistant

OS	Product name
Linux	FUJITSU Software Enterprise Postgres Advanced Edition 14 or 14 SP1

2.3 Excluded Software

There are no exclusive products.

2.4 Required Patches

There are no required patches.

2.5 Hardware Environment

The following hardware is required to use FUJITSU Enterprise Postgres.

Memory

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 FUJITSU Enterprise Postgres. If necessary, increase the size of the file system.

Table 2.3 Disk space required for installation

Directory	Required disk space (Unit: MB)
/etc	1
Installation destination of the server	211
Installation destination of WebAdmin	293
Installation destination of the client	121
Installation destination of PgPool	21

2.7 Supported System Environment

This section describes the supported system environment.

2.7.1 TCP/IP Protocol

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



Note

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

2.7.2 File System

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".

2.8 PostgreSQL Version Used for FUJITSU Enterprise Postgres

FUJITSU Enterprise Postgres is based on PostgreSQL 14.0.

2.9 Notes on Using Streaming Replication

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

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



Note

Streaming replication cannot be used in combination with Open Source PostgreSQL. It should also be used between instances running on same architecture (i.e. an instance running on Intel64 should not be used to stream instance running on ppc64le)

2.10 Key Management System Requirements

If you use a key management system as a keystore to use the Transparent Data Encryption feature, the following conditions must be met.

Protocol

Key management systems must use the Key Management Interoperability Protocol (KMIP) Version 1.4 protocol.

Encryption Key

The encryption key used must be able to be created or brought into the KMIP server under the following conditions.

- AES 256 bit symmetric key

A Managed Object that meets the following criteria:

- Cryptographic Algorithm : AES
- Cryptographic Length : 256
- Key not wrapped

Operation

The following operations using the KMIP protocol must be supported:

- Get operation

Encryption keys can be returned in Key Format Type: Raw format.

Client authentication

You must be able to authenticate and authorize clients in the following ways:

- The registered client certificate can authenticate the client and authorize access to the encryption key.

Quantitative Limits

FUJITSU Enterprise Postgres can receive a maximum response size of 8192 bytes from a key management system. Any further response results in an error.

Chapter 3 Installation

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

3.1 Pre-installation Tasks

Check the system environment below before installing FUJITSU Enterprise Postgres.

Check the disk space

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

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

Reconfigure the disk partition if disk space is insufficient.

Set JAVA_HOME

Ensure that Open JRE 8 is installed, and export the JAVA_HOME environment variable.

```
#export JAVA_HOME="OpenJre8InstallDir"
```

Refer to "[Appendix L Procedure when Modifying the JRE Installation](#)" for information on modifying JRE after installation.

Executable Users

Installation and uninstallation is performed by superuser.

On the system, run the following command to become superuser.

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

Determine the preferred WebAdmin configuration

WebAdmin can be installed in two configurations:

- Single-server
- Multiserver



See

.....
Refer to "[Appendix I Determining the Preferred WebAdmin Configuration](#)" for details.
.....

Determining the Pgpool-II System Configuration

The system configuration when using Pgpool-II is as follows:

- Place on database server
- Place on application server
- Place on dedicated server



See

.....
Refer to "[Appendix J System Configuration when using Pgpool-II](#)".
.....

3.2 Run Installation

Install according to the following procedure:



- The following characters can be used as input values:
Alphanumeric characters, hyphens, commas and forward slashes
- When reinstalling the product, back up the following folder in which the WebAdmin instance management information is stored:

`webAdminInstallFolder/data/fevpa`

Follow the procedure below to perform the backup.

1. Stop the WebAdmin server. Refer to "[B.1.3 Stopping the Web Server Feature of WebAdmin](#)" for details.
2. Back up the following folder:

`webAdminInstallFolder/data/fevpa`

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

1. Stop applications and programs

If the installation method is the following, all applications and programs that use the product must be stopped:

- Reinstallation

Before starting the installation, stop the following:

- Applications that use the product
- Connection Manager
- Instance
- Web server feature of WebAdmin

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

Example

If WebAdmin is installed in `/opt/fsepv<x>webadmin`:

```
# cd /opt/fsepv<x>webadmin/sbin
# ./WebAdminStop
```

- Mirroring Controller

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

Example

```
$ mc_ctl stop -M /mdir/inst1
```

- pgBadger
- Pgpool-II

2. Mount the DVD drive

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

Example

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

Here /dev/dvd is the device name for the DVD drive (which may vary depending on your environment), and /media/dvd is the mount point (which may need to be created before calling the command).



Note

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.

3. Run the installation

Select and install the following packages (rpm) with the rpm command.

Feature Name	Operating System	Package (Path)
Server	RHEL8	SERVER/Linux/packages/r80ppc64le/FJSVfsep-SV-*.rpm
	SLES 15	SERVER/Linux/packages/SUSE15ppc64le/FJSVfsep-SV-*.rpm
WebAdmin	RHEL8	WEBADMIN/Linux/packages/r80ppc64le/FJSVfsep-WAD-*.rpm
	SLES 15	WEBADMIN/Linux/packages/SUSE15ppc64le/FJSVfsep-WAD-*.rpm
Client	RHEL8	CLIENT64/Linux/packages/r80ppc64le/FJSVfsep-CL-*.rpm
	SLES 15	CLIENT64/Linux/packages/SUSE15ppc64le/FJSVfsep-CL-*.rpm
Pgpool-II	RHEL8	PGPOOL2/Linux/packages/r80ppc64le/FJSVfsep-POOL2-*.rpm
	SLES 15	PGPOOL2/Linux/packages/SUSE15ppc64le/FJSVfsep-POOL2-*.rpm

*is the version, OS, etc.

Example

In the following example, /media/dvd is the name of the mount point where the DVD is mounted.

Below is an example of new installation on RHEL8.

```
# cd /media/dvd/SERVER/Linux/packages/r80ppc64le
# rpm -ivh FJSVfsep-SV-14-1401-0.el8.ppc64le.rpm
```

Below is an example of new installation on SLES 15.

```
# cd /media/dvd/SERVER/Linux/packages/SUSE15ppc64le
# rpm -ivh FJSVfsep-SV-14-1401-0.s15.ppc64le.rpm
```

Below is an example of reinstallation on RHEL8.

```
# cd /media/dvd/SERVER/Linux/packages/r80ppc64le
# rpm -ivh --replacepks FJSVfsep-SV-14-1401-0.el8.ppc64le.rpm
```

Below is an example of reinstallation on SLES 15.

```
# cd /media/dvd/SERVER/Linux/packages/SUSE15ppc64le
# rpm -ivh --replacepks FJSVfsep-SV-14-1401-0.s15.ppc64le.rpm
```



Note

If you perform reinstallation, and an installation prefix (in the --prefix option of the rpm command) was used for the new installation, you must use the same prefix.

4. Setting Up WebAdmin

When using WebAdmin, use the WebAdminSetup command to set up WebAdmin. Refer to "[B.1 Setting Up WebAdmin](#)" for information on the WebAdmin setup procedure.

5. Set the installation environment to be used by Mirroring Controller

When using Database Multiplexing, use the mc_update_jre_env command to set the installation environment to be used by Mirroring Controller.

Example

```
# export JAVA_HOME="OpenJRE8InstallDir"
# /opt/fsepv<x>server64/bin/mc_update_jre_env
```

Chapter 4 Setup

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

4.1 Operating Method Types and Selection

This section describes how to operate FUJITSU Enterprise Postgres.

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

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

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

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

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

Usage method

Usage is started by using WebAdmin to create the instance.

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



Note

Do not use a server command other than `pgx_dmpall` and `pgx_keystore` or a server application. Operation modes that use server commands and server applications cannot be used in conjunction with WebAdmin. If used, WebAdmin will not be able to manage the instances correctly.

In addition, to perform a backup by copy command from the `pgx_dmpall` command, select the operating method using the server commands.

Refer to Reference and the PostgreSQL Documentation for information on server commands and server applications.

Advanced operation management using server commands

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

Refer to the Operation Guide for details.

Operation		Operation with the GUI	Operation with commands
Setup	Creating an instance	WebAdmin is used. The server machine capacity, and the optimum parameter for operations using WebAdmin, are set automatically.	The configuration file is edited directly using the <code>initdb</code> command.
	Creating a standby instance	WebAdmin is used. WebAdmin performs a base backup of the source instance and creates a standby instance.	A standby instance is created using the <code>pg_basebackup</code> command.
	Changing the configuration files	WebAdmin is used.	The configuration file is edited directly.
Starting and stopping an instance		WebAdmin is used.	The <code>pg_ctl</code> command is used.

Operation		Operation with the GUI	Operation with commands
Creating a database		None.	This is defined using the psql command or the application after specifying the DDL statement.
Backing up the database		WebAdmin, or the pgx_dmpall command, is used.	It is recommended that the pgx_dmpall command be used. Recovery to the latest database can be performed.
Database recovery		WebAdmin is used.	To use the backup that was performed using the pgx_dmpall command, the pgx_rcvall command is used.
Monitoring	Database errors	The status in the WebAdmin window can be checked.	The messages that are output to the database server log are monitored.
	Disk space	The status in the WebAdmin window can be checked. A warning will be displayed if the free space falls below 20%.	This is monitored using the df command of the operating system, for example.
	Connection status	None.	This can be checked referencing pg_stat_activity of the standard statistics view from psql or the application.



See

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

4.2 Preparations for Setup

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

4.2.1 Creating an Instance Administrator

Decide 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
```



Note

The following note applies if using WebAdmin for operations:

- If the password is changed for the user account of the instance administrator, set the changed password using ALTER ROLE WITH ENCRYPTED PASSWORD.

4.2.2 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. Specify a space other than that under the database cluster.
Transaction log	Stores log information in preparation for a crash recovery or rollback. This is the same as the WAL (Write Ahead Log).
Archive log	Stores log information for recovery
Mirrored transaction log (mirrored WAL)	Enables a database cluster to be restored to the state immediately before an error even if both the database cluster and transaction log fail when performing backup/recovery operations using the pgx_dmpall command or WebAdmin.
Corefile	FUJITSU Enterprise Postgres process corefile output when an error occurs with a FUJITSU Enterprise Postgres process.

Examples of disk deployment

The following are examples of disk deployment:

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

Proposal for disk deployment using WebAdmin

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

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

Preparing directories

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

The following shows the directories that need to be prepared:

Directory to be prepared	Using WebAdmin	Using the initdb command
Data storage destination	Y (*1)	Y
Backup data storage destination	O (*1) (*4)	O

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

Y: Required

O: Optional

N: Not required

*1: WebAdmin automatically creates a directory

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

*3: The corefile path is as follows:

```
/var/tmp/fsep_version/instanceAdminUser_instanceNamePortNumber/core
```

version: product version_WA_architecture

instanceAdminUser: operating system user name

PortNumber: port number specified when creating the instance

Example:

```
/var/tmp/fsep_140_WA_64/naomi_myinst27599/core
```

Note that resources placed in /var/tmp that have not been accessed for 30 days or more will be deleted by the default settings of the operating system. Consider excluding them from deletion targets or changing the output destination in the operating system settings.

To change the output destination, configure the core_directory and the core_contents parameters in postgresql.conf. Refer to "Parameters" in the Operation Guide for information on the settings for these parameters.

*4: This directory is required when instance backup is enabled.



Note

- The directories must meet the following conditions:
 - The directory owner must be the OS user account that you want to be the instance administrator
 - The directory must have write permission
 - The directory must be empty
- If you use WebAdmin, you cannot use directories mounted over the network. Examples include NFS (Network File System) and CIFS (Common Internet File System). Also, even if you are not using WebAdmin, do not use these directories unless you are creating tablespaces on a storage device on your network.

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 -p /database/inst1
# chown -R fsepuser:fsepuser /database/inst1
# chmod 700 /database/inst1
```

4.2.3 Estimating Resources

Estimate the resources to be used on the FUJITSU Enterprise Postgres.

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

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

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

4.2.4 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.

If you specify systemd-coredump as the core_pattern, the core file is not placed in the location specified by the core_directory parameter. See the systemd-coredump (8) man page for the location of core files.

Use coredumpctl to retrieve core files. For more information about using coredumpctl, see the coredumpctl (1) man page.

4.3 Creating Instances

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

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

Creating multiple instances

Multiple instances can be created.

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

Features that cannot be set up using WebAdmin

The "Storage data protection using transparent data encryption" feature cannot be set up using WebAdmin.

To set up this feature in an instance created by WebAdmin, perform the additional setup tasks detailed in "Storage Data Protection using Transparent Data Encryption" in the Operation Guide.



Note

- Instances created using the initdb command (command line instances) can be managed using WebAdmin, however, they must first be imported into WebAdmin. Refer to "[4.3.1.4 Importing Instances](#)" for details.
- Always use WebAdmin to delete instances that were created or imported using WebAdmin. Because WebAdmin management information cannot be deleted, WebAdmin will determine that the instance is abnormal.
- Databases with the names 'template0' and 'template1' are automatically created when an instance is created. These databases are used as the templates for databases created later. Furthermore, a default database with the name 'postgres' is automatically created, which will be used with FUJITSU Enterprise Postgres commands. It is important that you do not delete these databases created by default.

4.3.1 Using WebAdmin

This section describes how to create an instance using WebAdmin.

WebAdmin must be set up correctly before it can be used. Refer to "[B.1 Setting Up WebAdmin](#)" for details. Additionally, if WebAdmin needs to be configured to use an external repository database, refer to "[B.3 Using an External Repository for WebAdmin](#)" for details.

It is recommended to use the following browsers with WebAdmin:

- Microsoft Edge (Build41 or later)

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

Configure your browser to allow cookies and pop-up requests from the server on which FUJITSU Enterprise Postgres is installed.

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



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.
- If the same instance is operated from multiple WebAdmin versions, it will not work correctly. Always use the latest version of WebAdmin for instance operations.
- For efficient use of WebAdmin, it is recommended not to use the browser [Back] and [Forward] navigation buttons, the [Refresh] button, and context-sensitive menus, including equivalent keyboard shortcuts.
- Copying and pasting the WebAdmin URLs are not supported. Additionally, bookmarking of WebAdmin URLs is not supported.
- It is recommended to match the language between the instance server locale and WebAdmin.
- WebAdmin supports only one language: English.
- It is recommended to change the WebAdmin language setting from the instance details page only.
- It is recommended to operate WebAdmin using the WebAdmin launcher.
- WebAdmin uses the labels "Data storage path", "Backup storage path" and "Transaction log path" to indicate "data storage destination", "backup data storage destination" and "transaction log storage destination" respectively. In this manual these terms are used interchangeably.
- If the browser was not operated for a fixed period (about 30 minutes), the session will time out and the login page will be displayed again for the next operation.
- Port access permissions
If a port is blocked (access permissions have not been granted) by a firewall, enable use of the port by granting access. Refer to the vendor document for information on how to grant port access permissions.
Consider the security risks carefully when opening ports.
- When creating or importing an instance in WebAdmin, set the `log_directory` parameter in `postgresql.conf` to `'/var/tmp/fsep_version/instanceAdminUser_instanceNamePortNumber/log'`. Note that resources placed in `/var/tmp` that have not been accessed for 30 days or more will be deleted by the default settings of the operating system. Therefore, consider excluding instances created using WebAdmin from deletion targets in the operating system settings if you need to stop those instances for a long time.

4.3.1.1 Logging in to WebAdmin

This section describes how to log in to WebAdmin.

Startup URL for WebAdmin

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

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

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

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

Logging in to WebAdmin

Click [Launch WebAdmin] in the startup URL window to start WebAdmin and display the login window. Enter the instance administrator user ID (operating system user account name) and password, and log in to WebAdmin. User credential (instance administrator user ID and password) should not contain hazardous characters. Refer to "[Appendix C WebAdmin Disallow User Inputs Containing Hazardous Characters](#)".


4.3.1.2 Creating an Instance


This section describes how to create an instance.

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 H Configuring Kernel Parameters](#)" for information on configuring parameter names.

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

1. Start WebAdmin, and log in to the database server.
2. In the [Instances] tab, click .
3. Enter the information for the instance to be created.



Enter the following items:

- [Configuration type]: Whether to create a standalone instance or an instance that is part of a cluster.
- [Server product type]: Set the following:
 - FUJITSU Enterprise Postgres 14 Instances
- [Location]: Whether to create the instance in the server that the current user is logged into, or in a remote server. The default is "Local", which will create the instance in the server machine where WebAdmin is currently running.

- [Instance name]: Name of the database instance to manage

The name must meet the conditions below:

- Maximum of 16 characters
- The first character must be an ASCII alphabetic character
- The other characters must be ASCII alphanumeric characters
- [Instance port]: Port number of the database server
- [Data storage path]: Directory where the database data will be stored
- [Backup]: Whether to enable or disable the WebAdmin backup feature. The default is "Enabled". Select "Disabled" to disable all backup and restore functionality for the instance. If "Enabled" is selected, enter the following item:
 - [Backup storage path]: Directory where the database backup will be stored
- [Transaction log path]: Directory where the transaction log will be stored
- [Encoding]: Database encoding system
- [WAL file size]: Allow the WAL file size to be set when creating an instance. The default is 16 MB if the field is blank. The size specified must be a power of 2 between 1 and 1024. This option is not available for standby instances.

If "Remote" is selected for [Location], enter the following additional items:

- [Host name]: Name of the host where the instance is to be created
- [Operating system credential]: Operating system user name and password for the remote machine where the instance is to be created
- [Remote WebAdmin port for standalone]: Port in which WebAdmin is accessible in the remote machine




Note

- Refer to "[4.2.2 Preparing Directories for Resource Deployment](#)" - "Considerations when deploying resources" for information on points to consider when determining the data storage path, backup storage path, and transaction log path.
- The following items can be modified after the instance has been created.

- Instance name
- Port number
- Backup storage path

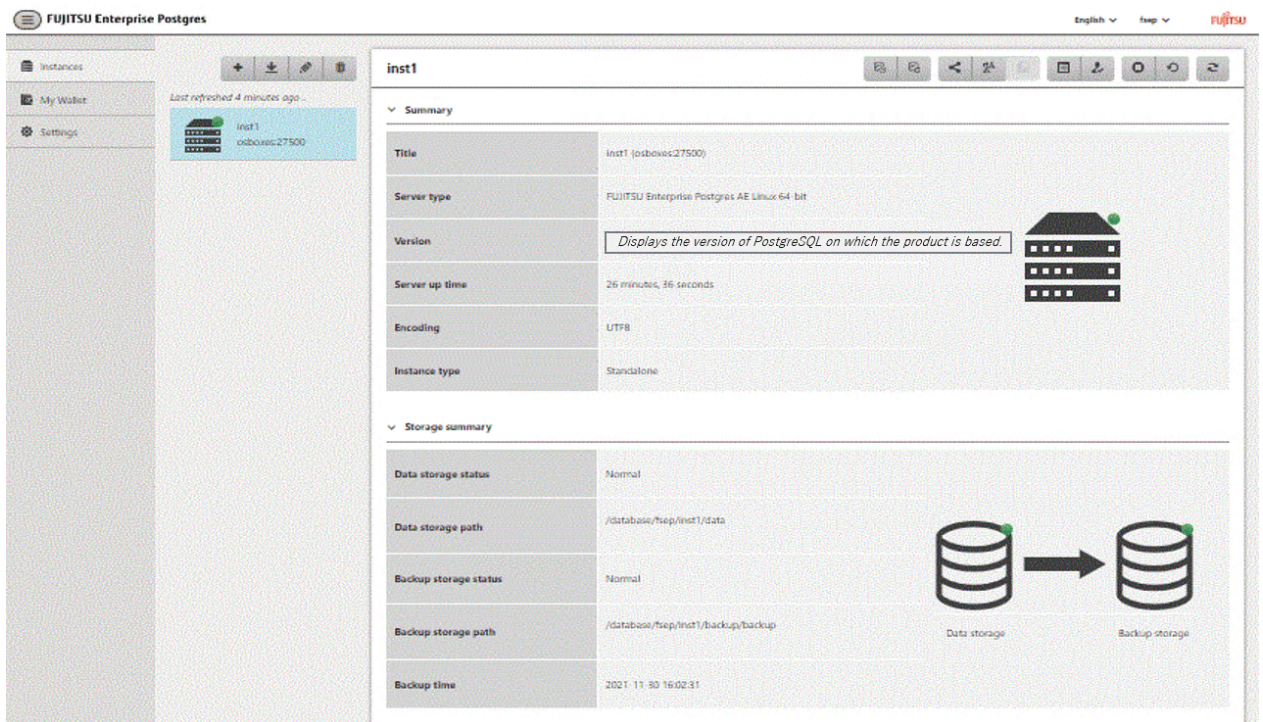
Refer to "[Editing instance information](#)" for details.

- Do not specify directories that include symbolic link or multibyte characters when specifying the data storage destination or backup data storage destination.
- In the instance that is created using WebAdmin, the locale of the character set to be used in the database, and the locale of the collating sequence, are fixed using C.
- For enhanced security, WebAdmin encrypts the superuser password using SCRAM-SHA-256 authentication.
- Host name and Operating system credential (Operating system user name and password) should not contain hazardous characters. Refer to "[Appendix C WebAdmin Disallow User Inputs Containing Hazardous Characters](#)".

4. Click  to create the instance.

If the instance is created successfully, a message indicating the same will be displayed.

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



6. Back up the basic information that was set

Back up the WebAdmin management information periodically to ensure operational continuity when a fault occurs on the system disk. Follow the procedure below to perform the backup.

- a. Stop the WebAdmin server. Refer to "[B.1.3 Stopping the Web Server Feature of WebAdmin](#)" for details.
- b. Back up the following directory:

```
webAdminInstallDir/data/feepwa
```

4.3.1.3 Changing Instance Settings

You can change the information that is set when an instance is created.

Change the following settings to suit the operating and management environment for FUJITSU Enterprise Postgres.

- [Instance configuration](#)
 - Character encoding
 - Communication
 - SQL options
 - Memory
 - Streaming replication
- [Changing client authentication information](#)
- [Editing instance information](#)



Information

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

- postgresql.conf

- pg_hba.conf






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

- postgresql.conf
- pg_hba.conf

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

- pg_hba.conf
- pg_ident.conf


Instance configuration


1. Start WebAdmin and log in to the database server.
2. In the [Instances] tab, click .
3. Click  to change the configuration.
4. Click  to save your changes.





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

Changing client authentication information

1. Start WebAdmin and log in to the database server.
2. In the [Instances] tab, click .

Click  to register new authentication information.

To change authentication information, select the information, and then click .

To delete authentication information, select the information, and then click .



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

Editing instance information

Use the [Edit instance] page to modify the following items for an instance:

- Instance name
- Port number

- Backup storage path

1. In the [Instances] tab, click . The [Edit instance] page is displayed.

2. Modify the relevant items.


If [Backup storage path] is changed, [Backup management] is enabled. Select the required option:

Retain existing backup: Create a backup in [Backup storage path] and retain the existing backup in its original location.

Copy existing backup to new path: Copy the existing backup to [Backup storage path]. A new backup will not be created. The existing backup will be retained in its original location.

Move existing backup to new path: Move the existing backup to [Backup storage path]. A new backup will not be created.

Remove existing backup: Create a backup in [Backup storage path]. The existing backup will be removed.

3. Click  to save your changes.



Note

- The [Edit instance] page is also displayed when the user selects 'Navigate to the "Edit instance" page' from the [Anomaly Error] dialog box. Refer to "Anomaly Detection and Resolution" in the Operation Guide for information on what takes place when an anomaly is detected.
- When [Instance name] or [Instance port] is modified, the log_directory and core_directory parameters in postgresql.conf are updated. Also, the specified directories are created if they do not exist.

Refer to [“4.3.1.4 Importing Instances”](#) for information on the format of these directories.


4.3.1.4 Importing Instances

Instances can be created using WebAdmin, or via the command line using the initdb command. Instances created using the initdb command (command line instances) can be managed using WebAdmin, however, they must first be imported into WebAdmin.

This section explains how to import command line instances into WebAdmin.

1. In the [Instances] tab, click . The [Import instance] page is displayed.

2. Enter the information for the instance being imported. Refer to [“4.3.1.2 Creating an Instance”](#) for information on the items that need to be entered.

3. Click  to import the instance.



Note

- Importing neither starts nor stops the instance.
- The following restrictions apply to instance import:
 - Any instance already managed by WebAdmin cannot be imported again.
 - The postgresql.conf file must be located in the same directory as [Data storage path].
 - Read/write permissions are required for [Data storage path].
 - The location specified in postgresql.conf for the following files must not have been changed:
 - hba_file
 - ident_file
 - If the following file contains records that span multiple lines, change the record to a single line before importing.
 - pg_hba.conf
 - pg_ident.conf

- If the instance is part of a cluster that is monitored by Mirroring Controller, WebAdmin will be unable to detect the Mirroring Controller settings.
- Instances making use of Mirroring Controller functionality should not be imported, because subsequent operations on those instances may cause unexpected and undesirable side-effects.
- It is not possible to import and operate an instance that uses a directory mounted by Network File System (NFS).
- You must make the following changes to the parameters in postgresql.conf prior to importing the instance in WebAdmin.

Parameter	Requirements
port	The port parameter should be uncommented.

The log_directory and core_directory parameters in postgresql.conf are updated during import. Also, the specified directories are created if they do not exist.

The format of these directories is as follows:

log_directory: '/var/tmp/fsep_version/instanceAdminUser_instanceNamePortNumber/log'

core_directory: '/var/tmp/fsep_version/instanceAdminUser_instanceNamePortNumber/core'

version: product version_WA_architecture

instanceAdminUser: operating system user name

PortNumber: port number specified when creating the instance

Examples:

log_directory: '/var/tmp/fsep_140_WA_64/naomi_myinst27599/log'

core_directory: '/var/tmp/fsep_140_WA_64/naomi_myinst27599/core'

- When a standby instance is imported, a valid entry, using the IP address of the standby instance, must exist in the pg_hba.conf file of the master instance to allow the standby instance to connect to the master instance.
- When a standby instance is imported, the value for "host" in the primary_conninfo parameter of postgresql.auto.conf should match the host name of the master instance.
- When a standby instance is imported, you cannot specify "passfile" in the primary_conninfo parameter of postgresql.auto.conf. Be sure to specify "password".
- Instances created by other operating systems cannot be imported.
- If a instance is being imported while it is running, WebAdmin will encrypt the superuser password using SCRAM-SHA-256 authentication.

4.3.2 Using the initdb Command

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



Note

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

Consider the security risks carefully when opening ports.

4.3.2.1 Editing Kernel Parameters

Refer to "[Appendix H 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 (command line instances) can be managed using WebAdmin, however, they must first be imported into WebAdmin. Refer to "4.3.1.4 Importing Instances" for details.
- If creating multiple instances, ensure that there is no duplication of port numbers or the directories that store database clusters.



See

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

The procedure to create an instance is described below.

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

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

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/fsepv<x>server64".

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

sh, bash

```
$ PATH=/opt/fsepv<x>server64/bin:$PATH ; export PATH
$ MANPATH=/opt/fsepv<x>server64/share/man:$MANPATH ; export MANPATH
$ LD_LIBRARY_PATH=/opt/fsepv<x>server64/lib:$LD_LIBRARY_PATH ; export LD_LIBRARY_PATH
```

csh, tcsh

```
$ setenv PATH /opt/fsepv<x>server64/bin:$PATH
$ setenv MANPATH /opt/fsepv<x>server64/share/man:$MANPATH
$ setenv LD_LIBRARY_PATH /opt/fsepv<x>server64/lib:$LD_LIBRARY_PATH
```

Note

If you execute any command other than FUJITSU Enterprise Postgres (OS commands, etc.) after LD_LIBRARY_PATH is set, remove the installation directory/lib from 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 --waldir=/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 FUJITSU 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.

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 `/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`.



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. Start an instance.

Start with the start mode of the `pg_ctl` command.

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 without the `-W` option specified

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

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

Example

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



See

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



Note

If the `-W` option is specified, the command will return without waiting for the instance to start. Therefore, it may be unclear as to whether instance startup was successful or failed.

4.4 Configuring Remote Connections

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

4.4.1 When an Instance was Created with WebAdmin

Settings related to connection

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

Change "listen_addresses" in postgresql.conf to modify the default behavior.
Refer to "[Appendix D Configuring Parameters](#)" for information on postgresql.conf.

Client Authentication Information settings

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

- Authentication of remote connections from local machines is performed.

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

4.4.2 When an Instance was Created with the initdb Command

Connection settings

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

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

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

Example

```
listen_addresses = '*'
```

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

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

Client authentication information settings

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



See

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

4.5 Other Settings

This section describes settings that are useful for operations.

4.5.1 Error Log Settings

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

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

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

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



Note

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



Point

- If you want fewer application errors being output to the 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 started in a failover operation, the cluster system will control the start or stop, therefore this feature should not be used. Also, when performing database multiplexing, refer to "Enabling Automatic Start and Stop of Mirroring Controller and Multiplexed Instances" in the Cluster Operation Guide (Database Multiplexing).

Note that "<x>" in paths indicates the product version.



Note

You should wait for time correction, network setup, and so on.

1. Create a unit file

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

`fujitsuEnterprisePostgresInstallDir/share/fsepsvoi.service.sample`

Example

In the following example, the installation directory is "/opt/fsepv<x>server64", and the instance name is "inst1".

```
# cp /opt/fsepv<x>server64/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	FUJITSU Enterprise Postgres <u><i>instanceName</i></u>	Specifies the feature overview. Specifies the name of the target instance. (*1)
Service	ExecStart	/bin/bash -c ' <u><i>installDir</i></u> /bin/pgx_symstd start <u><i>installDir dataStorageDestinationDir</i></u> '	Command to be executed when the service is started.
	ExecStop	/bin/bash -c ' <u><i>installDir</i></u> /bin/pgx_symstd stop <u><i>installDir dataStorageDestinationDir</i></u> '	Command to be executed when the service is stopped.
	ExecReload	/bin/bash -c ' <u><i>installDir</i></u> /bin/pgx_symstd reload <u><i>installDir dataStorageDestinationDir</i></u> '	Command to be executed when the service is reloaded
	User	<u>User</u>	OS user account of the instance administrator.
	Group	<u>Group</u>	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

To use the features compatible with Oracle databases, create a new instance and execute the following command for the "postgres" and "template1" databases:

```
CREATE EXTENSION oracle_compatible;
```

Features compatible with Oracle databases are defined as user-defined functions in the "public" schema created by default when database clusters are created, so they can be available for all users without the need for special settings.

For this reason, ensure that "public" (without the double quotation marks) is included in the list of schema search paths specified in the search_path parameter.

There are also considerations for use the features compatible with Oracle databases. Refer to "Precautions when Using the Features Compatible with Oracle Databases" in the Application Development Guide for details.

4.6 Setting Up and Removing OSS

This section explains how to set up OSS supported by FUJITSU Enterprise Postgres.

If you want to use OSS supported by FUJITSU Enterprise Postgres, follow the setup procedure.

If you decide not to use the OSS supported by FUJITSU Enterprise Postgres, follow the removing procedure.

To build and use OSS obtained from the web, etc., instead of OSS supported by FUJITSU Enterprise Postgres, see "[4.6.10 Build with PGXS](#)".



Information

- In this section, the applicable database that enables the features of each OSS is described as "postgres".
 - Execute CREATE EXTENSION for the "template1" database also, so that each OSS can be used by default when creating a new database.
-

Refer to "OSS Supported by FUJITSU Enterprise Postgres" in the General Description for information on OSS other than those described below.

4.6.1 oracle_fdw

4.6.1.1 Setting Up oracle_fdw

1. Add the path of the OCI library to the environment variable. The available version of the OCI library is 12.1 or later.
Add the installation path of the OCI library to the LD_LIBRARY_PATH environment variable.

2. As superuser, run the following command:

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

3. If a file named libclntsh.so.12.1 does not exist in your OCI library, create a symbolic link with the name libclntsh.so. 12.1 to libclntsh.so.xx.1 (xx is the version of the OCI library).

```
# ln -s libclntsh.so.12.2 libclntsh.so.12.1
```

4. Restart FUJITSU Enterprise Postgres.
5. Execute CREATE EXTENSION for the database that will use this feature.
Use the psql command to connect to the "postgres" database.

```
postgres=# CREATE EXTENSION oracle_fdw;  
CREATE EXTENSION
```



Information

- If the OCI library is not installed on the server, install it using the Oracle client or Oracle Instant Client.
Refer to the relevant Oracle manual for information on the installation procedure.
 - If the version of the OCI library is updated, change the path of the OCI library in the LD_LIBRARY_PATH environment variable to the updated path. Also, re-create the symbolic link named libclntsh.so.12.1 if necessary.
-



Note

This feature cannot be used on instances created in WebAdmin. It can only be used via server commands.

4.6.1.2 Removing oracle_fdw

1. Execute DROP EXTENSION for the database that will use this feature.
Use the psql command to connect to the "postgres" database.

```
postgres=# DROP EXTENSION oracle_fdw CASCADE;  
DROP EXTENSION
```

2. As superuser, run the following command:

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



Information

The files copied during setup can be checked below.

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

4.6.2 pg_bigm

4.6.2.1 Setting Up pg_bigm

1. Set the postgresql.conf file parameters.
Add "pg_bigm" to the shared_preload_libraries parameter.
2. As superuser, run the following command:

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

3. Restart FUJITSU Enterprise Postgres.
4. Execute CREATE EXTENSION for the database that will use this feature.
Use the psql command to connect to the "postgres" database.

```
postgres=# CREATE EXTENSION pg_bigm;  
CREATE EXTENSION
```

4.6.2.2 Removing pg_bigm

1. Execute DROP EXTENSION for the database that will use this feature.
Use the psql command to connect to the "postgres" database.

```
postgres=# DROP EXTENSION pg_bigm CASCADE;  
DROP EXTENSION
```

2. As superuser, run the following command:

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

Information

The files copied during setup can be checked below.

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

3. Set the postgresql.conf file parameters.
Delete "pg_bigm" to the shared_preload_libraries parameter.
4. Restart FUJITSU Enterprise Postgres.

4.6.3 pg_hint_plan

4.6.3.1 Setting Up pg_hint_plan

1. Set the postgresql.conf file parameters.
Add "pg_hint_plan" to the "shared_preload_libraries" parameter.
2. As superuser, run the following command:

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

3. Restart FUJITSU Enterprise Postgres.
4. Run CREATE EXTENSION for the database that uses this feature.
The target database is described as "postgres" here.
Use the psql command to connect to the "postgres" database.

```
postgres=# CREATE EXTENSION pg_hint_plan;  
CREATE EXTENSION
```

See

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

4.6.3.2 Removing pg_hint_plan

Note

Unsetting pg_hint_plan will cause hints registered in the hint_plan.hints table to be lost. Therefore, it is recommended that pg_dump back up the hint_plan.hints table for each database if it is likely that pg_hint_plan will be used again later.

1. Execute DROP EXTENSION for the database that will use this feature.
Use the psql command to connect to the "postgres" database.

```
postgres=# DROP EXTENSION pg_hint_plan CASCADE;  
DROP EXTENSION
```

2. As superuser, run the following command:

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

Information

The files copied during setup can be checked below.

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

3. Set the postgresql.conf file parameters.
Delete "pg_hint_plan" to the shared_preload_libraries parameter.
4. Restart FUJITSU Enterprise Postgres.

4.6.4 pg_dbms_stats

4.6.4.1 Setting Up pg_dbms_stats

1. Set the postgresql.conf file parameter.
Add "pg_dbms_stats" to the "shared_preload_libraries" parameter.
2. As superuser, run the following command:

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

3. Restart FUJITSU Enterprise Postgres.
4. Run CREATE EXTENSION for the database that will use this feature.
The target database is described as "postgres" here.
Use the psql command to connect to the "postgres" database.

```
postgres=# CREATE EXTENSION pg_dbms_stats;  
CREATE EXTENSION
```

See

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

4.6.4.2 Removing pg_dbms_stats

Note

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

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

1. Execute DROP EXTENSION for the database that will use this feature.
Use the psql command to connect to the "postgres" database.

```
postgres=# DROP EXTENSION pg_dbms_stats CASCADE;  
DROP EXTENSION
```

2. As superuser, run the following command:

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

Information

The files copied during setup can be checked below.

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

3. Set the postgresql.conf file parameters.
Delete "pg_dbms_stats" to the shared_preload_libraries parameter.
4. Restart FUJITSU Enterprise Postgres.

4.6.5 pg_repack

4.6.5.1 Setting Up pg_repack

1. As superuser, run the following command:

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

2. Execute CREATE EXTENSION for the database that will use this feature.
Use the psql command to connect to the "postgres" database.

```
postgres=# CREATE EXTENSION pg_repack;  
CREATE EXTENSION
```

4.6.5.2 Removing pg_repack

1. Execute DROP EXTENSION for the database that will use this feature.
Use the psql command to connect to the "postgres" database.

```
postgres=# DROP EXTENSION pg_repack CASCADE;  
DROP EXTENSION
```

2. As superuser, run the following command:

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

Information

The files copied during setup can be checked below.

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

4.6.6 pg_rman

4.6.6.1 Setting Up pg_rman

1. As superuser, run the following command:

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

2. Restart FUJITSU Enterprise Postgres.

Information

Before initialization of the backup catalog, it is recommended to set the parameters below in postgresql.conf. Refer to the pg_rman manual (http://ossc-db.github.io/pg_rman/index-ja.html) for details.

- log_directory
- archive_mode
- archive_command

Note

This feature cannot be used on instances created in WebAdmin. It can only be used via server commands.

4.6.6.2 Removing pg_rman

1. As superuser, run the following command:

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

Information

The files copied during setup can be checked below.

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

2. Restart FUJITSU Enterprise Postgres.

4.6.7 pg_statsinfo

4.6.7.1 Setting Up pg_statsinfo

1. Set the postgresql.conf file parameters.
 - Add "pg_statsinfo" to the shared_preload_libraries parameter.
 - Specify the log file name for the log_filename parameter.
2. As superuser, run the following command:

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

3. Restart FUJITSU Enterprise Postgres.

Information

Note that pg_statsinfo forcibly overwrites the settings below.

- log_destination
 - "csvlog" is added and "stderr" is deleted.

- logging_collector

"on" is set.



This feature cannot be used on instances created in WebAdmin. It can only be used via server commands.

4.6.7.2 Removing pg_statsinfo

1. As superuser, run the following command:

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



The files copied during setup can be checked below.

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

2. Set the postgresql.conf file parameters.
 - Delete "pg_statsinfo" to the shared_preload_libraries parameter.
 - Delete the log file name for the log_filename parameter.
3. Restart FUJITSU Enterprise Postgres.

4.6.8 pgBadger

4.6.8.1 Setting Up pgBadger

1. Set the postgresql.conf file parameters.

Set the parameters so that the information required for analysis is output to the server log.
Refer to "Documentation" in the pgBadger website (<https://pgbadger.darold.net/>) for details.
The pgBadger material is stored under /opt/fsepv<x>server64/OSS/pgbadger.
2. Restart FUJITSU Enterprise Postgres.

4.6.8.2 Removing pgBadger

1. Set the postgresql.conf file parameters.

Restores information you specified during Setup.
2. Restart FUJITSU Enterprise Postgres.

4.6.9 Pgpool-II

4.6.9.1 Setting Up Pgpool-II

1. As superuser, run the following command:

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

2. Execute CREATE EXTENSION for the database that will use this feature.
Use the psql command to connect to the "postgres" database.

```
postgres=# CREATE EXTENSION pgpool_recovery;  
CREATE EXTENSION
```

3. Set the postgresql.conf file parameters.
Specify the path to pg_ctl for the pgpool.pg_ctl parameter.
4. Restart FUJITSU Enterprise Postgres.



Note

The online recovery feature of Pgpool-II cannot be used on instances created in WebAdmin. It can only be used via server commands.

4.6.9.2 Removing Pgpool-II

1. Execute DROP EXTENSION for the database that will use this feature.
Use the psql command to connect to the "postgres" database.

```
postgres=# DROP EXTENSION pgpool_recovery CASCADE;  
DROP EXTENSION
```

2. As superuser, run the following command:

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



Information

The files copied during setup can be checked below.

```
# find /opt/fsepv<x>server64/OSS/Pgpool-II
```

3. Set the postgresql.conf file parameters.
Delete the pg_ctl path for the pgpool.pg_ctl parameter.
4. Restart FUJITSU Enterprise Postgres.

4.6.10 Build with PGXS

Many PostgreSQL extensions are built using a build base for extensions called PGXS. Building with PGXS also generates files related to llvm. Depending on which version of llvm you are using, follow these steps:

4.6.10.1 Using the Default Version of llvm

The default version of llvm is described in "[2.1 Required Operating System](#)". If you want to use the default version of llvm, use the OSS documentation to build and install OSS.

4.6.10.2 Using a Non-Default Version of llvm

1. As superuser, copy the Makefile.global corresponding to the version of llvm you want to use. The following is an example of using version 11 of llvm. Makefile.global is overwritten when an emergency fix is applied or removed from Fujitsu Enterprise Postgres, this procedure should be performed each time a build is performed.

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



```
# cp /opt/fsepv<x>server64/lib/pgxs/src/Makefile.global-vsn11 /opt/fsepv<x>server64/lib/pgxs/src/Makefile.global
```

2. Follow the OSS documentation to build and install OSS.
3. As superuser, run the following command:. The following is an example of using version 11 of llvm:.

```
$ su -  
Password:*****  
# mv /opt/fsepv<x>server64/lib/bitcode/<OSS名>* /opt/fsepv<x>server64/lib/bitcode-vsn11/
```

4.6.10.3 Without llvm

If you do not use llvm, use the with _ llvm = no option when performing the build, as shown below. For other options, follow the OSS documentation.

```
# make USE_PGXS=1 with_llvm=no
```

4.7 Integration with Message-Monitoring Software

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

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

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

Example

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



See

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

4.8 Deleting Instances

This section explains how to delete an instance.

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



Note

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

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

Example


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

4.8.1 Using WebAdmin

This section explains how to delete an instance using WebAdmin.

Use the following procedure to delete an instance.


1. Stop the instance

In the [Instances] tab, select the instance to stop and click .

2. Back up files.

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

3. Delete the instance

In the [Instances] tab, select the instance to delete and click .



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

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

4.8.2 Using Server Commands

This section explains how to delete an instance using server commands.

Use the following procedure to delete an instance.

1. Stop the instance

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 FUJITSU Enterprise Postgres.

5.1 Run Uninstallation

Uninstall according to the following procedure:

Note that "xSPz" in sample windows indicates the version and level of products to uninstall and "<x>" in paths indicates the product version.



- All files and directories in the installation directory are deleted during uninstallation. If user files have been placed in the installation directory, back them up before uninstallation if necessary.
- To reinstall FUJITSU 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 directory shown below in which the WebAdmin instance management information had been defined before uninstalling FUJITSU Enterprise Postgres, and then restore the backed up directory to its original location once FUJITSU Enterprise Postgres has been reinstalled.

Follow the procedure below to perform the backup.

1. Stop the WebAdmin server. Refer to "[B.1.3 Stopping the Web Server Feature of WebAdmin](#)" for details.
2. Back up the following directory:

```
webAdminInstallDir/data/fepwa
```

1. Delete the operation information

If the FUJITSU 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.

```
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 programs

Before starting the uninstallation, stop the following:

- Applications that use the product
- Connection Manager
- Instance

Using WebAdmin

In the [Instances] tab, select the instance to stop and click .

Using server commands

Execute the pg_ctl command in stop mode.

```
$ /opt/fsepv<x>server64/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

If WebAdmin is installed in /opt/fsepv<x>webadmin:

```
# cd /opt/fsepv<x>webadmin/sbin
# ./WebAdminStop
```

- Mirroring Controller

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

Example

```
$ mc_ctl stop -M /mdir/inst1
```

- pgBadger

- Pgpool-II

3. Remove WebAdmin setup

When uninstall WebAdmin feature, execute the WebAdminSetup command to remove WebAdmin setup.

Example

If WebAdmin is installed in /opt/fsepv<x>webadmin:

```
# cd /opt/fsepv<x>webadmin/sbin
# ./WebAdminSetup -d
```

4. Verifying Installation Features

Verify that the feature to be removed is installed by executing the following command.

Feature Name	Package Name
Server	FJSVfsep-SV-<x>
WebAdmin	FJSVfsep-WAD-<x>
Client	FJSVfsep-CL-<x>
Pgpool-II	FJSVfsep-POOL2-<x>

* Where *x* is a number indicating the version.

Example

```
# rpm -qi FJSVfsep-SV-14
```

5. Run the uninstallation

Run the following command.

Example

```
# rpm -e FJSVfsep-SV-14
```

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

Appendix A Recommended WebAdmin Environments

This appendix describes the recommended WebAdmin environment. The following explanation is based on the assumption that Microsoft Edge is used unless otherwise stated.



The displayed screen varies depending on your environment, so check and set according to the screen.

A.1 Recommended Browser Settings

- Use a display resolution of 1280 x 768 or higher, and 256 colors or more.
- Select [Setting] >> [Appearance] >> [Font size] >> [Medium (Recommended)].
- Select [Setting] >> [Appearance] >> [Zoom] >> [100%].

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

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

1. Click [Setting] >> [Cookie and site permissions] >> [All Permissions] >> [Pop-ups and redirects].
If the [Block (Recommended)] switch is not on (blue), the pop-up blocker is not working, and no further action is required.
2. Under [Pop-ups and Redirects], click the [Allow] >> [Add] button.
3. In [Add Site], in [Site], enter the address of the server where you installed FUJITSU Enterprise Postgres and click the [Add] button.
4. Close Microsoft Edge.

Appendix B Setting Up and Removing WebAdmin

This appendix describes how to set up and remove WebAdmin.

Note that "<x>" in paths indicates the product version.

B.1 Setting Up WebAdmin

This section explains how to set up WebAdmin.

B.1.1 Setting Up WebAdmin

Follow the procedure below to set up WebAdmin.

1. Change to the superuser

Acquire superuser privileges on the system.

Example

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

2. Set the JAVA_HOME environment variable

Set the JAVA_HOME environment variable to the installation destination of Open JRE 8.

Example

```
# export JAVA_HOME="OpenJRE8InstallDir"
```

3. Set up WebAdmin

Set up WebAdmin.

Example

If WebAdmin is installed in /opt/fsepv<x>webadmin:

```
# cd /opt/fsepv<x>webadmin/sbin  
# ./WebAdminSetup
```

4. 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 starting 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.

fsep_140_WA_64_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.

fsep_140_WA_64_WebAdmin_Port2

WebAdmin automatic start

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



Note

- 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.
- Changing port numbers
When using WebAdmin in multiserver mode, it is recommended not to change WebAdmin ports after creating instances. Otherwise, the created instances may not be accessible through WebAdmin after the port is changed.

B.1.2 Starting the Web Server Feature of WebAdmin

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

1. Change to the superuser

Acquire superuser privileges on the system.

Example

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

2. Start the Web server feature of WebAdmin

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

Example

If WebAdmin is installed in /opt/fsepv<x>webadmin:

```
# cd /opt/fsepv<x>webadmin/sbin  
# ./WebAdminStart
```

B.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 WebAdmin is installed in /opt/fsepv<x>webadmin:

```
# cd /opt/fsepv<x>webadmin/sbin  
# ./WebAdminStop
```



- For efficient operation of WebAdmin, it is recommended that the Web server feature be stopped only during a scheduled maintenance period.
- When WebAdmin is used to create and manage instances in a multiserver configuration, the Web server feature must be started and running on all servers at the same time.

B.2 Removing WebAdmin

This section explains how to remove WebAdmin.

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

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 WebAdmin is installed in /opt/fsepv<x>webadmin:

```
# cd /opt/fsepv<x>webadmin/sbin  
# ./WebAdminSetup -d
```


B.3 Using an External Repository for WebAdmin


WebAdmin can be configured to use an external database, where it can store the various metadata information it uses. WebAdmin will use this database as a repository to store the information it uses to manage all the created instances. This can be a FUJITSU Enterprise Postgres database or an Open Source PostgreSQL V9.2 or later database.

Using an external database as a WebAdmin repository provides you with more flexibility in managing WebAdmin. This repository can be managed, backed up and restored as needed using command line tools, allowing users to have greater flexibility and control.

Follow the procedure below to set up the repository.

1. Start WebAdmin, and log in to the database server.

2. Click the [Settings] tab, and then click  in the [WebAdmin repository configuration] section.




The screenshot shows the 'FUJITSU Enterprise Postgres' settings interface. On the left is a sidebar with 'Instances', 'My Wallet', and 'Settings' (selected). The main area is titled 'Settings' and contains the 'WebAdmin repository configuration' section. This section has a table with the following fields: 'Database type' (FUJITSU Enterprise Postgres), 'Host name *' (rdbHost), 'Port *' (27500), 'Database name *' (webadmindb), 'User name *' (postgres), and 'Password *' (masked with dots). A 'Test connection' button is located below the table. To the right of the table is a note: '* indicates required field(s)'. Below the table is the 'User preferences' section, which includes a toggle for 'Auto-refresh instance' set to 'Disabled'. There are checkmark and reset icons in the top right of the configuration section, and an edit icon in the top right of the user preferences section.

3. Enter the following items:

- [Host name]: Host name of the database server
- [Port]: Port number of the database server
- [Database name]: Name of the database
- [User name]: User name to access the database
- [Password]: Password of the database user

Note

- Database type
It is recommended to use a FUJITSU Enterprise Postgres database as a repository. A compatible PostgreSQL database can also be used as an alternative.
- It is recommended to click [Test connection] to ensure that the details entered are valid and WebAdmin is able to connect to the target database.
- Host name, Database name, User name, Password should not contain hazardous characters. Refer to "[Appendix C WebAdmin Disallow User Inputs Containing Hazardous Characters](#)".

4. Click  to register the repository details.





Note

- Once the repository is set up, it can be changed any number of times by the user logged into WebAdmin. When a repository is changed:
 - It is recommended to preload the backup into this database.
 - If the data is not preloaded, WebAdmin will create a new repository.
- The database repository can be set up even after WebAdmin was already used to create instances. In that scenario, the instances already created are retained and can continue to be operated on.
- If the instance used as a repository is stopped, WebAdmin will be unusable. For this reason, it is recommended to be familiar with starting an instance from the command line. If the instance is stopped for any reason, start it from the command line and WebAdmin will be usable again.

B.4 Using the WebAdmin Auto-Refresh Feature

The WebAdmin auto-refresh feature automatically refreshes the operating status of all instances in the Instance list at the specified interval. It also refreshes the details of the selected instance.

Follow the procedure below to configure the auto-refresh options.

1. Click the [Settings] tab, and then click  in the [User preferences] section.
2. Enter the following items:
 - [Auto-refresh instance]: To use the auto-refresh feature, select "Enabled". The default is "Disabled".
 - [Refresh interval (seconds)]: Number of seconds between each refresh. This is a countdown timer, which is reset every time the instance status is refreshed by any operation. Specify a value from 30 to 3600 (seconds). The default is 30.
3. Click  to save the auto-refresh settings.



Note

- Auto-refresh will run only if the [Instances] page is displayed and no user-initiated operation is in progress.
- A text indicator, which is independent of auto-refresh, is displayed at the top of the Instance list. It is dynamically updated to display when the page was last refreshed.

Appendix C WebAdmin Disallow User Inputs Containing Hazardous Characters

WebAdmin considers the following as hazardous characters, which are not allowed in user inputs.

- | (pipe sign)
- & (ampersand sign)
- ; (semicolon sign)
- \$ (dollar sign)
- % (percent sign)
- @ (at sign)
- ' (single apostrophe)
- " (quotation mark)
- \ ' (backslash-escaped apostrophe)
- \ " (backslash-escaped quotation mark)
- <> (triangular parenthesis)
- () (parenthesis)
- + (plus sign)
- CR (Carriage return, ASCII 0x0d)
- LF (Line feed, ASCII 0x0a)
- , (comma sign)
- \ (backslash)

Appendix D Configuring Parameters

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

- [postgresql.conf](#)

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

- [pg_hba.conf](#)

Contains various items of information related to client authentication.

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

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



See

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



Note

WebAdmin checks for port number and backup storage path anomalies when various operations are performed. An anomaly occurs when the value of [Port number] and/or [Backup storage path] in WebAdmin is different from the value of the corresponding parameter in postgresql.conf. Refer to "Anomaly Detection and Resolution" in the Operation Guide for details.

postgresql.conf

Parameters that can be changed in WebAdmin

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

Section	WebAdmin item	postgresql.conf file parameter
Instance Configuration		
Character encoding	Character set	client_encoding
	Message locale	lc_messages
Communication	Max connections	max_connections
SQL options	Transform NULL format	transform_null_equals
	Date output format	DateStyle (*1)
	Interval output format	IntervalStyle
	Number of digits for floating values	extra_float_digits
	Transaction isolation levels	default_transaction_isolation
	Currency format	lc_monetary
	Date and time format	lc_time
	Numerical value format	lc_numeric
Memory	Sort memory (KB)	work_mem
	Shared buffers (KB)	shared_buffers
Streaming replication	WAL level	wal_level

Section	WebAdmin item	postgresql.conf file parameter
	Maximum WAL senders	max_wal_senders
	WAL save size (MB)	wal_keep_size
	Hot standby	hot_standby
	Synchronous standby names	synchronous_standby_names
	WAL receiver timeout (ms)	wal_receiver_timeout
Edit instance		
	Instance name	n/a
	Instance port	port
	Backup storage path	backup_destination

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



Information

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

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

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

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

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

- Also check if the memory used exceeds the memory installed (refer to "[Parameters automatically set by WebAdmin according to the amount of memory](#)").
- When modifying "Shared buffers" or "Max connections", edit the kernel parameter. Refer to "[Appendix H Configuring Kernel Parameters](#)", and "Managing Kernel Resources" in "Server Administration" in the PostgreSQL Documentation for details.

Parameters set by WebAdmin

The following postgresql.conf parameters are set by WebAdmin during instance startup (they will be ignored even if specified in postgresql.conf):

Parameter	Value
listen_addresses	*
log_destination	'stderr,syslog'
logging_collector	on
log_line_prefix	'%e: %t [%p]: [%l-1] user = %u,db = %d,remote = %r app = %a '
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) (*1)

*1: The value will be capped at 2097151 KB.

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

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

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

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

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



See

Kernel parameters need to be tuned according to the parameters being changed. Refer to "[Appendix H 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):
 - 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
- 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 E Estimating Database Disk Space Requirements

This appendix describes how to estimate database disk space requirements.

E.1 Estimating Table Size Requirements

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

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

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

Item	Estimation formula (bytes)
(4) Number of pages required for storing records	Total number of records / (3) number of records per page - The result will be rounded up to the next integer.
(5) Amount of space	(4) Number of pages required for storing records x page length x safety factor (*1) *1: Specify 2.0 or higher. - This is the safety factor assumed if vacuuming is performed for garbage collection in tables and indexes.

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

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

E.2 Estimating Index Size Requirements

This section provides the formulas for estimating index size requirements.

FUJITSU Enterprise Postgres provides six index types: B-tree, Hash, GiST, GIN, SP-GiST, and VCI. If you do not specify the index type in the CREATE INDEX statement, a B-tree index is generated.

The following describes how to estimate a B-tree index. Refer to "[E.7 Estimating VCI Disk Space Requirements](#)" for information on how to estimate VCI.

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

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

Item	Estimation formula (bytes)
(1) Entry length	8 (*1) + key data length (*2) *1: Entry header *2: The key data length depends on its data type (refer to " E.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 E.4 Estimation formula when the key data length exceeds 512 bytes " to estimate the key data length.
(2) Page size requirement	8192 (*1) × fillfactor (*2) - 24 (*3) - 16 (*4) *1: Page length (8192) *2: Value of the fillfactor specified in the index definitions (if omitted, 90%) In the case of indexes of primary key constraints and unique constraints, the value of the fillfactor specified for each constraint in the table definitions (if omitted, 90%) *3: Page header (24) *4: Special data (16)

Item	Estimation formula (bytes)
	- The calculated (2) page size requirement will be rounded down to the nearest integer.
(3) Number of entries per page	(2) Page size requirement / ((1) entry length + 4 (*1)) *1: Pointer length - Result of (3) number of entries per page will be rounded down to the nearest integer.
(4) Number of pages required for storing indexes	Total number of records / (3) number of entries per page - Result of (4) number of pages required for storing indexes will be rounded up to the nearest integer.
(5) Space requirement	(4) Number of pages required for storing indexes x 8192 (*1) / usage rate (*2) *1: Page length *2: Specify 0.7 or lower.

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

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

E.3 Sizes of Data Types

This section lists the sizes of the data types.

E.3.1 Sizes of Fixed-Length Data Types

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

Data type	Size (bytes)
SMALLINT (INT2)	2
INTEGER (INT4)	4
BIGINT (INT8)	8
REAL	4
DOUBLE PRECISION	8
SERIAL (SERIAL4)	4
BIGSERIAL (SERIAL8)	8
MONEY	8
FLOAT	8
FLOAT (1-24)	4
FLOAT (25-53)	8
TIMESTAMP WITHOUT TIME ZONE	8
TIMESTAMP WITH TIME ZONE	8
DATE	4
TIME WITHOUT TIME ZONE	8
TIME WITH TIME ZONE	12

Data type	Size (bytes)
INTERVAL	12
BOOLEAN	1
CIDR	IPv4: 7 IPv6: 19
INET	IPv4: 7 IPv6: 19
MACADDR	6
MACADDR8	8
POINT	16
LINE	32
LSEG	32
BOX	32
CIRCLE	24

E.3.2 Sizes of Variable-Length Data Types

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

Data type	Size (bytes)	Remarks
path	Length of size portion + 12 + 16 x number of vertices	1) When carrying out division, round to the next integer. 2) If the real data length is less than 127, then the length of the size portion is 1 byte, otherwise it is 4 bytes. 3) The number of bytes per character depends on the character set (refer to "E.3.4 Number of Bytes per Character" for details).
polygon	Length of size portion + 36 + 16 x number of vertices	
decimal	Length of size portion + 2 + (integer precision / 4 + decimal precision / 4) x 2	
numeric		
bytea	Length of size portion + real data length	
character varying(<i>n</i>), varchar(<i>n</i>)	Length of size portion + number of characters x number of bytes per character	
character(<i>n</i>), char(<i>n</i>)	Length of size portion + <i>n</i> x number of bytes per character	
text	Length of size portion + number of characters x number of bytes per character	

E.3.3 Sizes of Array Data Types

The following table lists the sizes of array data types.

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

E.3.4 Number of Bytes per Character

The following table lists the number of bytes per character.

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

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

E.4 Estimating Transaction Log Space Requirements

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

```
Transaction log space requirements = max_wal_size
```

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

E.5 Estimating Archive Log Space Requirements

This section explains how to estimate archive log space requirements.

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

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

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

E.6 Estimating Backup Disk Space Requirements

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

```
Backup disk space requirements = size of the database cluster x 2 + transaction log space requirements  
+ archive log space requirements
```



Note

If the `pgx_dmpall` command performs a backup using a user exit, the backup disk size differs according to the database resources targeted for backup and the copy method.

E.7 Estimating VCI Disk Space Requirements

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

```
Disk space = (number of rows in tables) x (number of bytes per row) x (compression ratio) + (WOS size)
```

Number of bytes per row

$$\text{Number of bytes per row} = (19 + (\text{number of columns specified in CREATE INDEX}) / 8 + (\text{number of bytes per single column value})) \times 1.1$$

Note: Round up the result to the nearest integer.

Compression ratio

Specify a value between 0 and 1. Since compression ratio depends on the data being compressed, use actual data or test data that simulates it, then compare the value with the estimation result. As a guide, the compression ratio measured with the Fujitsu sample data is shown below:

- Data with high degree of randomness (difficult to compress): Up to approximately 0.9 times.
- Data with high degree of similarity (easy to compress): Up to approximately 0.5 times.

WOS size

$$\text{WOS size} = (\text{number of WOS rows}) / 185 \times 8096$$

One row is added to the number of WOS rows for each INSERT and DELETE, and two rows are added for UPDATE. On the other hand, the number decreases to 520,000 rows or less during conversion to ROS performed by the ROS control daemon.



Note

.....

VCI does not support retrieval of disk space usage using the database object size function `pg_indexes_size`. To find out the actual total VCI disk space, check the disk space of the storage directory using an OS command or other method.

.....

Appendix F Estimating Memory Requirements

This appendix explains how to estimate the memory.

F.1 FUJITSU Enterprise Postgres Memory Requirements

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

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

```
fujitsuEnterprisePostgresRequiredMemory = sharedMemoryAmount + localMemoryAmount
```

Shared memory amount

Refer to "Shared Memory and Semaphores" under "Server Administration" in the PostgreSQL Documentation for information on shared memory. If you enable the Global Meta Cache feature, you must also add the value of `pgx_global_metacache`. Refer to "Parameters" in the Operation Guide for the setting values.

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

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

Local memory amount

```
localMemoryAmount = processStackArea  
                    + memoryUsedInDbSessionsThatUseTempTables  
                    + memoryUsedInDbSessionsThatPerformSortAndHashTableOperations  
                    + memoryUsedInMaintenanceOperations  
                    + baseMemoryUsedInEachProcess  
                    + memoryUsedPreparingForDataAccess
```

Process stack area

```
processStackArea  
= max_stack_depth x (max_connections + autovacuum_max_workers + 9)
```

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

```
memoryUsedInDbSessionsThatUseTempTables  
= temp_buffers x max_connections
```

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

```
memoryUsedInDbSessionsThatPerformSortAndHashTableOperations  
= work_mem (*1) x max_connections
```


*1) For hash table operations, multiply work_mem by hash_mem_multiplier.

This formula evaluates to the maximum value.

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

Memory used in maintenance operations

```
memoryUsedInMaintenanceOperations
= maintenance_work_mem x (numOfSessionsPerformingMaintenance + autovacuum_max_workers)
```

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

Base memory used in each process

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

Use the result of the following formula for memory consumed per process. This formula evaluates to the memory used when server processes are running.

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

The amount of memory consumed per process is determined by the number of tables, indexes, and all columns of all tables that the process accesses. If your system has about 100 tables, you can estimate it to be 3 MB, but otherwise use the following estimate:

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

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

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

*1) Safety Factor (1.5)

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

Memory used preparing for data access

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

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

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

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

F.2 Database Multiplexing Memory Requirements

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

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

```
Memory usage of the database multiplexing feature for the database server
= Peak memory usage of the Mirroring Controller processes
+ Peak memory usage of the Mirroring Controller commands

Peak memory usage of the Mirroring Controller processes=150 MB

Peak memory usage of the Mirroring Controller commands=50 MB x Number of commands executed
simultaneously
```

F.3 VCI Memory Requirements

This section describes the formula for estimating VCI memory requirements.

Use the following formula to obtain a rough estimate of memory requirements:

```
memUsedByVci = memForData + memForEachProcess
```

Memory required to store data in memory

Secure the space estimated using the formula below on the stable buffer (part of shared_buffers).

```
memForData = (numOfRowsInTables) x (numOfBytesPerRow) + (wosSize)
```

Number of bytes per row

```
numOfBytesPerRow  
= (19 + (numOfColsInCreateIndexStatement) / 8 + (numOfBytesPerSingleColValue)) x 1.1
```

Note: Round up the result to the nearest integer.

WOS size

```
wosSize = (numOfWosRows) / 185 x 8096
```

One row is added to the number of WOS rows for each INSERT and DELETE, and two rows are added for UPDATE. On the other hand, the number decreases to 520,000 rows or less during conversion to ROS performed by the ROS control daemon.

Memory required for each process

```
memForEachProcess  
= memUsedPerScanning  
+ memUsedForVciMaintenance  
+ memUsedByCreateIndexStatement
```

Memory used per scanning

- Parallel scan

```
memUsedPerScanning  
= vci.shared_work_mem + (numOfParallelWorkers + 1) x vci.maintenance_work_mem
```

Note: The number of parallel workers used by VCI simultaneously in the entire instance is equal to or less than vci.max_parallel_degree.

- Non-parallel scan

```
memUsedPerScanning = vci.max_local_ros + vci.maintenance_work_mem
```



Note

- vci.shared_work_mem, and vci.max_local_ros are used to create local ROS. If local ROS exceeds these sizes, execute a query without using VCI according to the conventional plan.
- vci.maintenance_work_mem specifies the memory size to be secured dynamically. If it exceeds the specified value, a disk temporary file is used for operation.

Memory used for VCI maintenance

```
memUsedForVciMaintenance = vci.maintenance_work_mem x vci.control_max_workers
```

Memory used by CREATE INDEX

```
memUsedByCreateIndexStatement = vci.maintenance_work_mem
```



vci.maintenance_work_mem specifies the memory to be secured dynamically. If it exceeds the specified value, a disk temporary file is used for operation.

F.4 High-Speed Data Load Memory Requirements

This section describes the formula for estimating memory requirements for the high-speed data load feature.

Use the following formula to obtain a rough estimate of memory requirements:

```
Memory usage of high speed data load
= (Peak memory usage of pgx_loader processes + Peak memory usage of the pgx_loader commands)
x Number of commands executed simultaneously

Peak memory usage of pgx_loader processes
= Peak memory usage of the backend process      (6 MB)
+ Peak memory usage of parallel workers          (6 MB x number of parallel workers)
+ Peak memory usage of dynamic shared memory (80 MB x number of parallel workers)

Peak memory usage of the pgx_loader commands=9 MB
```



In addition to the size calculated using the formula above, the database cache on the shared memory estimated using the shared_buffers parameter is consumed according to the size of the data (table and index keys) loaded using this feature. Refer to ["E.1 Estimating Table Size Requirements"](#) and ["E.2 Estimating Index Size Requirements"](#) for information on estimating an appropriate shared buffers value.

F.5 Global Meta Cache Memory Requirements

This section describes the formula for estimating Global Meta Cache memory requirements.

The memory calculated by "Size of the GMC area" is allocated to the shared memory. The memory calculated by the per-process meta cache management information is allocated to the local memory. Refer to the graphic in "Architecture of Global Meta Cache Feature" in the "Memory usage reduction by Global Meta Cache" in the General Description for more information.

Use the following formula to obtain a rough estimate of memory requirements:

```
Amount of memory used by the Global Meta Cache feature
= Size of GMC area + Per-process meta cache management information

Size of GMC area = (All user tables x 0.4 KB
+ All user indexes x 0.3 KB
+ All user columns x 0.8 KB) x 1.5 (*1)

Per-process meta cache management information
= (All user tables + All user indexes + All user columns) x 0.1KB x max_connections x 1.5 (*1)
```

*1) Safety Factor (1.5)

This value takes into account the case where both GMC before and after the change temporarily exist at the same time in shared memory when the table definition is changed or the row of the system catalog is changed.

Appendix G Quantitative Limits

This appendix lists the quantitative limits of FUJITSU Enterprise Postgres.

Table G.1 Length of identifier

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

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

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

Table G.2 Database object

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

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

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

Table G.3 Schema element

Item	Limit
Number of columns that can be defined in one table	From 250 to 1600 (according to the data type)
Table row length	Up to 400 gigabytes
Number of columns comprising a unique constraint	Up to 32 columns
Data length comprising a unique constraint	Less than 2,000 bytes (*1) (*2)

Item	Limit
Table size	Up to 32 terabyte
Search condition character string length in a trigger definition statement	Up to 800 megabytes (*1) (*2)
Item size	Up to 1 gigabyte

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

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

Table G.4 Index

Item	Limit
Number of columns comprising a key (including VCI)	Up to 32 columns
Key length (other than VCI)	Less than 2,000 bytes (*1)

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

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

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

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

Table G.6 Function definition

Item	Limit
Number of arguments that can be specified	Up to 100
Number of variable names that can be specified in the declarations section	No limit
Number of SQL statements or control statements that can be specified in a function processing implementation	No limit

Table G.7 Data operation statement

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

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

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

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

Table G.8 Data size

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

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

Appendix H 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>
SEMMNI	Fourth parameter of kernel.sem	Specify <i>currentValue</i> + <i>calculatedValue</i>
SEMMNS	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	
SEVMX	No compatible parameter	

Remark 1: kernel.shmall specifies the number of pages.

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.

Appendix I Determining the Preferred WebAdmin Configuration

This appendix describes the two different configurations in which WebAdmin can be used and how to select the most suitable configuration.

I.1 WebAdmin Configurations

WebAdmin can be installed in two configurations:


- Single-server
- Multiserver

WebAdmin does not support encrypted communication between browser and server or between servers. Therefore, when using WebAdmin in either configuration, build the communication path with the browser or each server on a network that cannot be accessed externally.

I.1.1 Single-Server Configuration

A single-server configuration enables you to create and operate instances on a single server. In this configuration, WebAdmin must be installed on the same database server as the FUJITSU Enterprise Postgres Server component.


Single-server configuration




I.1.2 Multiserver Configuration

A multiserver configuration enables you to create and operate instances stored on multiple database servers. As shown in the figure below, WebAdmin can be installed on a dedicated WebAdmin server and used to collectively manage the instances stored on the database servers.

Multiserver configuration



Also, when setting up the arbitration server by WebAdmin during database multiplexing mode, install WebAdmin on the arbitration server.



I.2 Installing WebAdmin in a Single-Server Configuration

To install WebAdmin in a single-server configuration, the FUJITSU Enterprise Postgres Server component and WebAdmin must be installed on the same machine.

Select the following items when installing FUJITSU Enterprise Postgres in a single-server configuration:

- FUJITSU Enterprise Postgres Advanced Edition
- WebAdmin

I.3 Installing WebAdmin in a Multiserver Configuration

In a multiserver configuration, install WebAdmin on one server, and both WebAdmin and the FUJITSU Enterprise Postgres Server component on any number of database servers.

Select the following items when installing FUJITSU Enterprise Postgres in a multiserver configuration:

- WebAdmin server:
 - WebAdmin
- Database server:
 - FUJITSU Enterprise Postgres Advanced Edition
 - WebAdmin

Also, when setting up the arbitration server by WebAdmin during database multiplexing mode, select the following when installing FUJITSU Enterprise Postgres.

- Arbitration server
 - FUJITSU Enterprise Postgres Server Assistant
 - WebAdmin



See

.....
Refer to the Installation and Setup Guide for Server Assistant for details on how to install the Server Assistant.
.....

Appendix J System Configuration when using Pgpool-II

Describes the system configuration when using Pgpool-II.

The system configuration when using Pgpool-II is as follows:

Place on database server

System configuration to coexist the database server with Pgpool-II.

Place on application server

System configuration to coexist the application server with Pgpool-II.

Place on dedicated server

System configuration in which Pgpool-II resides on a dedicated server (Pgpool-II Server) that is separate from the database and application servers.

Select the system configuration that best meets your operational requirements.

J.1 Pgpool-II Configuration

In this example, Pgpool-II is deployed on a different Pgpool-II server than the database and application servers.

There are three configurations of Pgpool-II:

- Single-machine configuration
- Two-machine configuration
- Three-machine configuration


Although the Pgpool-II server can be operated on a single machine, to ensure business continuity, it is recommended to operate the Pgpool-II server using a three-machine configuration in FUJITSU Enterprise Postgres.

If employing a configuration of three or more machines, use an odd number of machines in the configuration.

J.1.1 Single-Machine Configuration

This is the basic configuration when running Pgpool-II.


Although the database server has redundancy, if an error occurs on the Pgpool-II server that accesses the database server, the job will stop.



J.1.2 Two-Machine Configuration

When an error occurs on the active server, the Pgpool-II monitoring feature that mutually monitors the status of the Pgpool-II servers enables jobs to continue uninterrupted by switching to the standby server.

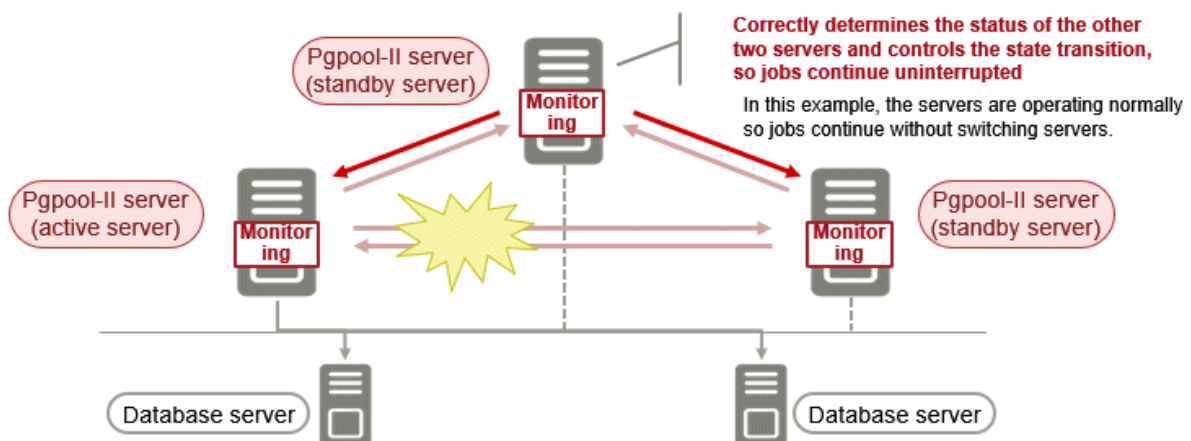
If the network between Pgpool-II servers is disconnected, even if the Pgpool-II servers are running correctly, which may lead to stoppage of jobs.



J.1.3 Three-Machine Configuration

The Pgpool-II monitoring feature enables a Pgpool-II server to monitor the other two Pgpool-II servers.

Even if any of the networks monitoring the Pgpool-II servers are disconnected, the status of servers on a network that is operating normally can be checked correctly, enabling accurate continuation of jobs.



J.2 Installing Pgpool-II

Pgpool-II is bundled with the server program and the client program. To use Pgpool-II, use the server program or the client program to install and set up Pgpool-II.

Depending on where Pgpool-II is installed, select the appropriate DVD for deployment:

Installing on Database Server (coexist)

Install the Pgpool-II program along with the server program from the server program DVD.

Installing on Application Server (coexist)

Install the Pgpool-II program along with the client program from the client program DVD.

Installing on Dedicated server different from the above (Pgpool-II server)

Install the Pgpool-II program along with the client program from the client program DVD.

J.3 Pgpool-II Setup

Describes how to set up Pgpool-II.

J.3.1 Setting Environment Variables

If you use the Pgpool-II command, set the following environment variables:

PATH environment variable

Add "Install Directory/bin".

LD_LIBRARY_PATH environment variable

Add "Install Directory/bin".

The following is an example of setting environment variables:

Example

The following is an example of setting environment variables when the installation directory is "/opt/fsepv <x> pgpool-II".

"<x>" indicates the product version.

```
$ PATH=/opt/fsepv<x>pgpool-II/bin:$PATH ; export PATH
$ LD_LIBRARY_PATH=/opt/fsepv<x>pgpool-II/lib:$LD_LIBRARY_PATH ; export LD_LIBRARY_PATH
```

J.3.2 Configuration file

Describes Pgpool-II configuration files.

J.3.2.1 Configuring pgpool.conf

To configure pgpool.conf, see the Pgpool-II documentation.

A sample configuration file is located under the installation directory/etc.

J.3.2.2 Using Configuration Files

The pgpool command makes use of configuration files such as pgpool.conf, pcp.conf, and pool_hba.conf.

To take advantage of these configuration files, specify the path to the files in the pgpool command options.

The following example shows how to configure options for the pgpool command:

Example

```
$ pgpool -f /usr/local/etc/pgpool.conf -F /usr/local/etc/pcp.conf -a /usr/local/etc/pool_hba.conf
```

Appendix K Supported contrib Modules and Extensions Provided by External Projects

FUJITSU Enterprise Postgres supports PostgreSQL contrib modules, and extensions provided by external projects.

Refer to the following for details on the supported contrib modules:

- "Additional Supplied Modules" in the PostgreSQL Documentation
- "Additional Supplied Programs" in the PostgreSQL Documentation



Information

You can also check the list of available extensions using the `pg_available_extensions` view.

Refer to "OSS Supported by FUJITSU Enterprise Postgres" in the General Description for information on supported extensions provided by external projects.

Appendix L Procedure when Modifying the JRE Installation

This appendix describes the procedures to follow when modifying the JRE installation.

The JRE is used by features such as WebAdmin and database multiplexing.

Therefore, when updating or reinstalling JRE after installing FUJITSU Enterprise Postgres, the procedures below must be performed.

L.1 When Using WebAdmin

WebAdmin must be set up again.

Follow the procedure below to modify the JRE installation:

1. Stop the Web server feature of WebAdmin

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

2. Remove WebAdmin

Refer to "[B.2 Removing WebAdmin](#)" for details.

3. Modify the JRE installation

4. Set the JAVA_HOME environment variable

Set the JAVA_HOME environment variable to the installation destination of Open JRE 8.

Example

```
# export JAVA_HOME="OpenJRE8InstallDir "
```

5. Set up WebAdmin

Refer to "[B.1.1 Setting Up WebAdmin](#)" for details.

6. Start the Web server feature of WebAdmin

Refer to "[B.1.2 Starting the Web Server Feature of WebAdmin](#)" for details.

L.2 When Performing Database Multiplexing

Mirroring Controller must be restarted.

Follow the procedure below to modify the JRE installation:

1. Stop Mirroring Controller

Refer to the Cluster Operation Guide (Database Multiplexing) for details.

2. Modify the JRE installation

3. Change the installation environment to be used by Mirroring Controller



Note

If database multiplexing is performed using WebAdmin, perform the procedure described in this procedure after performing step 4 "Set the JAVA_HOME environment variable" in "[L.1 When Using WebAdmin](#)".

Set the JAVA_HOME environment variable to the installation destination of Open JRE 8, and use the mc_update_jre_env command to change the installation environment to be used by Mirroring Controller.

This procedure must be executed by the superuser.

Example

/opt/fsepv<x>server64/bin is the installation directory where the server product is installed.

```
$ su -  
Password:*****  
# export JAVA_HOME="OpenJRE8InstallDir "  
# /opt/fsepv<x>server64/bin/mc_update_jre_env
```

4. Start Mirroring Controller

Refer to the Cluster Operation Guide (Database Multiplexing) for details.

Index

[C]		[T]	
Changing client authentication information.....	22	TCP/IP Protocol.....	6
Changing Instance Settings.....	21		
Check the disk space.....	8	[U]	
Client Authentication Information settings.....	28	Uninstallation.....	1,41
Creating an Instance.....	19,25	Uninstallation in Interactive Mode.....	41
Creating an Instance Administrator.....	13	Using the initdb Command.....	24
Creating Instances.....	17	Using WebAdmin.....	17
[D]		[W]	
Disk Space Required for Installation.....	5	WebAdmin automatic start.....	45
		Web server port number.....	44
[E]		When an Instance was Created with the initdb Command.....	28
Editing instance information.....	22	When an Instance was Created with WebAdmin.....	27
Excluded Software.....	5	When Performing Database Multiplexing.....	77
		When Using WebAdmin.....	77
[H]			
Hardware Environment.....	5		
How to Set Up the Pop-up Blocker.....	43		
[I]			
Importing Instances.....	23		
Installation.....	8		
Installation Types.....	1		
Instance configuration.....	22		
[L]			
Logging in to WebAdmin.....	18		
[N]			
New Installation.....	1		
[O]			
Operating Environment.....	2		
Operating Method Types and Selection.....	12		
[P]			
Port number to use when Tomcat is stopped.....	45		
postgresql.conf.....	50		
Pre-installation Tasks.....	8		
Preparations for Setup.....	13		
Procedure when Modifying the JRE Installation.....	77		
[R]			
Recommended Browser Settings.....	43		
Reinstallation.....	1		
Related Software.....	5		
Removing WebAdmin.....	46		
Required Operating System.....	2		
Required Patches.....	5		
[S]			
Settings related to connection.....	27		
Starting the Web Server Feature of WebAdmin.....	45		
Startup URL for WebAdmin.....	18		
Stopping the Web Server Feature of WebAdmin.....	45		
Supported contrib Modules and Extensions Provided by External Projects.....	76		