

Fujitsu Enterprise Postgres 16

Application Development Guide

Windows/Linux

J2UL-2958-01ENZ0(00) July 2024

Preface

Purpose of this document

This is a guide for the developers of Fujitsu Enterprise Postgres applications.

Intended readers

This document is intended for developers of applications that use Fujitsu Enterprise Postgres. Of the interfaces provided by Fujitsu Enterprise Postgres, this guide describes the PostgreSQL extended interface.

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

- L
- PostgreSQL
- SQL
- Linux
- w
- PostgreSQL
- SQL
- Windows

Structure of this document

This document is structured as follows:

Chapter 1 Overview of the Application Development Function

Provides an overview of Fujitsu Enterprise Postgres application development.

Chapter 2 JDBC Driver

Explains how to use JDBC drivers.

Chapter 3 ODBC Driver

Explains how to use ODBC drivers.

Chapter 4 .NET Data Provider

Explains how to use .NET Data Provider.

Chapter 5 C Library (libpq)

Explains how to use C applications.

Chapter 6 Embedded SQL in C

Explains how to use embedded SQL in C.

Chapter 7 Embedded SQL in COBOL

Explains how to use embedded SQL in COBOL.

Chapter 8 SQL References

Explains the SQL statements which were extended in Fujitsu Enterprise Postgres development.

Chapter 9 Compatibility with Oracle Databases

Explains features that are compatible with Oracle databases.

Chapter 10 Application Connection Switch Feature

Explains the application connection switch feature.

Chapter 11 Scan Using a Vertical Clustered Index (VCI)

Explains how to perform scan using a Vertical Clustered Index (VCI).

Appendix A Precautions when Developing Applications

Provides some points to note about application development.

Appendix B Conversion Procedures Required due to Differences from Oracle Database

Explains how to convert from an Oracle database to Fujitsu Enterprise Postgres, within the scope noted in "Compatibility with Oracle Databases" from the following perspectives.

Appendix C Tables Used by the Features Compatible with Oracle Databases

Explains the tables used by the features compatible with Oracle databases.

Appendix D ECOBPG - Embedded SQL in COBOL

Explains application development using embedded SQL in COBOL.

Appendix E Quantitative Limits

This appendix explains limitations.

Appendix F Reference

Provides a reference for each interface.

Export restrictions

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

Issue date and version

Edition 1.0: July 2024

Copyright

Copyright 2015-2024 Fujitsu Limited

Contents

Chapter 1 Overview of the Application Development Function	
1.1 Support for National Characters	2
1.1.1 Literal	2
1.1.2 Data Type	2
1.1.3 Functions and Operator	3
1.2 Compatibility with Oracle Database	3
1.3 Application Connection Switch Feature	3
1.3.1 Integration with Database Multiplexing	3
1.4 Notes on Application Compatibility	
1.4.1 Checking Execution Results	2
1.4.2 Referencing System Catalogs	4
1.4.3 Using Functions.	4
Chapter 2 JDBC Driver	ú
2.1 Development Environment	
2.1.1 Combining with JDK or JRE	
2.2 Setup	
2.2.1 Environment Settings.	
2.2.2 Message Language and Encoding System Used by Applications Settings	
2.2.3 Settings for Encrypting Communication Data	
2.3 Connecting to the Database	
2.3.1 Using the DriverManager Class.	
2.3.2 Using the PGConnectionPoolDataSource Class	
2.3.3 Using the PGXADataSource Class.	
2.4 Application Development	
2.4.1 Relationship between the Application Data Types and Database Data Types	
2.4.2 Statement Caching Feature	
2.4.3.1 Errors when an Application Connection Switch Occurs and Corresponding Actions	13
	13
2.4.3.1 Errors when an Application Connection Switch Occurs and Corresponding Actions. Chapter 3 ODBC Driver. 3.1 Development Environment.	13
2.4.3.1 Errors when an Application Connection Switch Occurs and Corresponding Actions	13
2.4.3.1 Errors when an Application Connection Switch Occurs and Corresponding Actions. Chapter 3 ODBC Driver. 3.1 Development Environment.	131414
2.4.3.1 Errors when an Application Connection Switch Occurs and Corresponding Actions. Chapter 3 ODBC Driver. 3.1 Development Environment. 3.2 Setup.	1314141414
2.4.3.1 Errors when an Application Connection Switch Occurs and Corresponding Actions. Chapter 3 ODBC Driver. 3.1 Development Environment. 3.2 Setup 3.2.1 Registering ODBC Drivers	
2.4.3.1 Errors when an Application Connection Switch Occurs and Corresponding Actions. Chapter 3 ODBC Driver	12141414141610
2.4.3.1 Errors when an Application Connection Switch Occurs and Corresponding Actions. Chapter 3 ODBC Driver	12141414141610
2.4.3.1 Errors when an Application Connection Switch Occurs and Corresponding Actions. Chapter 3 ODBC Driver	
2.4.3.1 Errors when an Application Connection Switch Occurs and Corresponding Actions. Chapter 3 ODBC Driver. 3.1 Development Environment. 3.2 Setup 3.2.1 Registering ODBC Drivers. 3.2.2 Registering ODBC Data Sources(for Windows(R)). 3.2.2.1 Registering Using GUI. 3.2.2.2 Registering Using Commands. 3.2.3 Registering ODBC Data Sources(for Linux).	
2.4.3.1 Errors when an Application Connection Switch Occurs and Corresponding Actions. Chapter 3 ODBC Driver. 3.1 Development Environment. 3.2 Setup 3.2.1 Registering ODBC Drivers 3.2.2 Registering ODBC Data Sources(for Windows(R)). 3.2.2.1 Registering Using GUI 3.2.2.2 Registering Using Commands 3.2.3 Registering ODBC Data Sources(for Linux) 3.2.4 Message Language and Encoding System Used by Applications Settings	
2.4.3.1 Errors when an Application Connection Switch Occurs and Corresponding Actions. Chapter 3 ODBC Driver	
2.4.3.1 Errors when an Application Connection Switch Occurs and Corresponding Actions. Chapter 3 ODBC Driver	
2.4.3.1 Errors when an Application Connection Switch Occurs and Corresponding Actions. Chapter 3 ODBC Driver	
2.4.3.1 Errors when an Application Connection Switch Occurs and Corresponding Actions. Chapter 3 ODBC Driver. 3.1 Development Environment. 3.2 Setup 3.2.1 Registering ODBC Drivers. 3.2.2 Registering ODBC Data Sources(for Windows(R)). 3.2.2.1 Registering Using GUI 3.2.2.2 Registering Using Commands. 3.2.3 Registering Using Commands. 3.2.4 Message Language and Encoding System Used by Applications Settings. 3.3 Connecting to the Database 3.4 Application Development 3.4.1 Compiling Applications (for Windows (R)). 3.4.2 Compiling Applications (for Linux)	
2.4.3.1 Errors when an Application Connection Switch Occurs and Corresponding Actions. Chapter 3 ODBC Driver. 3.1 Development Environment 3.2 Setup 3.2.1 Registering ODBC Drivers 3.2.2 Registering ODBC Data Sources(for Windows(R)) 3.2.2.1 Registering Using GUI 3.2.2.2 Registering Using Commands 3.2.3 Registering ODBC Data Sources(for Linux) 3.2.4 Message Language and Encoding System Used by Applications Settings 3.3 Connecting to the Database 3.4 Application Development 3.4.1 Compiling Applications (for Windows (R)) 3.4.2 Compiling Applications (for Linux) 3.4.3 Creating Applications While in Database Multiplexing Mode 3.4.3.1 Errors when an Application Connection Switch Occurs and Corresponding Actions.	
2.4.3.1 Errors when an Application Connection Switch Occurs and Corresponding Actions. Chapter 3 ODBC Driver	
2.4.3.1 Errors when an Application Connection Switch Occurs and Corresponding Actions Chapter 3 ODBC Driver 3.1 Development Environment 3.2 Setup 3.2.1 Registering ODBC Drivers 3.2.2 Registering ODBC Data Sources(for Windows(R)) 3.2.2.1 Registering Using GUI 3.2.2.2 Registering Using GUI 3.2.3.2 Registering ODBC Data Sources(for Linux) 3.2.4 Message Language and Encoding System Used by Applications Settings 3.3 Connecting to the Database 3.4 Application Development 3.4.1 Compiling Applications (for Windows (R)) 3.4.2 Compiling Applications (for Linux) 3.4.3 Creating Applications While in Database Multiplexing Mode 3.4.3.1 Errors when an Application Connection Switch Occurs and Corresponding Actions. Chapter 4 .NET Data Provider 4.1 Development Environment	15 14 14 14 16 16 18 18 19 19 19 19 19 19 19 19 19 19 19 19 19
2.4.3.1 Errors when an Application Connection Switch Occurs and Corresponding Actions. Chapter 3 ODBC Driver. 3.1 Development Environment. 3.2 Setup 3.2.1 Registering ODBC Drivers. 3.2.2 Registering ODBC Data Sources(for Windows(R)). 3.2.2.1 Registering Using GUI. 3.2.2.2 Registering Using Commands. 3.2.3 Registering ODBC Data Sources(for Linux). 3.2.4 Message Language and Encoding System Used by Applications Settings. 3.3 Connecting to the Database. 3.4 Application Development. 3.4.1 Compiling Applications (for Windows (R)). 3.4.2 Compiling Applications (for Linux). 3.4.3 Creating Applications While in Database Multiplexing Mode. 3.4.3.1 Errors when an Application Connection Switch Occurs and Corresponding Actions. Chapter 4 .NET Data Provider. 4.1 Development Environment. 4.2 Setup	
2.4.3.1 Errors when an Application Connection Switch Occurs and Corresponding Actions. Chapter 3 ODBC Driver	15 14 14 14 16 18 22 24 24 25 26 27 27 27
2.4.3.1 Errors when an Application Connection Switch Occurs and Corresponding Actions	15 14 14 14 16 16 18 22 24 24 22 25 26 27 27 27 27 28
2.4.3.1 Errors when an Application Connection Switch Occurs and Corresponding Actions	15 14 14 14 16 16 18 22 25 24 24 25 26 27 27 27 27 28
2.4.3.1 Errors when an Application Connection Switch Occurs and Corresponding Actions. Chapter 3 ODBC Driver	15 14 14 14 16 16 18 18 19 19 19 19 19 19 19 19 19 19 19 19 19
2.4.3.1 Errors when an Application Connection Switch Occurs and Corresponding Actions	15
2.4.3.1 Errors when an Application Connection Switch Occurs and Corresponding Actions. Chapter 3 ODBC Driver	

4.3.3 Connection String	32
4.4 Application Development	36
4.4.1 Data Types	36
4.4.2 Relationship between Application Data Types and Database Data Types	
4.4.3 Creating Applications while in Database Multiplexing Mode	41
4.4.3.1 Errors when an Application Connection Switch Occurs and Corresponding Actions	
4.4.4 Notes	42
4.5 Uninstallation	46
4.5.1 Uninstalling Npgsql	46
4.5.2 Uninstalling .NET Data Provider Type Plugins	47
4.5.3 Uninstalling Entity Framework Core	
4.5.4 Uninstalling Npgsql.OpenTelemetry	47
Chapter 5 C Library (libpq)	48
5.1 Development Environment	
5.2 Setup	
5.2.1 Environment Settings	
5.2.2 Message Language and Encoding System Used by Applications Settings	
5.2.3 Settings for Encrypting Communication Data	
5.3 Connecting with the Database	
5.4 Application Development	
5.4.1 Compiling Applications	
5.4.2 Creating Applications while in Database Multiplexing Mode	
5.4.2.1 Errors when an Application Connection Switch Occurs and Corresponding Actions	52
Chapter 6 Embedded SQL in C	53
6.1 Development Environment	53
6.2 Setup	53
6.2.1 Environment Settings.	53
6.2.2 Message Language and Encoding System Used by Applications Settings	53
6.2.3 Settings for Encrypting Communication Data	
6.3 Connecting with the Database	
6.4 Application Development	56
6.4.1 Support for National Character Data Types	
6.4.2 Compiling Applications	
6.4.3 Bulk INSERT	59
6.4.4 Creating Applications while in Database Multiplexing Mode	
6.4.4.1 Errors when an Application Connection Switch Occurs and Corresponding Actions	
6.4.5 Notes	63
Chapter 7 Embedded SQL in COBOL	
7.1 Development Environment	
7.2 Setup	
7.2.1 Environment Settings	
7.2.2 Message Language and Encoding System Used by Applications	
7.2.3 Settings for Encrypting Communication Data	
7.3 Connecting with the Database	
7.4 Application Development	
7.4.1 Support for National Character Data Types	
7.4.2 Compiling Applications	
7.4.3 Bulk INSERT	
7.4.4 DECLARE STATEMENT	
7.4.5 Creating Applications while in Database Multiplexing Mode	
7.4.5.1 Errors when an Application Connection Switch Occurs and Corresponding Actions	75
Chapter 8 SQL References	
8.1 Expanded Trigger Definition Feature	
8.1.1 CREATE TRIGGER	76

8.1.2 How to Define Triggers in pgAdmin	77
Chapter 9 Compatibility with Oracle Databases	78
9.1 Overview	
9.2 Precautions when Using the Features Compatible with Oracle Databases	78
9.2.1 Notes on search_path	79
9.2.2 Notes on SUBSTR	
9.2.3 Notes when Integrating with the Interface for Application Development	79
9.3 Queries.	
9.3.1 Outer Join Operator (+)	79
9.3.2 DUAL Table	
9.4 SQL Function Reference	82
9.4.1 DECODE	
9.4.2 SUBSTR	
9.4.3 NVL	
9.5 Package Reference	
9.5.1 DBMS_SQL	
9.5.1.1 Description	
9.5.1.2 Example	
Chapter 10 Application Connection Switch Feature	95
10.1 Connection Information for the Application Connection Switch Feature	95
10.2 Using the Application Connection Switch Feature	96
10.2.1 Using the JDBC Driver	96
10.2.2 Using the ODBC Driver	98
10.2.3 Using a .NET Data Provider	100
10.2.4 Using a Connection Service File	
10.2.5 Using the C Library (libpq)	
10.2.6 Using Embedded SQL	
10.2.7 Using the psql Command	106
Chapter 11 Scan Using a Vertical Clustered Index (VCI)	108
11.1 Operating Conditions	
11.2 Usage.	
11.2.1 Designing	
11.2.2 Checking.	
11.2.3 Evaluating	
11.3 Usage Notes	
Appendix A Precautions when Developing Applications	114
A.1 Precautions when Using Functions and Operators	
A.1.1 General rules of Functions and Operators.	
A.1.2 Errors when Developing Applications that Use Functions and/or Operators	
A.2 Notes when Using Temporary Tables	
A.3 Implicit Data Type Conversions.	
A.3.1 Function Argument.	
A.3.2 Operators.	
A.3.3 Storing Values.	
A.4 Notes on Using Index	
A.4.1 SP-GiST Index.	
A.5 Notes on Using Multibyte Characters in Definition Names	
A.6 How to Build and Run an Application that Uses Shared Libraries	
A.6.1 Setting DT_RUNPATH for Applications	
A.6.2 Direct Linking of Indirectly Used Libraries to Applications	
Appendix B Conversion Procedures Required due to Differences from Oracle Data	phase 124
B.1 Outer Join Operator (Perform Outer Join)	
B.1.1 Comparing with the ^= Comparison Operator.	

B.2 DECODE (Compare Values and Return Corresponding Results)	125
B.2.1 Comparing Numeric Data of Character String Types and Numeric Characters	125
B.2.2 Obtaining Comparison Result from more than 50 Conditional Expressions	125
B.2.3 Obtaining Comparison Result from Values with Different Data Types	126
B.3 SUBSTR (Extract a String of the Specified Length from Another String)	127
B.3.1 Specifying a Value Expression with a Data Type Different from the One that can be Specified for Function Argum	ents127
B.3.2 Extracting a String with the Specified Format from a Datetime Type Value	128
B.3.3 Concatenating a String Value with a NULL value.	128
B.4 NVL (Replace NULL)	129
B.4.1 Obtaining Result from Arguments with Different Data Types	129
B.4.2 Operating on Datetime/Numeric, Including Adding Number of Days to a Particular Day	130
B.4.3 Calculating INTERVAL Values, Including Adding Periods to a Date	
B.5 DBMS_OUTPUT (Output Messages)	131
B.5.1 Outputting Messages Such As Process Progress Status	131
B.5.2 Receiving a Return Value from a Procedure (PL/SQL) Block (For GET_LINES)	133
B.5.3 Receiving a Return Value from a Procedure (PL/SQL) Block (For GET_LINE)	135
B.6 UTL_FILE (Perform File Operation)	136
B.6.1 Registering a Directory to Load and Write Text Files	136
B.6.2 Checking File Information.	137
B.6.3 Copying Files	141
B.6.4 Moving/Renaming Files	142
B.7 DBMS_SQL (Execute Dynamic SQL)	143
B.7.1 Searching Using a Cursor.	143
Annondia C Tables Head by the Festives Competible with Oracle Databases	1.10
Appendix C Tables Used by the Features Compatible with Oracle Databases	
C.1 UTL_FILE.UTL_FILE_DIK	149
Appendix D ECOBPG - Embedded SQL in COBOL	150
D.1 Precautions when Using Functions and Operators	150
D.2 Managing Database Connections	151
D.2.1 Connecting to the Database Server	151
D.2.2 Choosing a Connection.	152
D.2.3 Closing a Connection.	153
D.3 Running SQL Commands	153
D.3.1 Executing SQL Statements.	153
D.3.2 Using Cursors.	154
D.3.3 Managing Transactions.	
D.3.4 Prepared Statements	
D.4 Using Host Variables.	155
D.4.1 Overview	
D.4.2 Declare Sections.	155
D.4.3 Retrieving Query Results	156
D.4.4 Type Mapping	
D.4.5 Handling Nonprimitive SQL Data Types	
D.4.6 Indicators.	
D.5 Dynamic SQL	
D.5.1 Executing Statements without a Result Set.	
D.5.2 Executing a Statement with Input Parameters	
D.5.3 Executing a Statement with a Result Set	
D.6 Using Descriptor Areas	
D.6.1 Named SQL Descriptor Areas.	
D.7 Error Handling.	
D.7.1 Setting Callbacks	
D.7.2 sqlca	
D.7.3 SQLSTATE vs. SQLCODE	
D.8 Preprocessor Directives	
D.8.1 Including Files.	
D.8.2 The define and undef Directives.	

D.8.3 itdet, itndet, else, elit, and endit Directives	
D.9 Processing Embedded SQL Programs	176
D.10 Large Objects	177
D.11 Embedded SQL Commands	177
D.11.1 ALLOCATE DESCRIPTOR	177
D.11.2 CONNECT	178
D.11.3 DEALLOCATE DESCRIPTOR	180
D.11.4 DECLARE	181
D.11.5 DESCRIBE	182
D.11.6 DISCONNECT	182
D.11.7 EXECUTE IMMEDIATE	
D.11.8 GET DESCRIPTOR	184
D.11.9 OPEN	186
D.11.10 PREPARE	186
D.11.11 SET AUTOCOMMIT	187
D.11.12 SET CONNECTION	
D.11.13 SET DESCRIPTOR	188
D.11.14 TYPE	189
D.11.15 VAR	190
D.11.16 WHENEVER	191
D.12 PostgreSQL Client Applications	192
D.12.1 ecobpg	192
Appendix E Quantitative Limits	195
Appendix F Reference	200
F.1 JDBC Driver	200
F.2 ODBC Driver	200
F.2.1 List of Supported APIs	200
F.3 .NET Data Provider	203
F.4 C Library (libpq)	203
F.5 Embedded SQL in C	203
Index	204

Chapter 1 Overview of the Application Development Function

The interface for application development provided by Fujitsu Enterprise Postgres is perfectly compatible with PostgreSQL.

Along with the PostgreSQL interface, Fujitsu Enterprise Postgres also provides the following extended interfaces:

- Support for National Characters


In order to secure portability from mainframes and databases of other companies, Fujitsu Enterprise Postgres provides data types that support national characters. The national characters are usable from the client application languages.

Refer to "1.1 Support for National Characters" for details.

- Compatibility with Oracle Databases

Compatibility with Oracle databases is offered. Use of the compatible features means that the revisions to existing applications can be isolated, and migration to open interfaces is made simpler.

Refer to "1.2 Compatibility with Oracle Database" for details.



- Application connection switch feature

The application connection switch feature is provided to enable automatic connection to the target server when there are multiple servers with redundant configurations.

Refer to "1.3 Application Connection Switch Feature" for details.

- Scanning using a Vertical Clustered Index (VCI)

Scans becomes faster during aggregation of many rows by providing the features below:

- Vertical clustered index (VCI)
- In-memory data

This feature can only be used in Advanced Edition.

Refer to "Chapter 11 Scan Using a Vertical Clustered Index (VCI)" for details.

1.1 Support for National Characters

NCHAR type is provided as the data type to deal with national characters.

The NCHAR type can be used with Fujitsu Enterprise Postgres pgAdmin.



- NCHAR can only be used when the character set of the database is UTF-8.
- NCHAR can be used in the places where CHAR can be used (function arguments, etc.).
- For applications handling NCHAR type data in the database, the data format is the same as CHAR type. Therefore, applications handling data in NCHAR type columns can also be used to handle data stored in CHAR type columns.



Note the following in order to cast NCHAR type data as CHAR type.

- When comparing NCHAR type data where the length differs, ASCII spaces are used to fill in the length of the shorter NCHAR type data so that it can be processed as CHAR type data.
- Depending on the character set, the data size may increase by between 1.5 and 2 times.
- Use the AS clause to specify "varchar" as the column alias.

1.1.1 Literal

Syntax

```
{ N \mid n }'[national character [ ...]]'
```

General rules

National character string literals consist of an 'N' or 'n', and the national character is enclosed in single quotation marks ('). Example: N'ABCDEF'

The data type is national character string type.

1.1.2 Data Type

Syntax

```
\{ \text{ NATIONAL CHARACTER } | \text{ NATIONAL CHAR } | \text{ NCHAR } \} [ \text{ VARYING }][(length) ]
```

The data type of the NCHAR type column is as follows:

Data type specification format	Explanation
NATIONAL CHARACTER(n)	National character string with a fixed length of <i>n</i> characters
NATIONAL CHAR(n)	This will be the same as (1) if (n) is omitted.
NCHAR(n)	n is a whole number larger than 0.
NATIONAL CHARACTER VARYING(n)	National character string with a variable length with a maximum of <i>n</i> characters
NATIONAL CHAR VARYING(n) NCHAR VARYING(n)	Any length of national character string can be accepted when this is omitted.

Data type specification format	Explanation
	n is a whole number larger than 0.

General rules

NCHAR is the national character string type data type. The length is the number of characters.

The length of the national character string type is as follows:

- When VARYING is not specified, the length of national character strings is fixed and will be the specified length.
- When VARYING is specified, the length of national character strings will be variable. In this case, the lower limit will be 0 and the upper limit will be the value specified for length.
- NATIONAL CHARACTER, NATIONAL CHAR, and NCHAR each have the same meaning.

When the national character string to be stored is shorter than the declared upper limit, the NCHAR value is filled with spaces, whereas NCHAR VARYING is stored as is.

The upper limit for character storage is approximately 1GB.

1.1.3 Functions and Operator

Comparison operator

When a NCHAR type or NCHAR VARYING type is specified in a comparison operator, comparison is only possible between NCHAR types or NCHAR VARYING types.

String functions and operators

All of the string functions and operators that can be specified by a CHAR type can also be specified by a NCHAR type. The behavior of these string functions and operators is also the same as with CHAR type.

Pattern matching (LIKE, SIMILAR TO regular expression, POSIX regular expression)

The patterns specified when pattern matching with NCHAR types and NCHAR VARYING types specify the percent sign (%) and the underline (_).

The underline (_) means a match with one national character. The percent sign (%) means a match with any number of national characters 0 or over.

1.2 Compatibility with Oracle Database

The following features have been extended in order to enhance compatibility with Oracle databases:

- Query (external join operator (+), DUAL table)
- Function (DECODE, SUBSTR, NVL)
- Built-in package (DBMS_OUTPUT, UTL_FILE, DBMS_SQL)

Refer to "Chapter 9 Compatibility with Oracle Databases" for information on the features compatible with Oracle databases.

1.3 Application Connection Switch Feature

The application connection switch feature enables automatic connection to the target server when there are multiple servers with redundant configurations.

Refer to "Chapter 10 Application Connection Switch Feature" for information on the application connection switch feature.

1.3.1 Integration with Database Multiplexing

The application connection switch feature is provided to enable automatic connection to the appropriate server when there are multiple servers with redundant configurations.



Refer to the Cluster Operation Guide (Database Multiplexing) for information on database multiplexing.

1.4 Notes on Application Compatibility

Fujitsu Enterprise Postgres upgrades contain feature improvements and enhancements that may affect the applications.

Accordingly, note the points below when developing applications, to ensure compatibility after upgrade.

- Checking execution results
- Referencing system catalogs
- Using functions

1.4.1 Checking Execution Results

Refer to SQLSTATE output in messages to check the SQL statements used in applications and the execution results of commands used during development.



See

Refer to Messages for information on the message content and number.

Refer to "PostgreSQL Error Codes" under "Appendixes" in the PostgreSQL Documentation for information on SQLSTATE.

1.4.2 Referencing System Catalogs

System catalogs can be used to obtain information about the Fujitsu Enterprise Postgres system and database objects.

However, system catalogs may change when the Fujitsu Enterprise Postgres version is upgraded. Also, there are many system catalogs that return information that is inherent to Fujitsu Enterprise Postgres.

Accordingly, reference the information schema defined in standard SQL (information_schema) wherever possible. Note also that queries specifying "*" in the selection list must be avoided to prevent columns being added.



See

Refer to "The Information Schema" under "Client Interfaces" in the PostgreSQL Documentation for details.

The system catalog must be referenced to obtain information not found in the information schema. Instead of directly referencing the system catalog in the application, define a view for that purpose. Note, however, that when defining the view, the column name must be clearly specified after the view name.

An example of defining and using a view is shown below.



CREATE VIEW my_tablespace_view(spcname) AS SELECT spcname FROM pg_tablespace; SELECT * FROM my_tablespace_view V1, pg_tables T1 WHERE V1.spcname = T1.tablespace;

If changes are made to a system catalog, the user will be able to take action by simply making changes to the view, without the need to make changes to the application.

The following shows an example of taking action by redefining a view as if no changes were made.

The pg_tablespace system catalog is redefined in response to the column name being changed from spcname to spacename.



DROP VIEW my_tablespace_view;
CREATE VIEW my_tablespace_view(spcname) AS SELECT spacename FROM pg_tablespace;

1.4.3 Using Functions

The default functions provided with Fujitsu Enterprise Postgres enable a variety of operations and manipulations to be performed, and information to be obtained, using SQL statements.

However, it is possible that internal Fujitsu Enterprise Postgres functions, such as those relating to statistical information or for obtaining system-related information, may change as Fujitsu Enterprise Postgres versions are upgraded.

Accordingly, when using these functions, define them as new functions and then use the newly-defined functions in the applications.

An example of defining and using a function is shown below.



CREATE FUNCTION my_func(relid regclass) RETURNS bigint LANGUAGE SQL AS 'SELECT pg_relation_size(relid)';
SELECT my_func(2619);

If changes are made to a function, the user will be able to take action by simply redefining the function, without the need to make changes to the application.

The following shows an example of taking action by redefining a function as if no changes were made.

The pg_relation_size function is redefined after arguments are added.



DROP FUNCTION my_func(regclass);
CREATE FUNCTION my_func(relid regclass) RETURNS bigint LANGUAGE SQL AS 'SELECT pg_relation_size(relid,\$\$main\$\$)';

Chapter 2 JDBC Driver

This section describes how to use JDBC drivers.

2.1 Development Environment

This section describes application development using JDBC drivers and the runtime environment.

2.1.1 Combining with JDK or JRE

Refer to Installation and Setup Guide for Client for information on combining with JDK or JRE where JDBC drivers can operate.

2.2 Setup

This section describes the environment settings required to use JDBC drivers and how to encrypt communication data.

2.2.1 Environment Settings

Configuration of the CLASSPATH environment variable is required as part of the runtime environment for JDBC drivers.

The name of the JDBC driver file is as follows:

postgresql-jdbc42.jar

Same name whether you use JDK 8, JRE 8, JDK 11, JRE 11, JDK 17 or JRE 17 environment.

The examples below show how to set the CLASSPATH environment variable.

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

- Linux
 - Setting example (TC shell)

setenv CLASSPATH /opt/fsepv<x>client64/jdbc/lib/postgresql-jdbc42.jar:\${CLASSPATH}

- Setting example (bash)

 ${\tt CLASSPATH=/opt/fsepv<} x > {\tt client64/jdbc/lib/postgresql-jdbc42.jar: \$CLASSPATH: export CLASSPATH} is {\tt CLASSPATH: export CLASSPATH: export CLASSPATH: export CLASSPATH is {\tt CLASSPATH: export CLASSPATH:$

- W
- Windows (32-bit)
 - Setting example

set CLASSPATH=C:\Program Files\Fujitsu\fsepv<x>client32\JDBC\lib\postgresql-jdbc42.jar;%CLASSPATH%

- Windows (64-bit)
 - Setting example (when Fujitsu Enterprise Postgres Client 32-bit is installed)

set CLASSPATH=C:\Program Files (x86)\Fujitsu\fsepv<x>client32\JDBC\lib\postgresql-jdbc42.jar;%CLASSPATH%

- Setting example (when Fujitsu Enterprise Postgres Client 64-bit is installed)

set CLASSPATH=C:\Program Files\Fujitsu\fsepv<x>client64\JDBC\lib\postgresql-jdbc42.jar;%CLASSPATH%

2.2.2 Message Language and Encoding System Used by Applications Settings

If the JDBC driver is used, it will automatically set the encoding system on the client to UTF-8, so there is no need to configure this.



See

Refer to "Automatic Character Set Conversion Between Server and Client" in "Server Administration" in the PostgreSQL Documentation for information on encoding systems.

Language settings

You must match the language settings for the application runtime environment with the message locale settings of the

Set language in the "user.language" system property.



Example of running a Java command with system property specified

```
java -Duser.language=en TestClass1
```

2.2.3 Settings for Encrypting Communication Data

When using the communication data encryption feature to connect to the database server, set as follows:

Settings for encrypting communication data for connection to the server

This section describes how to create applications for encrypting communication data.

Set the property of the SSL parameter to "true" to encrypt. The default for the SSL parameter is "false".

If ssl is set to "true", sslmode is internally treated as "verify-full".



- Setting example 1

```
String url = "jdbc:postgresql://sv1/test";
Properties props = new Properties();
props.setProperty("user","fsepuser");
props.setProperty("password","secret");
props.setProperty("ssl","true");
props.setProperty("sslfactory","org.postgresql.ssl.DefaultJavaSSLFactory");
Connection conn = DriverManager.getConnection(url, props);
```

- Setting example 2

```
String url = "jdbc:postgresql://svl/test?
user=fsepuser&password=secret&ssl=true&sslfactory=org.postgresql.ssl.DefaultJavaSSLF
actory";
Connection conn = DriverManager.getConnection(url);
```

To prevent spoofing of the database server, you need to use the keytool command included with Java to import the CA certificate to the Java keystore. In addition, specify "org.postgresql.ssl.DefaultJavaSSLFactory" for the sslfactory parameter.

Refer to JDK documentation and the Oracle website for details.



There is no need to set the ssl parameter if the connection string of the DriverManager class is specified, or if the sslmode parameter is specified in the data source, such as when the application connection switch feature is used. If the ssl parameter is set, the value in the sslmode parameter will be enabled.



Refer to "Secure TCP/IP Connections with SSL" in "Server Administration" in the PostgreSQL Documentation for information on encrypting communication data.

.............

.....

2.3 Connecting to the Database

This section explains how to connect to a database.

- Using the DriverManager Class
- Using the PGConnectionPoolDataSource Class
- Using the PGXADataSource Class



Do not specify "V2" for the "protocol Version" of the connection string.

2.3.1 Using the DriverManager Class

To connect to the database using the DriverManager class, first load the JDBC driver, then specify the connection string as a URI in the API of the DriverManager class.

Load the JDBC driver

Specify org.postgresql.Driver.

Connection string

URI connection is performed as follows:

jdbc:postgresql://host:port/database?
user=user&password=password1&loginTimeout=loginTimeout&socketTimeout=socketTimeout

Argument	Description	
host	Specify the host name for the connection destination.	
	The default is "localhost".	
port	Specify the port number for the database server.	
	The default is "27500".	
database	Specify the database name.	
user	Specify the username that will connect with the database.	
	If this is omitted, the username logged into the operating system that is executing the application will be used.	
password	Specify a password when authentication is required.	
loginTimeout	Specify the timeout for connections (in units of seconds).	

Argument	Description
	Specify a value between 0 and 2147483647. There is no limit set if you set 0 or an invalid value.
	An error occurs when a connection cannot be established within the specified time.
socketTimeout	Specify the timeout for communication with the server (in units of seconds).
	Specify a value between 0 and 2147483647. There is no limit set if you set 0 or an invalid value.
	An error occurs when data is not received from the server within the specified time.



Code examples for applications

```
import java.sql.*;
...
Class.forName("org.postgresql.Driver");
String url = "jdbc:postgresql://sv1:27500/mydb?
user=myuser&password=myuser01&loginTimeout=20&socketTimeout=20";
Connection con = DriverManager.getConnection(url);
```

2.3.2 Using the PGConnectionPoolDataSource Class

To connect to databases using data sources, specify the connection information in the properties of the data source.

Method description

Argument	Description	
setServerName	Specify the host name for the connection destination.	
	The default is "localhost".	
setPortNumber	Specify the port number for the database server.	
	The default is "27500".	
setDatabaseName	Specify the database name.	
setUser	Specify the username of the database.	
	By default, the name used will be that of the user on the operating system that is executing the application.	
setPassword	Specify a password for server authentication.	
setLoginTimeout	Specify the timeout for connections (in units of seconds).	
	Specify a value between 0 and 2147483647. There is no limit set if you set 0 or an invalid value.	
	An error occurs when a connection cannot be established within the specified time.	
setSocketTimeout	Specify the timeout for communication with the server (in units of seconds).	
	Specify a value between 0 and 2147483647. There is no limit set if you set 0 or an invalid value.	
	An error occurs when data is not received from the server within the specified time.	



Code examples for applications

```
import java.sql.*;
import org.postgresql.ds.PGConnectionPoolDataSource;
...
PGConnectionPoolDataSource source = new PGConnectionPoolDataSource();
source.setServerName("sv1");
source.setPortNumber(27500);
source.setDatabaseName("mydb");
source.setUser("myuser");
source.setPassword("myuser01");
source.setLoginTimeout(20);
source.setSocketTimeout(20);
...
Connection con = source.getConnection();
```

2.3.3 Using the PGXADataSource Class

To connect to databases using data sources, specify the connection information in the properties of the data source.

Method description

Argument	Description	
setServerName	Specify the host name for the connection destination.	
	The default is "localhost".	
setPortNumber	Specify the port number for the database server.	
	The default is "27500".	
setDatabaseName	Specify the database name.	
setUser	Specify the username that will connect with the database.	
	If this is omitted, the name used will be that of the user on the operating system that is executing the application.	
setPassword	Specify a password when authentication by a password is required.	
setLoginTimeout	Specify the timeout for connections.	
	The units are seconds. Specify a value between 0 and 2147483647. There is no limit set if you set 0 or an invalid value.	
	An error occurs when a connection cannot be established within the specified time.	
setSocketTimeout	Specify the timeout for communication with the server.	
	The units are seconds. Specify a value between 0 and 2147483647. There is no limit set if you set 0 or an invalid value.	
	An error occurs when data is not received from the server within the specified time.	



Code examples for applications

```
import java.sql.*;
import org.postgresql.xa.PGXADataSource;
...
PGXADataSource source = new PGXADataSource();
source.setServerName("sv1");
source.setPortNumber(27500);
source.setDatabaseName("mydb");
source.setUser("myuser");
source.setPassword("myuser01");
source.setLoginTimeout(20);
```

2.4 Application Development

This section describes the data types required when developing applications that will be connected with Fujitsu Enterprise Postgres.

.....

2.4.1 Relationship between the Application Data Types and Database Data Types

The following table shows the correspondence between data types in applications and data types in databases.

Data type on the server	Java data type	Data types prescribed by java.sql.Types
character	String	java.sql.Types.CHAR
national character	String	java.sql.Types.NCHAR
character varying	String	java.sql.Types.VARCHAR
national character varying	String	java.sql.Types.NVARCHAR
text	String	java.sql.Types.VARCHAR
bytea	byte[]	java.sql.Types.BINARY
smallint	short	java.sql.Types.SMALLINT
integer	int	java.sql.Types.INTEGER
bigint	long	java.sql.Types.BIGINT
smallserial	short	java.sql.Types.SMALLINT
serial	int	java.sql.Types.INTEGER
bigserial	long	java.sql.Types.BIGINT
real	float	java.sql.Types.REAL
double precision	double	java.sql.Types.DOUBLE
numeric	java.math.BigDecimal	java.sql.Types.NUMERIC
decimal	java.math.BigDecimal	java.sql.Types.DECIMAL
money	String	java.sql.Types.OTHER
date	java.sql.Date	java.sql.Types.DATE
time with time zone	java.sql.Time	java.sql.Types.TIME
time without time zone	java.sql.Time	java.sql.Types.TIME
timestamp without time zone	java.sql.Timestamp	java.sql.Types.TIMESTAMP
timestamp with time zone	java.sql.Timestamp	java.sql.Types.TIMESTAMP
interval	org.postgresql.util.PGInterval	java.sql.Types.OTHER
boolean	boolean	java.sql.Types.BIT
bit	boolean	java.sql.Types.BIT
bit varying	org.postgresql.util.Pgobject	java.sql.Types.OTHER
oid	long	java.sql.Types.BIGINT
xml	java.sql.SQLXML	java.sql.Types.SQLXML

Data type on the server	Java data type	Data types prescribed by java.sql.Types
array	java.sql.Array	java.sql.Types.ARRAY
uuid	java.util.UUID	java.sql.Types.OTHER
point	org.postgresql.geometric.Pgpoint	java.sql.Types.OTHER
box	org.postgresql.geometric.Pgbox	java.sql.Types.OTHER
lseg	org.postgresql.geometric.Pglseg	java.sql.Types.OTHER
path	org.postgresql.geometric.Pgpath	java.sql.Types.OTHER
polygon	org.postgresql.geometric.PGpolygon	java.sql.Types.OTHER
circle	org.postgresql.geometric.PGcircle	java.sql.Types.OTHER
json	org.postgresql.util.PGobject	java.sql.Types.OTHER
Network address type (inet,cidr,macaddr, macaddr8)	org.postgresql.util.PGobject	java.sql.Types.OTHER
Types related to text searches (svector, tsquery)	org.postgresql.util.PGobject	java.sql.Types.OTHER
Enumerated type	org.postgresql.util.PGobject	java.sql.Types.OTHER
Composite type	org.postgresql.util.PGobject	java.sql.Types.OTHER
Range type	org.postgresql.util.PGobject	java.sql.Types.OTHER

Although the getString() method of the ResultSet object can be used for all server data types, it is not guaranteed that it will always return a string in the same format for the same data type.

Strings in a format compatible with the JDBC specifications can be obtained using the Java toString() method of the appropriate data type (for example, getInt(), getTimestamp()) to conform to the data type on the server.

2.4.2 Statement Caching Feature

The statement caching feature caches SQL statements for each individual connection. This means that when an SQL statement with an identical string is next executed, the analysis and creation of the statement can be skipped. This improves performance in cases such as when an SQL statement with an identical string is executed within a loop or method that is executed repeatedly. Furthermore, the statement caching feature can be combined with the connection pooling feature to further enhance performance.

Cache registration controls

You can configure whether to cache SQL statements using the setPoolable(boolean) method of the PreparedStatement class when the statement caching feature is enabled.

Values that can be configured are shown below:

false

SQL statements will not be cached, even when the statement caching feature is enabled.

true

SQL statements will be cached if the statement caching feature is enabled.

2.4.3 Creating Applications while in Database Multiplexing Mode

This section explains points to consider when creating applications while in database multiplexing mode.



- Refer to the Cluster Operation Guide (Database Multiplexing) for information on database multiplexing mode.

•••••

- Refer to "Application Development" in the Cluster Operation Guide (PRIMECLUSTER) for points to consider when creating applications using the failover feature integrated with the cluster software.

2.4.3.1 Errors when an Application Connection Switch Occurs and Corresponding Actions

If an application connection switch occurs while in database multiplexing mode, explicitly close the connection and then reestablish the connection or reexecute the application.

The table below shows errors that may occur during a switch, and the corresponding action to take.

State		Error information (*1)	Action
Server failure or Fujitsu Enterprise Postgres system failure	Failure occurs during access	57P01 08006 08007	After the switch is complete, reestablish the connection, or reexecute the application.
	Accessed during system failure	08001	
Switch to the standby server	Switched during access	57P01 08006 08007	
	Accessed during switch	08001	

^{*1:} Return value of the getSQLState() method of SQLException.

Chapter 3 ODBC Driver

This section describes application development using ODBC drivers.

3.1 Development Environment

Applications using ODBC drivers can be developed using ODBC interface compatible applications, such as Access, Excel, and Visual Basic.

Refer to the manuals for the programming languages corresponding to the ODBC interface for information about the environment for development.

Fujitsu Enterprise Postgres supports ODBC 3.5.

3.2 Setup

You need to set up PsqlODBC, which is an ODBC driver, in order to use applications that use ODBC drivers with Fujitsu Enterprise Postgres. PsqlODBC is included in the Fujitsu Enterprise Postgres client package.

The following describes how to register the ODBC drivers and the ODBC data source.

3.2.1 Registering ODBC Drivers

When using the ODBC driver on Linux platforms, register the ODBC driver using the following procedure:

1. Install the ODBC driver manager (unixODBC)



- Fujitsu Enterprise Postgres supports unixODBC Version 2.3 or later.

You can download unixODBC from the following site:

http://www.unixodbc.org/

- To execute unixODBC, you must first install libtool 2.2.6 or later.

You can download libtool from the following website:

http://www.gnu.org/software/libtool/

[Note]

- ODBC driver operation is supported.
- unixODBC operation is not supported.
- 2. Register the ODBC drivers

Edit the ODBC driver manager (unixODBC) odbcinst.ini file.



[location of the odbcinst.ini file]

unixOdbcInstallDir/etc/odbcinst.ini

Set the following content:

Definition name	Description	Setting value	
[Driver name]	ODBC driver name	Set the name of the ODBC driver. Select the two strings below that correspond to the application type. Concatenate the strings with no spaces, enclose in "[]", and then specify this as the driver name. Note The placeholders shown below are enclosed in angle brackets '<>' to avoid confusion with literal text. Do not include the angle brackets in the string. - Application architecture "FujitsuEnterprisePostgres fujitsuEnterprisePostgresClientVers >x64" - Encoding system used by the application - In Unicode (only UTF-8 can be used) "unicode" - Other than Unicode "ansi" Example: The encoding system used by the application is Unicode: "[FujitsuEnterprisePostgres fujitsuEnterprisePostgresClientVers>x6	
Descriptio n	Descriptio n of the ODBC driver	4unicode]" Specify a supplementary description for the current data source. Any description may be set.	
Driver64	Path of the ODBC driver (64- bit)	Set the path of the ODBC driver (64-bit). - If the encoding system is Unicode: fujitsuEnterprisePostgresClientInstallDir/ odbc/lib/psqlodbcw.so - If the encoding system is other than Unicode: fujitsuEnterprisePostgresClientInstallDir/ odbc/lib/psqlodbca.so	
FileUsage	Use of the data source file	Specify 1.	
Threading	Level of atomicity secured for connectio n pooling	Specify 2.	



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

[Fujitsu Enterprise Postgres<x>x64unicode]

Description = Fujitsu Enterprise Postgres <x> x64 unicode driver

Driver64 = /opt/fsepv<x>client64/odbc/lib/psqlodbcw.so

FileUsage = 1

FileUsage = 1 Threading = 2

w

3.2.2 Registering ODBC Data Sources(for Windows(R))

This section describes how to register ODBC data sources.

There are the following two ways to register ODBC data sources on Windows(R).

3.2.2.1 Registering Using GUI

This section describes how to start the [ODBC Data Source Administrator] and register ODBC data sources.

Use the following procedure to register ODBC data sources:

1. Start the [ODBC Data Source Administrator].

Select [Start] >> [Control Panel] >> [Administrative Tools] >> [ODBC Data Source Administrator].




To register data sources for 32-bit applications in Windows(R) for 64-bit, execute the ODBC administrator (odbcad32.exe) for 32-bit, as shown below.

%SYSTEMDRIVE%\WINDOWS\SysWOW64\odbcad32.exe

- 2. When only the current user is to use the ODBC data source, select [User DSN]. When all users using the same computer are to use the ODBC data source, select [System DSN].
- 3. Click [Add].
- 4. Select one of the following drivers from the list of available ODBC drivers displayed in [Create New Data Source], and then click [Finish]. The notation "x" indicates the version of the Fujitsu Enterprise Postgres client feature.
 - Fujitsu Enterprise Postgres Unicode x
 Select this driver if using Unicode as the application encoding system.
 - Fujitsu Enterprise Postgres ANSI *x* Select this driver if using other than Unicode as the application encoding system.

5. The [PostgreSQL Unicode ODBC Driver (psqlODBC) Setup] window is displayed. Enter or select the required items, then click [Save].



Set the following content:

Definition name	Setting value
Data Source	Specify the data source name to be registered in the ODBC driver manager. The application will select the name specified here and connect with the Fujitsu Enterprise Postgres database. This parameter cannot be omitted. Specify the following characters up to 32 bytes.
	- National characters
	- Alphanumerics
	- "_", "<", ">", "+", "`", " ", "~", """ " &" , """, "#" , "\$" , "%" , "-" , "^" , ":" , "/" , "."
Description	Specify a supplementary description for the current data source. Specify characters up to 255 bytes.
	- National characters
	- Alphanumerics
Database	Specify the database name to be connected.
SSLMode	Specify to encrypt communications. The default is "disable".
	The setting values for SSLMode are as follows:
	- disable: Connect without SSL
	- allow: Connect without SSL, and if it fails, connect using SSL
	- prefer: Connect using SSL, and if it fails, connect without SSL
	- require: Connect always using SSL
	- verify-ca: Connect using SSL, and use a certificate issued by a trusted CA (*1)

Definition name	Setting value
	- verify-full: Connect using SSL, and use a certificate issued by a trusted CA to verify if the server host name matches the certificate (*1)
Server	Specify the host name of the database server to connect to, using up to 63 bytes. This parameter cannot be omitted.
Port	Specify the port number to be used for remote access. The default value is "27500".
Username(*2)	Specify the user that will access the database.
Password(*2)	Specify the password for the user that will access the database.

^{*1:} If specifying either "verify-ca" or "verify-full", use the system environment variable PGSSLROOTCERT of your operating system to specify the CA certificate file as shown below.

Example:

```
Variable name: PGSSLROOTCERT
Variable value: cACertificateFile
```

3.2.2.2 Registering Using Commands

This section describes how to use commands to register ODBC data sources.

Use the following tools from Microsoft to register ODBC data sources.

- ODBCConf.exe
- Add-OdbcDsn

Refer to the Microsoft Developer Network (MSDN) Library for information on how to use these tools.

When using ODBCConf.exe

ODBCConf.exe is a tool supported on all Windows(R) platforms.

Specification format

```
ODBCConf.exe /A { dataSourceType "odbcDriverName" "optionName=value[| optionName=value...]"} [/Lv fileName]
```

Refer to the Microsoft MSDN library for information on the format and parameters.

Description

Set the following content:

Definition name	Setting value
Data source type	Specify the data source type. - "CONFIGSYSDSN": A system data source is created. This requires user admin rights. The data source can be used by all users of the same computer. - "CONFIGDSN": A user data source is created. The data source can be used by the current user only.

^{*2:} In consideration of security, specify the Username and the Password by the application.

Definition name	Setting value
	€ Note
	When CONFIGSYSDSN is specified as the data source type, it is necessary to execute the command in the command prompt in administrator mode.
ODBC driver name	Specify an ODBC driver name that has already been registered on the system. Specify one of the following.
	₲ Note
	The placeholders shown below are enclosed in angle brackets '<>' to avoid confusion with literal text. Do not include the angle brackets in the string.
	- "Fujitsu Enterprise Postgres Unicode <fujitsuenterprisepostgresclientvers>" Specify this driver name if using Unicode as the application encoding system.</fujitsuenterprisepostgresclientvers>
	 "Fujitsu Enterprise Postgres ANSI < fujitsuEnterprisePostgresClientVers>" Specify this driver name if using other than Unicode as the application encoding system.
Option name	The following items must be set:
	- "DSN": Specify the data source name.
	- "Servername": Specify the host name for the database server.
	- "Port": Specify the port number for connection to the database
	- "Database": Specify the database name.
	Specify the following values as required:
	- "UID": User ID
	- "Password": Password
	- "SSLMode": Specify to encrypt communications. The default is "disable". Refer to the SSLMode explanation in the table under step 5 of "3.2.2.1 Registering Using GUI" for information on how to configure SSLMode.
File Name	You can output process information to a file when creating a data source. This operand can be omitted.



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

ODBCConf.exe /A {CONFIGSYSDSN "Fujitsu Enterprise Postgres Unicode <x>" "DSN=odbcconf1| Servername=sv1|Port=27500|Database=db01|SSLMode=verify-ca"} /Lv log.txt

.....



In consideration of security, specify the UID and the Password by the application.

When using Add-OdbcDsn

Add-OdbcDsn is used in the PowerShell command interface.

Specification format

```
Add-OdbcDsn dataSourceName -DriverName "odbcDriverName" -DsnType dataSourceType - Platform oSArchitecture -SetPropertyValue @("optionName=value" [,"optionName=value"...])
```

Refer to the Microsoft MSDN library for information on the format and parameters.

Description

Set the following content:

Definition name	Setting value	
Data source name	Specify any name for the data source name.	
ODBC driver name	Specify an ODBC driver name that has already been registered on the system. Specify one of the following.	
	The placeholders shown below are enclosed in angle brackets '<>' to avoid confusion with literal text. Do not include the angle brackets in the string.	
	- "Fujitsu Enterprise Postgres Unicode <fujitsuenterprisepostgresclientvers>" Specify this driver name if using Unicode as the application encoding system.</fujitsuenterprisepostgresclientvers>	
	- "Fujitsu Enterprise Postgres ANSI < fujitsuEnterprisePostgresClientVers>" Specify this driver name if using other than Unicode as the application encoding system.	
Data source	Specify the data source type.	
type	"System": A system data source is created. Requires user admin rights. The data source can be used by all users of the same computer.	
	- "User": A user data source is created. The data source can be used by the current user only.	
	G Note	
	When System is specified as the data source type, it is necessary to execute the command in the administrator mode of the command prompt.	
OS	Specify the OS architecture of the system.	
architecture	- "32-bit": 32-bit system	
	- "64-bit": 64-bit system	
Option name	The following items must be set:	
Орион паше	- "Servername": Specify the host name for the database server.	
	- "Port": Specify the port number for connection to the database	
	- "Database": Specify the database name.	
1	Specify the following values as required:	

Definition name	Setting value
	- "SSLMode": Specify to encrypt communications. The default is "disable". Refer to the SSLMode explanation in the table under step 5 of "3.2.2.1 Registering Using GUI" for information on how to configure SSLMode.
	When using Add-OdbcDsn, the strings "UID" and "Password" cannot be set as option names. These can only be used when using ODBCConf.exe.



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

Add-OdbcDsn odbcps1 -DriverName "Fujitsu Enterprise Postgres Unicode <x>" -DsnType System - Platform 32-bit -SetPropertyValue @("Servername=sv1", "Port=27500", "Database=db01", "SSLMode=verify-ca")

.....

3.2.3 Registering ODBC Data Sources(for Linux)

This section describes how to register ODBC data sources on Linux.

1. Register the data sources

Edit the odbc.ini definition file for the data source.



Edit the file in the installation directory for the ODBC driver manager (unixODBC)

unixOdbcInstallDir/etc/odbc.ini

Or

Create a new file in the HOME directory

~/.odbc.ini



If unixOdbcInstallDir is edited, these will be used as the shared settings for all users that log into the system. If created in the HOME directory (\sim /), the settings are used only by the single user.

Set the following content:

Definition name	Setting value
[Data source name]	Set the name for the ODBC data source.
Description	Set a description for the ODBC data source. Any description may be set.
Driver	Set the following as the name of the ODBC driver. Do not change this value. Select the two strings below that correspond to the application type. Concatenate the strings with no spaces and then specify this as the driver name.

Definition name	Setting value
	The placeholders shown below are enclosed in angle brackets '<>' to avoid confusion with literal text. Do not include the angle brackets in the string. - Application architecture "Fujitsu Enterprise Postgres < fujitsuEnterprisePostgresClientVers>x64" - Encoding system used by the application - In Unicode (only UTF-8 can be used) "unicode"
	- Other than Unicode "ansi"
	Example: The encoding system used by the application is Unicode:
	"Fujitsu Enterprise Postgres fujitsuEnterprisePostgresClientVers >x64unicode"
Database	Specify the database name to be connected.
Servername	Specify the host name for the database server.
Username	Specify the user ID that will connect with the database.
Password	Specify the password for the user that will connect to the database.
Port	Specify the port number for the database server.
	The default is "27500".
SSLMode	Specify the communication encryption method. The setting values for SSLMode are as follows:
	- disable: Connect without SSL
	- allow: Connect without SSL, and if it fails, connect using SSL
	- prefer: Connect using SSL, and if it fails, connect without SSL
	- require: Connect always using SSL
	- verify-ca: Connect using SSL, and use a certificate issued by a trusted CA (*1)
	 verify-full: Connect using SSL, and use a certificate issued by a trusted CA to verify if the server host name matches the certificate (*1)
ReadOnly	Specify whether to set the database as read-only.
	- 1: Set read-only
	- 0: Do not set read-only

^{*1:} If specifying either "verify-ca" or "verify-full", use the environment variable PGSSLROOTCERT to specify the CA certificate file as shown below.

Example

export PGSSLROOTCERT=cACertificateFileStorageDir/root.crt



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

[MyDataSource]

Description = Fujitsu Enterprise Postgres

Driver = Fujitsu Enterprise Postgres<x>x64unicode

Database = db01Servername = sv1Port = 27500ReadOnly = 0

G Note

In consideration of security, specify the UserName and the Password by the application.

2. Configure the environment variable settings

To execute applications that use ODBC drivers, all of the following settings must be configured in the LD_LIBRARY_PATH environment variable:

- unixOdbcInstallDir(*1)/lib
- libtoolInstallDir(*1)/lib

*1: If the installation directory is not specified when unixODBC and libtool are installed, they will be installed in /usr/local.

3.2.4 Message Language and Encoding System Used by Applications Settings

This section explains the language settings for the application runtime environment and the encoding settings for the application.

Language settings

You must match the language settings for the application runtime environment with the message locale settings of the database server.

Messages output by an application may include text from messages sent from the database server. In the resulting text, the text of the application message will use the message locale of the application, and the text of the message sent by the database server will use the message locale of the database server. If the message locales do not match, more than one language or encoding system will be used. Moreover, if the encoding systems do not match, characters in the resulting text can be garbled.

- Linux

Set the locale for messages (LC_MESSAGES category) to match the message locale of the database server. This can be done in a few different ways, such as using environment variables. Refer to the relevant manual of the operating system for information on the setlocale function.

🜠 Example

Example of specifying "en_US.UTF-8" with the setlocale function

setlocale(LC_ALL, "en_US.UTF-8");

Specifying the locale of the LC_ALL category propagates the setting to LC_MESSAGE.

- Windows(R)

W

Align the locale of the operating system with the message locale of the database server.

Encoding System Settings

Ensure that the encoding system that is embedded in the application and passed to the database, and the encoding system setting of the runtime environment, are the same. The encoding system cannot be converted correctly on the database server.

Use one of the following methods to set the encoding system for the application:

- Set the PGCLIENTENCODING environment variable in the runtime environment.
- Set the client_encoding keyword in the connection string.
- Use the PQsetClientEncoding function.



See

Refer to "Supported Character Sets" in "Server Administration" in the PostgreSQL Documentation for information on the strings that represent the encoding system that can be set.

...........

For example, when using "Unicode" and "8 bit", set the string "UTF8".



Setting the "PGCLIENTENCODING" environment variable

An example of setting when the encoding of the client is "UTF8" (Bash)

> PGCLIENTENCODING=UTF8; export PGCLIENTENCODING

W

An example of setting when the encoding of the client is "UTF8"

> set PGCLIENTENCODING=UTF8



Text may be garbled when outputting results to the command prompt. Review the font settings for the command prompt if this occurs.

......

3.3 Connecting to the Database

Refer to the manual for the programming language corresponding to the ODBC interface, i.e. Access, Excel, or Visual Basic, for example.

3.4 Application Development

This section describes how to develop applications using ODBC drivers.

3.4.1 Compiling Applications (for Windows (R))

Refer to the manual for the programming language corresponding to the ODBC interface, i.e. Access, Excel, or Visual Basic, for example.



The cl command expects input to be a program that uses one of the following code pages, so convert the program to these code pages and then compile and link it (refer to the Microsoft documentation for details).

- ANSI console code pages (example: UTF8)
- UTF-16 little-endian with or without BOM (Byte Order Mark)
- UTF-16 big-endian with or without BOM
- UTF-8 with BOM

The cl command converts strings in a program to an ANSI console code page before generating a module, so the data sent to and received from the database server becomes an ANSI console code page. Therefore, set the coding system corresponding to the ANSI console code page as the coding system of the client.

Refer to "Character Set Support" in "Server Administration" in the PostgreSQL Documentation for information on how to set the client encoding system.

3.4.2 Compiling Applications (for Linux)

Specify the following options when compiling applications.

Table 3.1 Include file and library path

Option	How to specify the option
Path of the include file	-I unixOdbc64bitIncludeFileDir
Path of the library	-L unixOdbc64bitLibraryDir

Table 3.2 ODBC library

Type of library	Library name	
Dynamic library	libodbc.so	



Specify -m64 when creating a 64-bit application.



The following are examples of compiling ODBC applications:

gcc -m64 -I/usr/local/include(*1) -L/usr/local/lib(*1) -lodbc testproc.c -o testproc

*1: This is an example of building and installing from the source without specifying an installation directory for unixODBC. If you wish to specify a location, set the installation directory.

3.4.3 Creating Applications While in Database Multiplexing Mode

This section explains points to consider when creating applications while in database multiplexing mode.



- Refer to the Cluster Operation Guide (Database Multiplexing) for information on database multiplexing mode.

- Refer to "Application Development" in the Cluster Operation Guide (PRIMECLUSTER) for points to consider when creating applications using the failover feature integrated with the cluster software.

3.4.3.1 Errors when an Application Connection Switch Occurs and Corresponding Actions

If an application connection switch occurs while in database multiplexing mode, explicitly close the connection and then reestablish the connection or reexecute the application.

The table below shows errors that may occur during a switch, and the corresponding action to take.

State		Error information (*1)	Action
or Fujitsu Enterprise Postgres system during	Failure occurs during access	57P01 08S01	After the switch is complete, reestablish the connection, or reexecute the application.
	Accessed during system failure	08001	
Switch to the standby server	Switched during access	57P01 08S01	
	Accessed during switch	08001	

^{*1:} Return value of SQLSTATE.

Chapter 4 .NET Data Provider

This chapter describes how to configure for the purpose of creating .NET applications.

4.1 Development Environment

.NET Data Provider can operate in the following environments:

Environment required for development and operation of .NET applications	.NET 8.0 .NET 7.0 .NET 6.0 .NET Standard 2.0 .NET Standard 2.1
Integrated development environment for applications running in a .NET environment	Visual Studio 2022 Visual Studio 2019 Visual Studio 2017 Visual Studio 2015
Available development languages	C# Visual Basic .NET

4.2 Setup

This section explains how to set up .NET Data Provider and Entity Framework Core.

4.2.1 Setting Up .NET Data Provider

The .NET Data Provider of Fujitsu Enterprise Postgres can be installed as a local NuGet package. Follow the steps below to install.

Location of NuGet package

The NuGet package are stored in the following location. "<x>" indicates the product version.

- Linux

- Windows(R)

fujitsuEnterprisePostgresClientInstallDir\DOTNET\Npgsql.<x>.0.0.nupkg

Add a local package source

Add a NuGet local package source if one does not exist.

Linux

Add the above Nuget package location as a local package source.

 ${\tt dotnet\ nuget\ add\ source\ fujitsuEnterprisePostgresClientInstallDir/DOTNET\ -n\ sourceName}$

- Windows(R)

- 1. Start Visual Studio, click [Tools] >> [Options] >> [NuGet Package Manager], and then select [Package Sources].
- 2. Click [+] in the upper-right corner, and then set [Name] to "Local Package Source".
- 3. Click [···] and navigate to the folder above. Select this folder, and then click [OK].

Install the NuGet package

Install the NuGet package from the local package source.

- Linux

Add a package reference to the project file.

dotnet add projectFilePath package Npgsql

w

- Windows(R)
 - 1. Start Visual Studio, click [Tools] >> [NuGet Package Manager] >> [Manage NuGet Packages for Solution].
 - 2. In the upper-right corner, select "Local Package Source" from [Package Source].
 - 3. Once the local package source is set, all available NuGet packages in this local location will be displayed. Select "Npgsql", and then select the projects for which this package is to be installed.
 - 4. Click [Install].

4.2.2 Setting Up .NET Data Provider Type Plugins

Fujitsu Enterprise Postgres .NET Data Provider now comes packaged with 4 Type Plugins that provide additional support for more data type mappings (eg date time support via the Npgsql.NodaTime). The plugins modify how Npgsql maps the PostgreSQL values to CLR types.

The type plugins are available for installation as a local NuGet packages. There are 4 pacakages available for installation. "<x>" indicates the product version.

- Npgsql.NodaTime.<x>.0.0.nupkg
- Npgsql.Json.NET.<x>.0.0.nupkg
- Npgsql.NetTopologySuite.<x>.0.0.nupkg: Spatial type(*1)
- Npgsql.GeoJSON.<x>.0.0.nupkg: Spatial type(*1)

*1: Please note that the Spatial Type plugins require the PostGIS extension installed on the server.

Also, Refre to "Type Plugins" and "Additional Notes on each Type Plugin" on the each type plugins.

To install any of the plugins please follow the procedure below:

Location of NuGet package

The Plugin NuGet package are stored in the following location. "<x>" indicates the product version.

- Linux

fujitsuEnterprisePostgresClientInstallDir/DOTNET/Npgsql.*.<x>.0.0.nupkq

W

- Windows(R)

 $\verb|fujitsuEnterprisePostgresClientInstallDir\DOTNET\Npgsql.*.<\!x>.0.0.nupkg$

Add a local package source

Add a NuGet local package source if one does not exist.

- Linux

Add the above Nuget package location as a local package source.

 ${\tt dotnet\ nuget\ add\ source\ fujitsuEnterprisePostgresClientInstallDir/DOTNET\ -n\ sourceName}$

- W
- Windows(R)
 - 1. Start Visual Studio, click [Tools] >> [Options] >> [NuGet Package Manager], and then select [Package Sources].
 - 2. Click [+] in the upper-right corner, and then set [Name] to "Local Package Source".
 - 3. Click [···] and navigate to the folder above. Select this folder, and then click [OK].

Install the NuGet package

Install the NuGet package from the local package source.

Lin - Lin

Add a package reference to the project file (Example: Npgsql.NodaTime).

dotnet add projectFilePath package Npgsql.NodaTime

W

- Windows(R)
 - 1. Start Visual Studio, click [Tools] >> [NuGet Package Manager] >> [Manage NuGet Packages for Solution].
 - 2. In the upper-right corner, select "Local Package Source" from [Package Source].
 - Once the local package source is set, all available NuGet packages in this local location will be displayed. Select the plugin to be installed (eg "Npgsql.NodaTime", and then select the projects for which this package is to be installed.
 - 4. Click [Install].

4.2.3 Setting Up Entity Framework Core

Entity Framework Core is supplied as a NuGet package file. To install it locally, follow the procedure below.

Location of NuGet package

The EntityFramework Core NuGet package is stored in the following location.

LΙ

- Linux

 $fujits u {\tt EnterprisePostgresClientInstallDir/DOTNET/Npgsql.EntityFrameworkCore.PostgreSQL.8.0.2.nupkg$

W

- Windows(R)

fujitsuEnterprisePostgresClientInstallDir\DOTNET
\Npgsql.EntityFrameworkCore.PostgreSQL.8.0.2.nupkg

Add a local package source

Add a NuGet local package source if one does not exist.

- Linux

Add the above Nuget package location as a local package source.

 $\label{thm:control} \mbox{dotnet nuget add source fujitsuEnterprisePostgresClientInstallDir/DOTNET -n} \\ sourceName$

W

- Windows(R)
 - 1. Start Visual Studio, click [Tools] >> [Options] >> [NuGet Package Manager], and then select [Package Sources].
 - 2. Click [+] in the upper-right corner, and then set [Name] to "Local Package Source".
 - 3. Click [···] and navigate to the folder above. Select this folder, and then click [OK].

Install the NuGet package

Install the NuGet package from the local package source.

Linux

Add a package reference to the project file.

dotnet add projectFilePath package Npgsql.EntityFrameworkCore.PostgreSQL

w

- Windows(R)
 - 1. Start Visual Studio, click [Tools] >> [NuGet Package Manager] >> [Manage NuGet Packages for Solution].
 - 2. In the upper-right corner, select "Local Package Source" from [Package Source].
 - Once the local package source is set, all available NuGet packages in this local location will be displayed. Select "Npgsql.EntityFrameworkCore.PostgreSQL", and then select the projects for which this package is to be installed.
 - 4. Click [Install].

4.2.4 Setting Up Npgsql.OpenTelemetry

Fujitsu Enterprise Postgres .NET Data Provider comes with a plug-in that supports tracing by the observability framework OpenTelemetry. This will allow Npgsql to emit trace data.

OpenTelemetry plugin is supplied as a local NuGet package. To install it locally, follow the procedure below.

Location of NuGet package

The NuGet packages for OpenTelemetry are stored in the following location. "<x>" indicates the product version.

- Linux

 $fujitsu \textit{EnterprisePostgresClientInstallDir/DOTNET/Npgsql.OpenTelemetry.< x>.0.0.nupkg$

W

L

- Windows(R)

fujitsuEnterprisePostgresClientInstallDir\DOTNET\Npgsql.OpenTelemetry.<x>.0.0.nupkg

Add a local package source

Add a NuGet local package source if one does not exist.

L Linux

Add the above Nuget package location as a local package source.

 ${\tt dotnet\ nuget\ add\ source\ fujitsuEnterprisePostgresClientInstallDir/DOTNET\ -n\ sourceName}$

W

- Windows(R)
 - 1. Start Visual Studio, click [Tools] >> [Options] >> [NuGet Package Manager], and then select [Package Sources].
 - 2. Click [+] in the upper-right corner, and then set [Name] to "Local Package Source".
 - 3. Click [...] and navigate to the folder above. Select this folder, and then click [OK].

Install the NuGet package

Install the NuGet package from the local package source.

- Linux

Add a package reference to the project file.

dotnet add projectFilePath package Npgsql.OpenTelemetry

W

- Windows(R)
 - 1. Start Visual Studio, click [Tools] >> [NuGet Package Manager] >> [Manage NuGet Packages for Solution].
 - 2. In the upper-right corner, select "Local Package Source" from [Package Source].
 - 3. Once the local package source is set, all available NuGet packages in this local location will be displayed. Select "Npgsql.OpenTelemetry", and then select the projects for which this package is to be installed.
 - 4. Click [Install].

4.3 Connecting to the Database

This section explains how to connect to a database.

- Using NpgsqlConnection
- Using NpgsqlConnectionStringBuilder

4.3.1 Using NpgsqlConnection

Connect to the database by specifying the connection string.



Code examples for applications

```
using Npgsql;

NpgsqlConnection conn = new NpgsqlConnection("Server=sv1;Port=27500;Database=mydb;
Username=myuser;Password=myuser01; Timeout=20;CommandTimeout=20;");
```

Refer to "4.3.3 Connection String" for information on database connection strings.

4.3.2 Using NpgsqlConnectionStringBuilder

Generate connection strings by specifying the connection information in the properties of the NpgsqlConnectionStringBuilder object.



Code examples for applications

```
using Npgsql;

NpgsqlConnectionStringBuilder sb = new NpgsqlConnectionStringBuilder();
sb.Host = "sv1";
sb.Port = 27500;
sb.Database = "mydb";
sb.Username = "myuser";
sb.Password = "myuser01";
sb.Timeout = 20;
sb.CommandTimeout = 20;
NpgsqlConnection conn = new NpgsqlConnection(sb.ConnectionString);
```

Refer to "4.3.3 Connection String" for information on database connection strings.

4.3.3 Connection String

Specify the following connection information to connect to the database.

Server=127.0.0.1;Port=27500;Database=mydb;Username=myuser;Password=myuser01;...;
(1) (2) (3) (4) (5) (6)

- (1) Specify the host name or IP address of the server to be connected. This must be specified.
- (2) Specify the port number for the database server. The default is "27500".
- (3) Specify the database name to be connected.
- (4) Specify the username that will connect with the database.
- (5) Specify the password for the user that will connect to the database.
- (6) Refer to the following for information on how to specify other connection information.

The table below shows keywords that are available to specify in the connection string in .NET Data Provider (Npgsql):

Note that some settings require care if using an Oracle database-compatible feature (refer to "9.2.3 Notes when Integrating with the Interface for Application Development" for details).

Basic connection

Keyword	Description	Default
Host	Specify the host name of the server to be connected. Multiple hosts may be specified. A host name must be specified.	
Port	Specify the TCP port number of the server.	27500
Database	Specify the database name to connect to. Sa Us	
Username	Specify the username for the username that will connect to the database. Not required if using Integrated Security.	PGUSER
Password	Specify the password for the username that will connect to the database. Not required if using Integrated Security.	PGPASSWO RD
Passfile	Specify the path of the password file (PGPASSFILE) to obtain the password.	PGPASSFILE

Security and encryption

Keyword	Description	Default
SSL Mode	Specify one of the following values to indicate whether or not to use SSL, depending on your server's support.	Prefer
	Disable: No SSL connection is made.	
	Allow: SSL connection will be made only if the server requires it.	
	Prefer: Make an SSL connection if possible.	
	Require: Perform SSL connection. It does not verify the authenticity of Connection Server at this time.	
	VerifyCA: Make an SSL connection. At this time, verify the authenticity of the connection server.	

Keyword	Description	Default
	VerifyFull: Make an SSL connection. This time it verifies the authenticity of the Connection Server and verifies that it is the server you specified.	
SSL Certificate	Specify the location of the client certificate sent to the server.	PGSSLCERT
SSL Key	Specify the location of the client key for the client certificate sent to the server.	PGSSLKEY
SSL Password	Specify the password for the client certificate key.	
Root Certificate	Specify the location of the CA certificate used to validate server certificates.	PGSSLROOT CERT
Check Certificate Revocation	Specify whether to check the certificate revocation list during authentication.	false
Integrated Security	Specify whether to log in using integrated security (GSS/SSPI).	false
Persist Security Info	Gets or sets a Boolean value that indicates whether security- sensitive information, such as passwords, should not be returned as part of a connection if it is connected or has ever been connected.	false
Kerberos Service Name	Specify the Kerberos service name used for authentication.	postgres
Include Realm	Specify the Kerberos realm used for authentication.	
Include Error Detail	When enabled, database error and notification details are included in PostgresException.Detail and PostgresNotice.Detail. These can contain confidential information.	
Log Parameters	When enabled, parameter values are logged when the command is executed.	false

Pooling

Keyword	Description	Default
Pooling	Specify whether to use connection pooling.	true
Minimum Pool Size	Minimum connection pool size.	0
Maximum Pool Size	Maximum size of the connection pool. 100	
Connection Idle Lifetime	Specify the time (in seconds) to wait before disconnecting idle connections in the pool when the number of all connections exceeds the Minimum Pool Size.	
Connection Pruning Interval	Specify the number of seconds the pool waits before attempting to prune an idle connection that has exceeded its lifetime. See Connection Idle Lifetime.	
ConnectionLifetime	Specify the total maximum lifetime of the connection in seconds. Connections beyond this value are discarded rather than returned from the pool. This is useful in cluster configurations that force load balancing between running servers and servers that have just come online.	0 (disabled)

Timeouts and keepalive

Keyword	Description	Default
Timeout	Specify the amount of time (in seconds) to wait before terminating the attempt and generating an error during connection establishment.	15
Command Timeout	Specify how long to wait (in seconds) while executing a command before terminating the attempt and generating an error. Set to 0 for infinity.	30
Internal Command Timeout	Specify how long to wait (in seconds) before terminating the attempt and generating an error when attempting to execute an internal command1 uses CommandTimeout, 0 means no timeout.	-1
Cancellation Timeout	Specify the number of milliseconds to wait while trying to read a response to a cancel request for a timed-out or canceled query before terminating the attempt and generating an error1 skips the wait, 0 means infinite wait.	2000
Keepalive	Npgsql will send his keep-alive queries if the connection has been inactive for more than the number of seconds specified here.	0 (disabled)
Tcp Keepalive	Specify whether TCP keepalives are used by the system default if no override is specified.	false
Tcp Keepalive Time	A TCP keep-alive query is sent when the connection inactivity exceeds the number of milliseconds specified here. Use of this option is not recommended. We recommend using Keepalive. This is supported on Windows only.	
Tcp Keepalive Interval	Specify the interval (in milliseconds) between successive keepalive packets sent if no acknowledgment is received. Tcp KeepAlive Time must also be non-zero. This is supported on Windows only.	TCP Keepalive Time value

Performance

Keyword	Description	Default
Max Auto Prepare	Maximum number of SQL statements that can be automatically prepared at any given time. When this number is exceeded, the last used statement is reused. A value of 0 disables autoprepare.	0
Auto Prepare Min Usages	A SQL statement is automatically prepared when its usage count exceeds the value specified here.	5
Read Buffer Size	The size of the internal buffer Npgsql uses when reading. Increasing this value may improve performance when transferring large amounts of data from the database.	8192
Write Buffer Size	Determines the size of the internal buffer Npgsql uses when writing. Increasing this value may improve performance when transferring large amounts of data to the database.	8192
Socket Receive Buffer Size	The size of the socket receive buffer.	System dependent
Socket Send Buffer Size	The size of the socket send buffer.	System dependent
No Reset On Close	When true, do not reset the connection state when the connection is returned to the pool. In some cases, this improves performance at the cost of leak conditions. Use it only if your benchmark shows performance improvements.	false

Failover and load balancing

Keyword	Description	Default
Target Session Attributes	Specifies the server type to preferentially connect to. any: Allows any successful connection. primary: Connect to the primary server. standby: Connect to the standby server. prefer-primary: Attempt to connect to the primary server first, but	PGTARGETS ESSIONATT RS, Any
	retry in any mode if none of the listed hosts has a primary server. prefer-standby: Attempt to connect to the standby server first, but retry in any mode if there is no standby server on the listed host. read-write: Connect to a server whose default transaction mode is read/write. read-only: Connect to a server whose default transaction mode is read-only.	
Load Balance Hosts	Specify whether to perform load balancing among multiple hosts using the round-robin method.	false
Host Recheck Seconds	Controls how long a host's cached state is considered valid.	10
Enable Fdw Acs	Specifies whether to switch the connection destination type of the data node with on or off when using the connection routing function. This cannot be set to on when "Target Session Attributes=prefer-primary" is specified.	off

Others

Keyword	Description	Default
Options	Specify valid database connection options. PG	
	(Example: Options=-c synchronous_commit=local)	
Application Name	Specify the application name sent to the backend when opening a connection.	
Enlist	Specify whether the connection should participate in transactions declared in the transaction scope.	true
Search Path	Sets the schema search path.	None
Client Encoding	Gets or sets the client_encoding parameter.	PGCLIENTE NCODING
Encoding	Gets or sets the .NET encoding used to encode/decode string data in the database.	UTF8
Timezone	Gets or sets the session timezone.	PGTZ
EF Template Database	A database template that is specified when creating a database with Entity Framework.	template1
EF Admin Database	The managed database that you specify when creating and deleting databases in Entity Framework.	
Load Table Composites	Specify whether to load complex type definitions for tables in addition to independent complex types.	false
Array Nullability Mode	Sets how to return an array of value types when requested as an object instance. Possible values are:	Never

Keyword	Description	Default
	Never: Always return as a non-nullable array.	
	Always: Always return as a nullable array.	
	PerInstance: The returned array is determined at runtime.	

4.4 Application Development

4.4.1 Data Types

A variety of data types can be used with Fujitsu Enterprise Postgres.

List of supported data types

Data Type		
boolean		
smallint		
integer		
bigint		
real		
double precision		
numeric		
money		
text		
character varying		
character		
national character varying		
national character		
citext		
json		
jsonb		
xml		
uuid		
bytea		
timestamp without time zone		
timestamp with time zone		
date		
time without time zone(*1)		
time with time zone(*1)		
interval(*2)		
cidr		
inet		
macaddr		
tsquery		

Data Type	
tsvector	
bit(1)	
bit(n)	
bit varying	
point	
lseg	
path	
polygon	
line	
circle	
box	
hstore	
oid	
xid	
cid	
oidvector	
name	
(internal) char	
geometry (PostGIS)	
record	
composite types	
range types	
multirange types (PG14)	
enum types	
array types	

*1: As shown below, "time with time zone" and "time without time zone" values display the date portion as additional information. However, the actual data comprises the time data only, so with the exception of this displayed format, there are no other resulting issues.

Example:

Composition of table (t1)

col1 (time with time zone)	col2 (time without time zone)
1/01/0001 10:21:30 +08:00	10:21:30
1/01/0001 23:34:03 +08:00	23:34:03
1/01/0001 17:23:54 +08:00	17:23:54

"time with time zone" values display a fixed value of "2/01/0001" in the date portion, while "time without time zone" values display just the time data.

1/01/0001	10:21:30	+08:00	10:21:30
1/01/0001	23:34:03	+08:00	23:34:03
1/01/0001	17:23:54	+08:00	17:23:54

^{*2:} The format is d.hh:mm:ss, where d is an integer and hh:mm:ss is a maximum of 23.59.59 (23 hours, 59 minutes, and 59 seconds).

4.4.2 Relationship between Application Data Types and Database Data Types

The data types available for SQL data types are as follows:

Read mappings

PostgreSQL Type	Default.NET Type	Non-default Type
boolean	bool	
smallint	short	byte, sbyte, int, long, float, double, decimal
integer	int	byte, short, long, float, double, decimal
bigint	long	long, byte, short, int, float, double, decimal
real	float	double
double precision	double	
numeric	decimal	byte, short, int, long, float, double, BigInteger
money	decimal	
text	string	char[]
character varying	string	char[]
character	string	char[]
national character varying	string	char[]
national character	string	char[]
citext	string	char[]
json	string	char[]
jsonb	string	char[]
xml	string	char[]
uuid	Guid	
bytea	byte[]	
timestamp without time zone	DateTime (Unspecified)	
timestamp with time zone	DateTime (Utc)(*1)	DateTimeOffset (Offset=0)(*2)
date	DateTime	DateOnly
time without time zone	TimeSpan	TimeOnly
time with time zone	DateTimeOffset	
interval	TimeSpan(*3)	NpgsqlInterval
cidr	NpgsqlCidr	

PostgreSQL Type	Default.NET Type	Non-default Type
inet	IPAddress	NpgsqlInet
macaddr	PhysicalAddress	
tsquery	NpgsqlTsQuery	
tsvector	NpgsqlTsVector	
bit(1)	bool	BitArray
bit(n)	BitArray	
bit varying	BitArray	
point	NpgsqlPoint	
lseg	NpgsqlLSeg	
path	NpgsqlPath	
polygon	NpgsqlPolygon	
line	NpgsqlLine	
circle	NpgsqlCircle	
box	NpgsqlBox	
hstore	Dictionary <string, string=""></string,>	
oid	uint	
xid	uint	
cid	uint	
oidvector	uint[]	
name	string	char[]
(internal) char	char	byte, short, int, long
geometry (PostGIS)	PostgisGeometry	
record	object[]	
composite types	Т	
range types	NpgsqlRange <telement></telement>	
multirange types (PG14)	NpgsqlRange <telement>[]</telement>	
enum types	TEnum	
array types	Array (of element type)	

^{*1:} When Npgsql.EnableLegacyTimestampBehavior is enabled, reading a timestamp with time zone returns a Local DateTime instead of Utc.

Write mappings

PostgreSQL Type	Default.NET Type	Non-default.NET Type	NpgsqlDbType	DbType
boolean	bool		Boolean	Boolean

^{*2:} When Npgsql.EnableLegacyTimestampBehavior is enabled, reading a timestamp with time zone as a DateTimeOffset returns a local offset based on the timezone of the server where Npgsql is running.

^{*3:} PostgreSQL intervals with month or year components cannot be read as TimeSpan. Consider using NodaTime's Period type, or NpgsqlInterval.

PostgreSQL Type	Default.NET Type	Non-default.NET Type	NpgsqlDbType	DbType
smallint	short, byte, sbyte		Smallint	Int16
integer	int		Integer	Int32
bigint	long		Bigint	Int64
real	float		Real	Single
double precision	double		Double	Double
numeric	decimal, BigInteger		Numeric	Decimal, VarNumeric
money		decimal	Money	Currency
text	string, char[], char		Text	String, StringFixedLength, AnsiString, AnsiStringFixedLengt
character varying		string, char[], char	Varchar	
character		string, char[], char	Char	
citext		string, char[], char	Citext	
json		string, char[], char	Json	
jsonb		string, char[], char	Jsonb	
xml		string, char[], char	Xml	
uuid	Guid		Uuid	
bytea	byte[]	ArraySegment >	Bytea	Binary
timestamp without time zone	DateTime (Utc) (*1), DateTimeOffset		TimestampTz	DateTime(*2), DateTimeOffset
timestamp with time zone	DateTime (Local/ Unspecified) (*1)		Timestamp	DateTime2
date	DateOnly	DateTime	Date	Date
time without time zone	TimeOnly	TimeSpan	Time	Time
time with time zone		DateTimeOffset	TimeTz	
interval	TimeSpan	NpgsqlInterval	Interval	
cidr		ValueTuple <ipadd ress, int>, IPAddress</ipadd 	Cidr	
inet	IPAddress	ValueTuple <ipadd int="" ress,=""></ipadd>	Inet	
macaddr	PhysicalAddress		MacAddr	
tsquery	NpgsqlTsQuery		TsQuery	
tsvector	NpgsqlTsVector		TsVector	
bit		bool, BitArray, string	Bit	

PostgreSQL Type	Default.NET Type	Non-default.NET Type	NpgsqlDbType	DbType
bit varying	BitArray	bool, BitArray, string	Varbit	
point	NpgsqlPoint		Point	
lseg	NpgsqlLSeg		LSeg	
path	NpgsqlPath		Path	
polygon	NpgsqlPolygon		Polygon	
line	NpgsqlLine		Line	
circle	NpgsqlCircle		Circle	
box	NpgsqlBox		Box	
hstore	IDictionary <stri ng, string></stri 		Hstore	
oid		uint	Oid	
xid		uint	Xid	
cid		uint	Cid	
oidvector		uint[]	Oidvector	
name		string, char[], char	Name	
(internal) char		byte	InternalChar	
composite types	Pre-mapped type		Composite	
range types	NpgsqlRange <t Subtype></t 		Range NpgsqlDbType	
enum types	Pre-mapped type		Enum	
array types	T[], List <t></t>		Array NpgsqlDbType	

^{*1:} UTC DateTime is written as timestamp with time zone, Local/Unspecified DateTimes are written as timestamp without time zone. When Npgsql.EnableLegacyTimestampBehavior is enabled, DateTime is always written as timestamp without time zone.

4.4.3 Creating Applications while in Database Multiplexing Mode

This section explains points to consider when creating applications while in database multiplexing mode.



- Refer to the Cluster Operation Guide (Database Multiplexing) for information on database multiplexing mode.
- Refer to "Application Development" in the Cluster Operation Guide (PRIMECLUSTER) for points to consider when creating applications using the failover feature integrated with the cluster software.

4.4.3.1 Errors when an Application Connection Switch Occurs and Corresponding Actions

If an application connection switch occurs while in database multiplexing mode, explicitly close the connection and then reestablish the connection or reexecute the application.

The table below shows errors that may occur during a switch, and the corresponding action to take.

^{*2:} When Npgsql.EnableLegacyTimestampBehavior is enabled, DbType.DateTime is mapped to timestamp without time zone.

State		Error information	Action
Server failure or Fujitsu Enterprise Postgres	Failure occurs during access	57P01 (*1) Empty string (*1)	After the switch is complete, reestablish the connection, or reexecute
system failure	Accessed during system failure	Empty string (*1)	the application.
Switch to the standby server	Switched during access	57P01 (*1) Empty string (*1)	
	Accessed during switch	Empty string (*1)	

^{*1:} This is the return value of the PostgresException attribute SqlState.

4.4.4 Notes

Type Plugins

- These type libraries include:
 - NodaTime the recommended way to interact with PostgrSQL date/time types
 - Json.NET allows Npgsql to use the NewtonSoft Json.NET library when reading and writing JSON data (both json and jsonb)
 - NetTopologySuite allows Npgsql to map PostGIS spatial types directly to the NetTopology suite types (the leading spatial library in .NET)
 - GeoJSON allows Npgsql to read and write PostGIS spatial types as GeoJSON types via the GeoJSON.NET library
- Setup the plugin in your application simply by adding a dependency on the plugin (this should have been done automatically when installed to the project) and set it up. See the following code snippet for an example of setting up the Npgsql.NodaTime plugin:

```
using Npgsql;
   // Place this at the beginning of your program to use NodaTime everywhere
(recommended)
   NpgsqlConnection.GlobalTypeMapper.UseNodaTime();
   // Or to temporarily use NodaTime on a single connection only:
   conn.TypeMapper.UseNodaTime();
```

Once the plugin is setup, you can read and write NodaTime objects as per the code snippet below:

```
// Write NodaTime Instant to PostgreSQL "timestamp without time zone"
    using (var cmd = new NpgsqlCommand(@"INSERT INTO mytable (my_timestamp) VALUES
(@p)", conn))
    {
        cmd.Parameters.Add(new NpgsqlParameter("p", Instant.FromUtc(2011, 1, 1, 10,
30)));
        cmd.ExecuteNonQuery();
    }

    // Read timestamp back from the database as an Instant
    using (var cmd = new NpgsqlCommand(@"SELECT my_timestamp FROM mytable", conn))
    using (var reader = cmd.ExecuteReader())
    {
        reader.Read();
        var instant = reader.GetFieldValue<Instant>(0);
    }
}
```

- To apply the type plugin updates, do one of the following:
 - After uninstalling the type plugin (Refer to "4.5.2 Uninstalling .NET Data Provider Type Plugins"), setup the type plugin (Refer to "4.2.2 Setting Up .NET Data Provider Type Plugins").
 - Remove the type plugin directory from the packages directory of the solution, and then restore it using the nuget restore command.
- When you deploy an application with a type plugin, the type plugin is included in the distribution. Therefore, after applying the type plugin updates, you must rebuild the application and deploy the updated application.

Additional Notes on each Type Plugin

Describe notes about each type plugin.

NodaTime

Describes the mapping of PostgreSQL data types to NodaTime data types.

Mapping Table

PostgreSQL Type	Default NodaTime Type	Additional NodaTime Type	Note
timestamp with time zone	Instant	ZonedDateTime(*1), OffsetDateTime(*1)	UTC timestamp in database. Only UTC ZonedDateTime and OffsetDateTime are supported.
timestamp without time zone	LocalDateTime(*2)		A timestamp in an unknown or implied timezone.
date	LocalDate		A simple date with no timezone or offset information.
time without time zone	LocalTime		A simple time with no timezone or offset information.
time with time zone	OffsetTime		A type that stores a time and an offset. Generally not recommended for use.
interval	Period	Duration	A time interval from fractional seconds to years. NodaTime Duration is supported for days and smaller intervals, but not for years or months (because they have no absolute duration). Duration can be used in any interval unit.
tszrange	Interval	NpgsqlRange <instant>,etc</instant>	The time interval between two time instances (start and end).
strange	NpgsqlRange <localdatetim e=""></localdatetim>		The time interval between two

PostgreSQL Type	Default NodaTime Type	Additional NodaTime Type	Note
			timestamps in unknown or implied timezones.
daterange	DateInterval	NpgsqlRange <localdate>,e</localdate>	The interval between two dates.

^{*1:} When Npgsql.EnableLegacyTimestampBehavior is enabled, writing or reading ZonedDateTime or OffsetDateTime is automatically converted to UTC.

Json.NET

Once the JSON plugin has been setup, users can transparently read and write CLR objects as JSON values and the plugin will automatically serialize/deserialize them

See the code snippet below:

```
// Write arbitrary CLR types as JSON
using (var cmd = new NpgsqlCommand(@"INSERT INTO mytable (my_json_column) VALUES (@p)",
conn))
{
    cmd.Parameters.Add(new NpgsqlParameter("p", NpgsqlDbType.Jsonb) { Value =
    MyClrType });
    cmd.ExecuteNonQuery();
}

// Read arbitrary CLR types as JSON
using (var cmd = new NpgsqlCommand(@"SELECT my_json_column FROM mytable", conn))
using (var reader = cmd.ExecuteReader())
{
    reader.Read();
    var someValue = reader.GetFieldValue<MyClrType>(0);
}
```

NetTopologySuite (spatial)

By default the plugin handles only ordinates provided by the DefaultCoordinateSequenceFactory of GeometryServiceProvider.Instance. If GeometryServiceProvider is initialized automatically the X and Y ordinates are handled. To change the behavior specify the handleOrdinates parameter like in the following example:

```
conn.TypeMapper.UseNetTopologySuite(handleOrdinates: Ordinates.XYZ);
```

To process the M ordinate, you must initialize GeometryServiceProvider.Instance to a new NtsGeometryServices instance with coordinateSequenceFactory set to a DotSpatialAffineCoordinateSequenceFactory. Or you can specify the factory when calling UseNetTopologySuite.

```
// Place this at the beginning of your program to use the specified settings everywhere
(recommended)
GeometryServiceProvider.Instance = new NtsGeometryServices(
    new DotSpatialAffineCoordinateSequenceFactory(Ordinates.XYM),
    new PrecisionModel(PrecisionModels.Floating),
    -1);

// Or specify settings for Npgsql only
conn.TypeMapper.UseNetTopologySuite(
    new DotSpatialAffineCoordinateSequenceFactory(Ordinates.XYM));
```

^{*2:} When Npgsql.EnableLegacyTimestampBehavior is enabled, timestamp without time zone is mapped to Instant instead of LocalDateTime by default.

Reading and Writing Geometry Values

When reading PostGIS values from the database, Npgsql will automatically return the appropriate NetTopologySuite types: Point, LineString, and so on. Npgsql will also automatically recognize NetTopologySuite's types in parameters, and will automatically send the corresponding PostGIS type to the database. The following code demonstrates a roundtrip of a NetTopologySuite Point to the database:

```
var point = new Point(new Coordinate(ld, ld));
conn.ExecuteNonQuery("CREATE TEMP TABLE data (geom GEOMETRY)");
using (var cmd = new NpgsqlCommand("INSERT INTO data (geom) VALUES (@p)", conn))
{
    cmd.Parameters.AddWithValue("@p", point);
    cmd.ExecuteNonQuery();
}

using (var cmd = new NpgsqlCommand("SELECT geom FROM data", conn))
using (var reader = cmd.ExecuteReader())
{
    reader.Read();
    Assert.That(reader[0], Is.EqualTo(point));
}
```

You may also explicitly specify a parameter's type by setting NpgsqlDbType.Geometry.

Geography (geodetic) Support

PostGIS has two types:geometry (for Cartesian coordinates) and geography (for geodetic or spherical coordinates). You can read about the geometry/geography distinction in the PostGIS docs. In a nutshell, geography is much more accurate when doing calculations over long distances, but is more expensive computationally and supports only a small subset of the spatial operations supported by geometry.

Npgsql uses the same NetTopologySuite types to represent both geometry and geography - the Point type represents a point in either Cartesian or geodetic space. You usually don't need to worry about this distinction because PostgreSQL will usually cast types back and forth as needed. However, it's worth noting that Npgsql sends Cartesian geometry by default, because that's the usual requirement. You have the option of telling Npgsql to send geography instead by specifying NpgsqlDbType.Geography:

```
using (var cmd = new NpgsqlCommand("INSERT INTO data (geog) VALUES (@p)", conn))
{
   cmd.Parameters.AddWithValue("@p", NpgsqlDbType.Geography, point);
   cmd.ExecuteNonQuery();
}
```

If you prefer to use geography everywhere by default, you can also specify that when setting up the plugin:

```
{\tt NpgsqlConnection.GlobalTypeMapper.UseNetTopologySuite(geographyAsDefault: true);}
```

GeoJSON (spatial)

Using the GeoJSON plugin is the same as the NetTopologuSuite.

Notes on the Query Builder

- Prefix named parameters with "@".
- Uppercase object names cannot be used, even when enclosed in double quotation marks.
 To use uppercase object names enclosed in double quotation marks, include them in SQL statements and enter these in the [Generate the SQL statements] window rather than in the Query Builder.
- SQL statements cannot be correctly generated if the SQL statement specified in Filter matches any of the conditions below:
 - It uses PostgreSQL intrinsic operators such as << or ::.

- It uses functions with keywords such as AS, FROM, IN, OVER.
 - Example: extract(field from timestamp), RANK() OVER
- It uses functions with the same names as those prescribed in SQL conventions, but that require different arguments.

Notes on metadata

- The CommandBehavior.KeyInfo argument must be specified if executing ExecuteReader before obtaining metadata using GetSchemaTable.

.....



NpgsqlDataReader ndr=cmd.ExecuteReader(CommandBehavior.KeyInfo);
DataTable dt = dr.GetSchemaTable();

Notes on automatically generating update-type SQL statements

- If the SQL statement includes a query (which cannot be updated) that matches any of the conditions below, an update-type SQL statement will be generated (note that it may not be possible to execute this SQL statement in some cases):

.....

- It includes derived tables
- It includes the same column name as the select list

Update-type SQL statements will be automatically generated in the following cases:

- If update statements are obtained using NpgsqlCommandBuilder
- If data is updated using NpgsqlDataAdapter

Notes on distributed transactions

- Applications using transaction scope can use distributed transactions by linking with Microsoft Distributed Transaction Coordinator (MSDTC). In this case, note the following:
 - Ensure that the value of max_prepared_transactions is greater than max_connection, so that "PREPARE TRANSACTION" can be issued for each transaction that simultaneously connects to the database server.
 - If each transaction in the transaction scope accesses the same resource using different connections, the database server will perceive it as requests from different applications, and a deadlock may occur. By configuring a timeout value for the transaction scope beforehand, the deadlock can be broken.

4.5 Uninstallation

This section explains how to uninstall .NET Data Provider and Entity Framework Core.

4.5.1 Uninstalling Npgsql

Npgsql is installed per project. To uninstall it, follow the procedure below:

- Linux

Remove a package reference to the project file.

dotnet remove projectFilePath package Npgsql

W - Windows(R)

L

1. In Visual Studio, open a project for which Npgsql is installed.

- 2. Click [Tools] >> [NuGet Package Manager] >> [Manage NuGet Packages for Solution].
- 3. Select all the projects that have Npgsql installed, and then click [Uninstall]. Alternatively, if the plugin packages have been removed and is no longer installed, in the solution explorer open the Dependencies/NuGet node and delete the plugins that require uninstallation.

4.5.2 Uninstalling .NET Data Provider Type Plugins

The .NET Data Provider type plugins are installed per project. To uninstall it, follow the procedure below:

- Linux

Remove a package reference to the project file (Example: Npgsql.NodaTime).

dotnet remove projectFilePath package Npgsql.NodaTime

W - Windows(R)

- 1. In Visual Studio, open a project for which Npgsql the Plugin to be removed is installed.
- 2. Click [Tools] >> [NuGet Package Manager] >> [Manage NuGet Packages for Solution].
- 3. Select all the projects that have Npgsql Plugin(s) installed, and then click [Uninstall]. Alternatively, if the plugin packages have been removed and are no longer installed, in the solution explorer open the Dependencies/NuGet node and delete the plugins that require uninstallation.

4.5.3 Uninstalling Entity Framework Core

Entity Framework Core is installed per project. To uninstall it, follow the procedure below:

- Linux

Remove a package reference to the project file.

dotnet remove projectFilePath package Npgsql.EntityFrameworkCore.PostgreSQL

W - Windows(R)

- 1. In Visual Studio, open a project for which Entity Framework Core is installed.
- 2. Click [Tools] >> [NuGet Package Manager] >> [Manage NuGet Packages for Solution].
- 3. Select all the projects that have Entity Framework Core installed, and then click [Uninstall].
 Alternatively, if the Entity Framework package has been removed and is no longer installed, in the solution explorer open the Dependencies/NuGet node and delete the packages that require uninstallation.

4.5.4 Uninstalling Npgsql.OpenTelemetry

Npgsql.OpenTelemetry is installed per project. To uninstall it, follow the procedure below:

- Linux

Remove a package reference to the project file.

dotnet remove projectFilePath package Npgsql.OpenTelemetry

W - Windows(R)

- 1. In Visual Studio, open a project for which Npgsql.OpenTelemetry is installed.
- 2. Click [Tools] >> [NuGet Package Manager] >> [Manage NuGet Packages for Solution].
- 3. Select all the projects that have Npgsql.OpenTelemetry installed, and then click [Uninstall]. Alternatively, if the Npgsql.OpenTelemetry package has been removed and is no longer installed, in the solution explorer open the Dependencies/NuGet node and delete the packages that require uninstallation.

Chapter 5 C Library (libpq)

This chapter describes how to use C libraries.

5.1 Development Environment

Install Fujitsu Enterprise Postgres Client Package for the architecture to be developed and executed.



See

Refer to Installation and Setup Guide for Client for information on the C compiler required for C application development.

5.2 Setup

This section describes the environment settings required to use C libraries and how to encrypt data for communication.

5.2.1 Environment Settings

To execute an application that uses libpq, set the environment variable as shown below.

L

Linux

- Required for execution of the application
 - PGLOCALEDIR fujitsuEnterprisePostgresClientInstallDin/share/locale



"<x>" indicates the product version.

> PGLOCALEDIR=/opt/fsepv<x>client64/share/locale;export PGLOCALEDIR

W

Windows (R)

- Required for compile/link
 - LIB

fujitsuEnterprisePostgresClientInstallDir\lib

INCLUDE

fujitsuEnterprisePostgresClientInstallDir\include

- Required for execution of the application
 - PATH

fujitsuEnterprisePostgresClientInstallDir\lib

- PGLOCALEDIR



When the 32-bit version client package is installed on a 64-bit operating system. "<x>" indicates the product version.

- > SET PATH=%ProgramFiles(x86)%\Fujitsu\fsepv<x>client32\lib;%PATH%
- > SET LIB=%ProgramFiles(x86)%\Fujitsu\fsepv<x>client32\lib;%LIB%

5.2.2 Message Language and Encoding System Used by Applications Settings

This section explains the language settings for the application runtime environment and the encoding settings for the application.

Language settings

You must match the language settings for the application runtime environment with the message locale settings of the database server.

Messages output by an application may include text from messages sent from the database server. In the resulting text, the text of the application message will use the message locale of the application, and the text of the message sent by the database server will use the message locale of the database server. If the message locales do not match, more than one language or encoding system will be used. Moreover, if the encoding systems do not match, characters in the resulting text can be garbled.

- Linux

Set the locale for messages (LC_MESSAGES category) to match the message locale of the database server. This can be done in a few different ways, such as using environment variables. Refer to the relevant manual of the operating system for information on the setlocale function.



Example of specifying "en_US.UTF-8" with the setlocale function

setlocale(LC_ALL, "en_US.UTF-8");

Specifying the locale of the LC_ALL category propagates the setting to LC_MESSAGE.

Windows(R)

Align the locale of the operating system with the message locale of the database server.

Encoding System Settings

Ensure that the encoding system that is embedded in the application and passed to the database, and the encoding system setting of the runtime environment, are the same. The encoding system cannot be converted correctly on the database server.

Use one of the following methods to set the encoding system for the application:

- Set the PGCLIENTENCODING environment variable in the runtime environment.
- Set the client_encoding keyword in the connection string.
- Use the PQsetClientEncoding function.



Refer to "Supported Character Sets" in "Server Administration" in the PostgreSQL Documentation for information on the strings that represent the encoding system that can be set.

For example, when using "Unicode" and "8 bit", set the string "UTF8".

W



Text may be garbled when outputting results to the command prompt. Review the font settings for the command prompt if this occurs.

5.2.3 Settings for Encrypting Communication Data

Set in one of the following ways when performing remote access using communication data encryption:

When setting from outside with environment variables

Specify "require", "verify-ca", or "verify-full" in the PGSSLMODE environment variable.

In addition, the parameters for the PGSSLROOTCERT and PGSSLCRL environment variables need to be set to prevent spoofing of the database server.



See

Refer to "Environment Variables" in "Client Interfaces" in the PostgreSQL Documentation for information on environment variables.

When specifying in the connection URI

Specify "require", "verify-ca", or "verify-full" in the "sslmode" parameter of the connection URI.

In addition, the parameters for the sslcert, sslkey, sslrootcert, and sslcrl need to be set to prevent spoofing of the database server.



Refer to "Secure TCP/IP Connections with SSL" in "Server Administration" in the PostgreSQL Documentation for information on encrypting communication data.

5.3 Connecting with the Database



Point

Use the connection service file to specify the connection destination. In the connection service file, a name (service name) is defined as a set, comprising information such as connection destination information and various types of tuning information set for connections. By using the service name defined in the connection service file when connecting to databases, it is no longer necessary to modify applications when the connection information changes.

Refer to "Client Interfaces", "The Connection Service File" in the PostgreSQL Documentation for details.



See

Refer to "Database Connection Control Functions" in "Client Interfaces" in the PostgreSQL Documentation.

In addition, refer to "6.3 Connecting with the Database" in "Embedded SQL in C " for information on connection string.

5.4 Application Development



Refer to "libpq - C Library" in "Client Interfaces" in the PostgreSQL Documentation for information on developing applications.

However, if you are using the C library, there are the following differences to the PostgreSQL C library (libpq).

5.4.1 Compiling Applications

Specify the following paths when compiling applications.

Refer to your compiler documentation for information on how to specify the path.

Also, when using a shared library, refer to "A.6 How to Build and Run an Application that Uses Shared Libraries".

Linux

L

W

Table 5.1 Include file and library path

Type of path	Path name	
Path of the include file	fujitsuEnterprisePostgresClientInstallDir/include	
Path of the library	fujitsuEnterprisePostgresClientInstallDir/lib	

Table 5.2 C Library (libpg library)

Type of library	Library name	
Dynamic library	libpq.so	
Static library	libpq.a	

W - Windows(R)

If the include file and the library path have been set in the environment variable, there is no need to specify the paths shown below for the compile.

Table 5.3 Include file and library path

Type of path	Path name
Path of the include file	fujitsuEnterprisePostgresClientInstallDir\include
Path of the library	fujitsuEnterprisePostgresClientInstallDir\lib

Table 5.4 C Library (libpg library)

Type of library	Library name
Library for links	libpq.lib
Dynamic library	libpq.dll

5.4.2 Creating Applications while in Database Multiplexing Mode

This section explains points to consider when creating applications while in database multiplexing mode.



- Refer to the Cluster Operation Guide (Database Multiplexing) for information on database multiplexing mode.
- Refer to "Application Development" in the Cluster Operation Guide (PRIMECLUSTER) for points to consider when creating applications using the failover feature integrated with the cluster software.

5.4.2.1 Errors when an Application Connection Switch Occurs and Corresponding Actions

If an application connection switch occurs while in database multiplexing mode, explicitly close the connection and then reestablish the connection or reexecute the application.

The table below shows errors that may occur during a switch, and the corresponding action to take.

State		Error information	Action
Server failure or Fujitsu Enterprise Postgres system failure	Failure occurs during access	PGRES_FATAL_ERRO R(*1) 57P01(*2) NULL(*2)	After the switch is complete, reestablish the connection, or reexecute the
	Accessed during system failure	CONNECTION_BAD(* 3)	application.
Switch to the standby server	Switched during access	PGRES_FATAL_ERRO R(*1) 57P01(*2) NULL(*2)	
	Accessed during switch	CONNECTION_BAD(* 3)	

^{*1:} Return value of PQresultStatus().

^{*2:} Return value of PQresultErrorField() PG_DIAG_SQLSTATE.

^{*3:} Return value of PQstatus().

Chapter 6 Embedded SQL in C

This chapter describes application development using embedded SQL in C.

6.1 Development Environment

Install Fujitsu Enterprise Postgres Client Package for the architecture to be developed and executed.



See

Refer to Installation and Setup Guide for Client for information on the C compiler required for C application development.



C++ is not supported. Create a library by implementing embedded SQL in C, and call it from C++.

6.2 Setup

6.2.1 Environment Settings

When using embedded SQL in C, the same environment settings as when using the C library (libpq) are required.

Refer to "5.2.1 Environment Settings" in "C Library (libpq)" for information on the environment settings for the library for C.

Additionally, set the following path for the precompiler ecpg in the PATH environment variable:



Linux

fujitsuEnterprisePostgresClientInstallDir/bin



Windows(R)

fujitsuEnterprisePostgresClientInstallDir\bin

6.2.2 Message Language and Encoding System Used by Applications Settings

The message language and the encoding System Settings Used by Applications settings are the same as when using the library for C.

However, in embedded SQL, the PQsetClientEncoding function cannot be used in the encoding system settings. In embedded SQL, use the SET command to specify the encoding system in client_encoding.

Refer to "5.2.2 Message Language and Encoding System Used by Applications Settings" in "C Library (libpq)" for information on the settings for the library for C.

6.2.3 Settings for Encrypting Communication Data

When encrypting the communication data, the same environment settings as when using the C library (libpq) are required.

Refer to "5.2.3 Settings for Encrypting Communication Data" in "C Library (libpq)" for information on the environment settings for the C library.

6.3 Connecting with the Database



- It is recommended to use a connection service file to specify connection destinations. In the connection service file, a name (service name) is defined as a set, comprising information such as connection destination information and various types of tuning information set for connections. By using the service name defined in the connection service file when connecting to databases, it is no longer necessary to modify applications when the connection information changes. Refer to "The Connection Service File" in "Client Interfaces" in the PostgreSQL Documentation for information.
- If using a connection service file, perform either of the procedures below:
 - Set the service name as a string literal or host variable, as follows:
 tcp:postgresql://?service=my_service
 - Set the service name in the environment variable PGSERVICE, and use CONNECT TO DEFAULT

Use the CONNECT statement shown below to create a connection to the database server.

Format

EXEC SQL CONNECT TO target [AS connection-name] [USER user-name];

target

Write in one of the following formats:

- dbname@host:port
- tcp:postgresql://host:port/dbname[?options]
- unix:postgresql://host[:port][/dbname][?options]
 (Definition method when using the UNIX domain socket)
- SQL string literal containing one of the above formats
- Reference to a character variable containing one of the above formats
- DEFAULT

user-name

Write in one of the following formats:

- username
- username/password
- username IDENTIFIED BY password
- username USING password

Description of the arguments

Argument	Description		
dbname	Specify the database name.		
host	Specify the host name for the connection destination.		
port	Specify the port number for the database server.		
	The default is "27500".		
connection-name	Specify connection names to identify connections when multiple connections are to be processed within a single program.		

Argument	Description		
username	Specify the user that will connect with the database.		
	If this is omitted, the name used will be that of the user on the operating system that is executing the application.		
password	Specify a password when authentication is required.		
options	Specify the following parameter when specifying a time for timeout. Connect parameters with & when specifying more than one. The following shows the values specified for each parameter.		
	- connect_timeout		
	Specify the timeout for connections.		
	Specify a value between 0 and 2147483647 (in seconds). There is no limit set if you set 0 or an invalid value. If "1" is specified, the behavior will be the same as when "2" was specified. An error occurs when a connection cannot be established within the specified time.		
	- keepalives		
	This enables keepalive.		
	Keepalive is disabled if 0 is specified. Keepalive is enabling when any other value is specified. The default is keepalive enabled. Keepalive causes an error to occur when it is determined that the connection with the database is disabled.		
	- keepalives_idle		
	Specify the time until the system starts sending keepalive messages when communication with the database is not being performed.		
	- Linux		
	Specify a value between 1 and 32767 (in seconds). The default value of the system is used if this is not specified.		
	- Windows(R)		
	Specify a value between 1 and 2147483647 (in seconds). 7200 will be set as default if a value outside this range is specified or if nothing is specified.		
	- keepalives_interval		
	Specify the interval between resends when there is no response to keepalive messages.		
	- Linux		
	Specify a value between 1 and 32767 (in seconds). The default value of the system is used if this is not specified.		
	- Windows(R)		
	Specify a value between 1 and 2147483647 (in seconds). 1 will be set as default if a value outside this range is specified or if nothing is specified.		
	- keepalives_count		
	Specify the number of resends for keepalive messages.		
	- Linux		
	Specify a value between 1 and 127. The default value of the system is used if this is not specified.		
	- Windows(R)		
	The system default value is used irrespective of what is specified for this parameter.		
	- tcp_user_timeout		

Argument	Description
	After establishing the connection, when sending from the client to the server, if the TCP resend process operates, specify the time until it is considered to be disconnected.
	- Linux
	Specify a value between 0 and 2147483647 (in millseconds). The default value of the system is used if 0. 0 will be set as default if nothing is specified.
	- Windows(R)
	Cannot be specified.





If a value other than 0 is specified for the tcp_user_timeout parameter, the waiting time set by the tcp_keepalives_idle parameter and tcp_keepalives_interval parameter will be invalid and the waiting time specified by the tcp_user_timeout parameter will be used.

Code examples for applications

EXEC SQL CONNECT TO tcp:postgresql://sv1:27500/mydb?
connect_timeout=20&keepalives_idle=20&keepalives_interval=5&keepalives_count=2&keepalives=
1 USER myuser/myuser01;

6.4 Application Development

Refer to "ECPG - Embedded SQL in C" in "Client Interfaces" in the PostgreSQL Documentation for information on developing applications.

However, when using embedded SQL in C, there are the following differences to the embedded SQL (ECPG) in PostgreSQL C

6.4.1 Support for National Character Data Types

This section describes how to use the national character data types using the SQL embedded C preprocessor.

The following explains the C language variable types corresponding to the NCHAR type:

Specify the number of characters specified for the NCHAR type multiple by 4, plus 1 for the length of the host variable.

Data Type	Host variable type	
NATIONAL CHARACTER(n)	NCHAR variable name [nx4+1]	
NATIONAL CHARACTER VARYING(n)	NVARCHAR variable name [nx4+1]	



See

Refer to "Handling Character Strings" in "Client Interfaces" in the PostgreSQL documentation for information on using character string types.

6.4.2 Compiling Applications

Append the extension "pgc" to the name of the source file for the embedded SQL in C.

When the pgc file is precompiled using the ecpg command, C source files will be created, so use the C compiler for the compile.

Precompiling example

ecpg testproc.pgc

If an optimizer hint block comment is specified for the SQL statement, specify the following option in the ecpg command:

--enable-hint

Enables the optimizer hint block comment (hereafter, referred to as the "hint clause"). If this option is not specified, the hint clause will be removed as a result of the ecpg precompile and be disabled.

The SQL statements that can be specified in the hint clause are SELECT, INSERT, UPDATE, and DELETE.

The locations in which the hint clause can be specified are immediately after one of the SELECT, INSERT, UPDATE, DELETE, or WITH keywords. A syntax error will occur if any other location is specified.

Example of specifying the hint clause

EXEC SQL SELECT /*+ IndexScan(prod ix01) */ name_id INTO :name_id FROM prod WHERE id = 1;



See

For basic usage of pg_hint_plan and pg_dbms_stats, see below.

Control execution plans with pg_hint_plan
 https://www.postgresql.fastware.com/postgresql-insider-tun-hint-plan



Take the following points into account when using embedded SQL source files:

- Multibyte codes expressed in SJIS or UTF-16 cannot be included in statements or host variable declarations specified in EXEC SQL.
- Do not use UTF-8 with a byte order mark (BOM), because an error may occur during compilation if the BOM character is incorrectly recognized as the source code.
- Multibyte characters cannot be used in host variable names.
- It is not possible to use a TYPE name that contains multibyte characters, even though it can be defined.

Specify the following paths when compiling a C application output with precompiling.

Refer to your compiler documentation for information on how to specify the path.

Also, when using a shared library, refer to "A.6 How to Build and Run an Application that Uses Shared Libraries".

Linux

Table 6.1 Include file and library path

	- 7 1	
Type of path	ath Path name	
Path of the include file	fujitsuEnterprisePostgresClientInstallDir/include	
Path of the library	fujitsuEnterprisePostgresClientInstallDir/lib	

L

Table 6.2 C Library

Type of library	Library name	Note
Dynamic library	libecpg.so	
	libpgtypes.so	When using the pgtypes library
Static library	libecpg.a	
	libpgtypes.a	When using the pgtypes library

W

Windows(R)

If the include file and the library path have been set in the environment variable, there is no need to specify the paths shown below for the compile.

W

Table 6.3 Include file and library path

Table die molade me am	Table die Heldde me did nerdly path		
Type of path	Path name		
Path of the include file	fujitsuEnterprisePostgresClientInstallDir include		
Path of the library	fujitsuEnterprisePostgresClientInstallDir lib		

W

Table 6.4 C Library

Type of library	Library name	Note
Library for links	libecpg.lib	
	libpgtypes.lib	When using the pgtypes library
Dynamic library	libecpg.dll	
	libpgtypes.dll	When using the pgtypes library



Note

- The libecpg library in Windows(R) is created by "release" and "multithreaded" options. When using the ECPGdebug function included in this library, compile using the "release" and "multithreaded" flags in all programs that use this library. When you do this, use the "dynamic" flag if you are using libecpg.dll, and use the "static" flag if you are using libecpg.lib.

Refer to "Library Functions" in "Client Interfaces" in the PostgreSQL Documentation for information on the ECPGdebug function.

- The cl command expects input to be a program that uses one of the following code pages, so convert the program to these code pages and then compile and link it (refer to the Microsoft documentation for details).
 - ANSI console code pages (example: Shift-JIS for Japanese)
 - UTF-16 little-endian with or without BOM (Byte Order Mark)
 - UTF-16 big-endian with or without BOM
 - UTF-8 with BOM

The cl command converts strings in a program to an ANSI console code page before generating a module, so the data sent to and received from the database server becomes an ANSI console code page. Therefore, set the coding system corresponding to the ANSI console code page as the coding system of the client.

Refer to "Character Set Support" in "Server Administration" in the PostgreSQL Documentation for information on how to set the client encoding system.

(Example: To use environment variables in Japanese, set SJIS in PGCLIENTENCODING.)

6.4.3 Bulk INSERT

This section describes the bulk INSERT.

Synopsis

```
EXEC SQL [ AT conn ] [ FOR { numOfRows | ARRAY_SIZE } ]
    INSERT INTO tableName [ ( colName [, ...] ) ]
    { VALUES ( { expr | DEFAULT } [, ...] ) [, ...] | query }
    [ RETURNING * | outputExpr [ [ AS ] outputName ] [, ...]
    INTO outputHostVar [ [ INDICATOR ] indicatorVar ] [, ...] ];
```

Description

Bulk INSERT is a feature that inserts multiple rows of data in bulk.

By specifying the array host variable that stored the data in the VALUES clause of the INSERT statement, the data for each element in the array can be inserted in bulk. This feature is used by specifying the insertion count in the FOR clause immediately before the INSERT statement.

FOR Clause

Specify the insertion count using numOfRows or ARRAY_SIZE in the FOR clause. The FOR clause can be specified only in the INSERT statement, not in other update statements.

numOfRows and ARRAY_SIZE

Insertion processing will be executed only for the specified count. However, if the count is 1, it will be assumed that the FOR clause was omitted when the application is executed. In this case, proceed according to the INSERT specification in the PostgreSQL Documentation.

Specify the FOR clause with a short, int, or long long host variable or with a literal.

Specify ARRAY_SIZE to insert all elements of the array in the table. When specifying ARRAY_SIZE, specify at least one array in *expr*.

If two or more arrays were specified in expr, it will be assumed that ARRAY_SIZE is the minimum number of elements in the array.

numOfRows or ARRAY_SIZE must exceed the minimum number of elements in all arrays specified in expr, outputHostVar, and indicatorVal.

The following example shows how to specify the FOR clause.

expr

Specify the value to be inserted in the table. Array host variables, host variable literals, strings, and pointer variables can be specified. Structure type arrays and pointer variable arrays cannot be specified.

Do not use pointer variables and ARRAY_SIZE at the same time. The reason for this is that the number of elements in the area represented by the pointer variable cannot be determined.

query

A query (SELECT statement) that supplies the rows to be inserted. The number of rows returned by *query* must be 1. If two or more rows are returned, an error will occur. This cannot be used at the same time as ARRAY_SIZE.

outputHostVar, indicatorVal

These must be array host variables or pointer variables.

Error Messages

Given below are the error messages that are output when bulk INSERT functionality is not used correctly.

Message

invalid statement name "FOR value should be positive integer"

Cause

The value given for *numOfRows* is less than or equal to 0.

Solution

Specify a value that is more than or equal to 1 for numOfRows.

Message

invalid statement name "Host array variable is needed when using FOR ARRAY_SIZE"

Cause

A host array is not specified in the values clause when using the ARRAY_SIZE keyword.

Solution

At least one host array variable should be included in the values clause

Message

SELECT...INTO returns too many rows

Cause

The number of rows returned by the 'SELECT ... INTO' query in the INSERT statement is more than one.

Solution

When the value of *numOfRows* is more than one, the maximum number of rows that can be returned by the 'SELECT ... INTO' query in the INSERT statement is one.

Limitations

The limitations when using bulk INSERT are given below.

- Array of structures should not be used as an input in the 'VALUES' clause. Attempted use will result in junk data being
 inserted into the table.
- Array of pointers should not be used as an input in the 'VALUES' clause. Attempted use will result in junk data being inserted into the table.
- ECPG supports the use of 'WITH' clause in single INSERT statements. 'WITH' clause cannot be used in bulk INSERT statements
- ECPG does not calculate the size of the pointer variable. So when a pointer variable is used that includes multiple elements, *numOfRows* should be less than or equal to the number of elements in the pointer. Otherwise, junk data will be inserted into the table.
- If an error occurs, all bulk INSERT actions will be rolled back, therefore, no rows are inserted. However, if the RETURNING clause was used, and the error occurred while obtaining the rows after the insertion was successful, the insertion processing will not be rolled back.

Samples

Given below are some sample usages of the bulk INSERT functionality.

Basic Bulk INSERT

```
int in_f1[4] = {1,2,3,4};
...
EXEC SQL FOR 3 INSERT INTO target (f1) VALUES (:in_f1);
```

The number of rows to insert indicated by the FOR clause is 3, so the data in the first 3 elements of the host array variable are inserted into the table. The contents of the target table will be:

```
f1
----
1
2
3
(3 rows)
```

Also a host integer variable can be used to indicate the number of rows that will be inserted in FOR clause, which will produce the same result as above:

```
int num = 3;
int in_f1[4] = {1,2,3,4};
...
EXEC SQL FOR :num INSERT INTO target (f1) VALUES (:in_f1);
```

Inserting constant values

Constant values can also be bulk INSERTed into the table as follows:

```
EXEC SQL FOR 3 INSERT INTO target (f1,f2) VALUES (DEFAULT, 'hello');
```

Assuming the 'DEFAULT' value for the 'f1' column is '0', the contents of the target table will be:

Using ARRAY_SIZE

'FOR ARRAY_SIZE' can be used to insert the entire contents of a host array variable, without explicitly specifying the size, into the table.

```
int in_f1[4] = {1,2,3,4};
...
EXEC SQL FOR ARRAY_SIZE INSERT INTO target (f1) VALUES (:in_f1);
```

In the above example, four rows are inserted into the table.



If there are multiple host array variables specified as input values, then the number of rows inserted is same as the smallest array size. The example given below demonstrates this usage.

```
int in_f1[4] = {1,2,3,4};
char in_f3[3][10] = {"one", "two", "three"};
```

```
...
EXEC SQL FOR ARRAY_SIZE INSERT INTO target (f1,f3) VALUES (:in_f1,:in_f3);
```

In the above example, the array sizes are 3 and 4. Given that the smallest array size is 3, only three rows are inserted into the table. The table contents are given below.

Using Pointers as Input

Pointers that contain multiple elements can be used in bulk INSERT.

```
int *in_pf1 = NULL;
in_pf1 = (int*)malloc(4*sizeof(int));
in_pf1[0]=1;
in_pf1[1]=2;
in_pf1[2]=3;
in_pf1[3]=4;
...
EXEC SQL FOR 4 INSERT INTO target (f1) values (:in_pf1);
```

The above example will insert four rows into the target table.

Using SELECT query

When using bulk INSERT, the input values can be got from the results of a SELECT statement. For example,

```
EXEC SQL FOR 4 INSERT INTO target(f1) SELECT age FROM source WHERE name LIKE 'foo';
```

Assuming that the 'SELECT' query returns one row, the same row will be inserted into the target table four times.



If the 'SELECT' query returns more than one row, the INSERT statement will throw an error.

```
EXEC SQL FOR 1 INSERT INTO target(f1) SELECT age FROM source;
```

In the above example, all the rows returned by the 'SELECT' statement will be inserted into the table. In this context '1' has the meaning of 'returned row equivalent'.

Using RETURNING clause

Bulk INSERT supports the same RETURNING clause syntax as normal INSERT. An example is given below.

```
int out_f1[4];
int in_f1[4] = {1,2,3,4};
...
EXEC SQL FOR 3 INSERT INTO target (f1) VALUES (:in_f1) RETURNING f1 INTO :out_f1;
```

After the execution of the above INSERT statement, the 'out_f1' array will have 3 elements with the values of '1','2' and '3'.

6.4.4 Creating Applications while in Database Multiplexing Mode

This section explains points to consider when creating applications while in database multiplexing mode.



- Refer to the Cluster Operation Guide (Database Multiplexing) for information on database multiplexing mode.
- Refer to "Application Development" in the Cluster Operation Guide (PRIMECLUSTER) for points to consider when creating applications using the failover feature integrated with the cluster software.

......

6.4.4.1 Errors when an Application Connection Switch Occurs and Corresponding Actions

If an application connection switch occurs while in database multiplexing mode, explicitly close the connection and then reestablish the connection or reexecute the application.

The table below shows errors that may occur during a switch, and the corresponding action to take.

State		Error information (*1)	Action
Server failure or Fujitsu Enterprise Postgres system failure	Failure occurs during access	57P01 57P02 YE000 26000 40001	After the switch is complete, reestablish the connection, or reexecute the application.
	Accessed during node/system failure	08001	
Switch to the standby server	Switched during access	57P01 57P02 YE000 26000 40001	
	Accessed during switch	08001	

^{*1:} Return value of SQLSTATE.

6.4.5 Notes

Notes on creating multithreaded applications

In embedded SQL in C, DISCONNECT ALL disconnects all connections within a process, and therefore it is not thread-safe in all operations that use connections. Do not use it in multithreaded applications.

Chapter 7 Embedded SQL in COBOL

This chapter describes application development using embedded SQL in COBOL.

7.1 Development Environment

Install Fujitsu Enterprise Postgres Client Package for the architecture to be developed and executed.



See

Refer to the Installation and Setup Guide for Client for information on the COBOL compiler required for COBOL application development.

7.2 Setup

7.2.1 Environment Settings

When using embedded SQL in COBOL, the same environment settings as when using the C library (libpq) are required. Refer to "5.2.1 Environment Settings" in "C Library (libpq)" for information on the environment settings for the library for C.

Additionally, set the following path for the precompiler ecobpg in the PATH environment variable:



Linux

fujitsuEnterprisePostgresClientInstallDir/bin



Windows(R)

fujitsuEnterprisePostgresClientInstallDir\bin

7.2.2 Message Language and Encoding System Used by Applications

The settings for the message language and the encoding system used by applications should be the same as those required when using the library for C.

However, in embedded SQL, the PQsetClientEncoding function cannot be used in the encoding system settings. In embedded SQL, use the SET command to specify the encoding system in client_encoding.

Refer to "5.2.2 Message Language and Encoding System Used by Applications Settings" in "C Library (libpq)" for information on the settings for the library for C.

7.2.3 Settings for Encrypting Communication Data

When encrypting the communication data, the same environment settings as when using the C library (libpq) are required.

Refer to "5.2.3 Settings for Encrypting Communication Data" in "C Library (libpq)" for information on the environment settings for the C library.

7.3 Connecting with the Database

Use the CONNECT statement shown below to create a connection to the database server.

Format

EXEC SQL CONNECT TO target [AS connection-name] [USER user-name]END-EXEC.

target

Write in one of the following formats:

- dbname@host:port
- tcp:postgresql://host:port/dbname[?options]
- unix:postgresql://host[:port][/dbname][?options]
 (Definition method when using the UNIX domain socket)
- SQL string literal containing one of the above formats
- Reference to a character variable containing one of the above formats
- DEFAULT



If target is DEFAULT, the AS clause and USER clause cannot be specified.

user-name

Write in one of the following formats:

- username
- username/password
- username IDENTIFIED BY password
- username USING password

Description of the arguments

Argument	Description			
dbname	Specify the database name.			
host	Specify the host name for the connection destination.			
port	Specify the port number for the database server.			
	The default is "27500".			
connection-name	Specify connection names to identify connections when multiple connections are to be processed within a single program.			
username	Specify the user that will connect with the database.			
	If this is omitted, the name used will be that of the user on the operating system that is executing the application.			
password	Specify a password when authentication is required.			
options	Specify the following parameter when specifying a time for timeout. Connect parameters with & when specifying more than one. The following shows the values specified for each parameter.			
	- connect_timeout			
	Specify the timeout for connections.			
	Specify a value between 0 and 2147483647 (in seconds). There is no limit set if you set 0 or an invalid value. If "1" is specified, the behavior will be the same as when "2" was specified. An error occurs when a connection cannot be established within the specified time.			
	- keepalives			
	This enables keepalive.			

Argument	Description		
	Keepalive is disabled if 0 is specified. Keepalive is enabled when any other value is specified. The default is keepalive enabled. Keepalive causes an error to occur when it is determined that the connection with the database is disabled.		
	- keepalives_idle		
	Specify the time until the system starts sending keepalive messages when communication with the database is not being performed.		
	- Linux		
	Specify a value between 1 and 32767 (in seconds). The default value of the system is used if this is not specified.		
	- Windows(R)		
	Specify a value between 1 and 2147483647 (in seconds). 7200 will be set as default if a value outside this range is specified or if nothing is specified.		
	- keepalives_interval		
	Specify the interval between resends when there is no response to keepalive messages.		
	- Linux		
	Specify a value between 1 and 32767 (in seconds). The default value of the system is used if this is not specified.		
	- Windows(R)		
	Specify a value between 1 and 2147483647 (in seconds). 1 will be set as default if a value outside this range is specified or if nothing is specified.		
	- keepalives_count		
	Specify the number of resends for keepalive messages.		
	- Linux		
	Specify a value between 1 and 127. The default value of the system is used if this is not specified.		
	- Windows(R)		
	The system default value is used irrespective of what is specified for this parameter.		

Code examples for applications

EXEC SQL CONNECT TO tcp:postgresql://sv1:27500/mydb? connect_timeout=20&keepalives_idle=20&keepalives_interval=5&keepalives_count=2&keepalives=1 USER myuser/myuser01 END-EXEC.

7.4 Application Development

Refer to "Appendix D ECOBPG - Embedded SQL in COBOL" for information on developing applications.

7.4.1 Support for National Character Data Types

This section describes how to use the national character data types using the SQL embedded COBOL preprocessor.

The table below lists the COBOL variable types supporting the national character data types. The number of characters specified for the national character data type must be specified for the length of the host variable.

National character data type	COBOL variable type
CHARACTER(n)	varName PIC N(n)

National character data type	COBOL variable type
NATIONAL CHARACTER(n)	
CHARACTER VARYING(n)	varName PIC N(n) VARYING
NATIONAL CHARACTER VARYING(n)	

To use COBOL variable types that support national character data types, it is necessary to specify the ECOBPG_NCHAR environment variable.

ECOBPG_NCHAR={ UTF16LE | UTF16BE | UTF32LE | UTF32BE | SJIS }

In SQL embedded COBOL, specify the encoding of the COBOL variable types that support national character data types.

- UTF16LE: UTF-16 little-endian

- UTF16BE: UTF-16 big-endian

- UTF32LE: UTF-32 little-endian

- UTF32BE: UTF-32 big-endian

- SJIS: Shift JIS

If this environment variable is omitted, the encoding will be determined according to the encoding system of the client.

- If UTF8 is used: UTF16 (endians will be encoded in accordance with endians of the client system)

- If SJIS is used: SJIS

If encoding is specified for the translation option when compiling with NetCOBOL, the encoding specified for the national character data types should be used for the environment variable ECOBPG_NCHAR.

The list below shows NetCOBOL translation options and their corresponding environment variable ECOBPG_NCHAR values.

NetCOBOL translation options	Environment variable ECOBPG_NCHAR	
ENCODE (UTF-8,UTF16,LE)	UTF-16LE	
RCS (UTF-16,LE)		
ENCODE (UTF-8,UTF-16,BE)	UTF-16BE	
RCS (UTF-16,BE)	UIF-10BE	
ENCODE (UTF-8,UTF-32,LE)	UTF-32LE	
ENCODE (UTF-8,UTF-32,BE)	UTF-32BE	
ENCODE (SJIS,SJIS)	SJIS	
Not specified	No need to specify	

Also, if the post-compiling encoding for an application differs from the locale of the execution environment, then the client encoding must be used for the application.

The list below shows the values supported for the combinations of application encoding, locale of the execution environment, and client encodings.

Application encoding	Locale used when executing an application	Client encoding
UTF-8	UTF-8	UTF-8
U1F-8	SJIS	UTF-8
SJIS	UTF-8	SJIS

Application encoding	Locale used when executing an application	Client encoding
	SJIS	SJIS

Refer to "7.2.2 Message Language and Encoding System Used by Applications" for information on how to set client encoding systems.

The following example shows host variable declaration of a national character data type.

```
01 DATA1 PIC N(10).
01 DATA2 PIC N(10) VARYING.
```



- Halfwidth characters should not be used for the national character data type COBOL variable.
- The national character data type column attribute obtained by applications should be the CHAR type.
- Encoding cannot be specified using the ENCODING clause, which is a feature of NetCOBOL.

7.4.2 Compiling Applications

Append the extension "pco" to the name of the source file for the embedded SQL in COBOL.

When the pco file is precompiled using the ecobpg command, COBOL source files will be created, so use the COBOL compiler for the compile.

Precompiling example

```
ecobpg testproc.pco
```

For COBOL code notation, "fixed" or "variable" format can be specified as an ecobpg command option. If not specified, "fixed" format is used.

Refer to "D.1 Precautions when Using Functions and Operators" and "D.12.1 ecobpg" for information on COBOL code notation and how to specify options for ecobpg.

If an optimizer hint block comment is specified for the SQL statement, specify the following option in the ecobpg command:

--enable-hint

Enables the optimizer hint block comment (hereafter, referred to as the "hint clause"). If this option is not specified, the hint clause will be removed as a result of the ecobpg precompile and be disabled.

The SQL statements that can be specified in the hint clause are SELECT, INSERT, UPDATE, and DELETE.

The locations in which the hint clause can be specified are immediately after one of the SELECT, INSERT, UPDATE, DELETE, or WITH keywords. A syntax error will occur if it is specified in any other location.

Example of specifying the hint clause

```
EXEC SQL SELECT /*+ IndexScan(prod ix01) */ name_id
INTO :name_id FROM prod WHERE id = 1 END-EXEC.
```



For basic usage of pg_hint_plan and pg_dbms_stats, see below.

Control execution plans with pg_hint_plan
 https://www.postgresql.fastware.com/postgresql-insider-tun-hint-plan

If the encoding used for embedded SQL source files differs from that of the locale when precompiling was executed, set the encoding for the embedded SQL source files by specifying the following option for ecobpg.

-E-encode

Specify "UTF8", "SJIS", or "EUC_JP".

If this option is omitted, the encoding is processed based on the locale.

Path of the library file

The ecobpg command defines a group item "sqlca_t" to handle errors, which is defined in the library file stored in the following path:

Linux

Library file name	The storage destination of library file
SQLCA-COBOL.cob	fujitsuEnterprisePostgresClientInstallDir/include

Windows(R)

Library file name	The storage destination of library file
SQLCA-COBOL.cob	fujitsuEnterprisePostgresClientInstallDir\include

When the ecobpg command generates a COBOL file, it inserts a COPY statement with no options to copy the library file. Therefore, specify the path of the storage destination of library file when compiling. How to specify the path must conform to your compiler's specifications.

There is also a library file with the same contents without the extension"cob".



Refer "D.7.2 sqlca" for information on the sqlca_t.

Libraries to use

The applications generated by ecobpg connect to PostgreSQL through the ECPG library. The ECPG library internally loads the libpq library.

Path of library

Refer to "6.4.2 Compiling Applications" for information on the location and name of the ECPG library. And Refer to "Chapter 5 C Library (libpq)" for information on the location and name of the libpq library.

Also, when using a shared library, refer to "A.6 How to Build and Run an Application that Uses Shared Libraries".

The COBOL compiler provides the how to link various libraries, so be sure to specify the path and libraries according to the specifications of your compiler.

Entry information of subprogram

If you use the ECPG library with a dynamic program structure, copy the entry information stored below. For details, follow the specifications of your compiler.

Linux

fujitsuEnterprisePostgresClientInstallDir/share/cobol_entry.info

Windows(R)

fujitsuEnterprisePostgresClientInstallDir\share\cobol_entry.info



The examples of compiling the applications that dynamically links the COBOL language library. Note that "<*x*>" indicates the product version.

Linux

- Linux 64-bit application:

```
cobol -M -o testproc -I/opt/fsepv<x>client64/include -L/opt/fsepv<x>client64/lib -
lecpg -lpq testproc.cob
```

- Linux 64-bit application(When to set DT_RUNPATH):

```
cobol -c -M -o testproc.o -I/opt/fsepv<x>client64/include testproc.cob
ld -dynamic-linker /lib64/ld-linux-x86-64.so.2 -rpath /opt/fsepv<x>client64/lib --
enable-new-dtags -o testproc /usr/lib64/crti.o /usr/lib64/crt1.o /usr/lib64/crtn.o
testproc.o -L/opt/fsepv<x>client64/lib -lecpg -lpq -Bdynamic -L/opt/FJSVcbl64/lib -
lrcobol -ldl -lc
```

Windows(R)

The examples of compiling on a 64-bit operating system.

- 64-bit application

```
> SET LIB=%ProgramFiles%\Fujitsu\fsepv<x>client64\lib;%LIB%
> cobol -I "%ProgramFiles%\Fujitsu\fsepv<x>client64\include" -M testproc.cob
> link testproc.obj F4AGCIMP.LIB LIBCMT.LIB LIBECPG.LIB LIBPQ.LIB /OUT:testproc.exe
```

- 32-bit application

[NetCOBOL V10.5 or earlier]

```
> SET LIB=%ProgramFiles(x86)%\Fujitsu\fsepv<x>client32\lib;%LIB%
> cobol32 -I "%ProgramFiles(x86)%\Fujitsu\fsepv<x>client32\include" -M testproc.cob
> link testproc.obj LIBC.LIB F3BICIMP.LIB LIBECPG.LIB LIBPQ.LIB /OUT:testproc.exe
```

[NetCOBOL V11.0 or later]

```
> SET LIB=%ProgramFiles(x86)%\Fujitsu\fsepv<x>client32\lib;%LIB%
> cobol32 -I "%ProgramFiles(x86)%\Fujitsu\fsepv<x>client32\include" -M testproc.cob
> link testproc.obj MSVCRT.LIB F3BICIMP.LIB LIBECPG.LIB LIBPQ.LIB /OUT:testproc.exe
```

7.4.3 Bulk INSERT

This section describes the bulk INSERT.

Synopsis

```
EXEC SQL [ AT conn ] [ FOR { numOfRows | ARRAY_SIZE } ]
   INSERT INTO tableName [ ( colName [, ...] ) ]
   { VALUES ( { expr | DEFAULT } [, ...] ) [, ...] | query }
   [ RETURNING * | outputExpr [ [ AS ] outputName ] [, ...]
   INTO outputHostVar [ [ INDICATOR ] indicatorVar ] [, ...] ] END-EXEC
```

Description

Bulk INSERT is a feature that inserts multiple rows of data in bulk.

By specifying the array host variable that stored the data in the VALUES clause of the INSERT statement, the data for each element in the array can be inserted in bulk. This feature is used by specifying the insertion count in the FOR clause immediately before the INSERT statement.

FOR Clause

Specify the insertion count using numOfRows or ARRAY_SIZE in the FOR clause. The FOR clause can be specified only in the INSERT statement, not in other update statements.

numOfRows and ARRAY_SIZE

Insertion processing will be executed only for the specified count. However, if the count is 1, it will be assumed that the FOR clause was omitted when the application is executed. In this case, proceed according to the INSERT specification in the PostgreSQL Documentation.

Specify the FOR clause with a host variable defined as an integer item or with a literal.

Specify ARRAY_SIZE to insert all elements of the array in the table. When specifying ARRAY_SIZE, specify at least one array in *expr*.

If two or more arrays were specified in *expr*, it will be assumed that ARRAY_SIZE is the minimum number of elements in the array.

numOfRows or ARRAY_SIZE must exceed the minimum number of elements in all arrays specified in expr, outputHostVar, and indicatorVal.

The following example shows how to specify the FOR clause.

```
01 NUMBER-OF-ROWS PIC S9(9) COMP VALUE 10.
01 GROUP-ITEM.
05 ID1 PIC S9(9) OCCURS 25.
05 NAME PIC X(10) OCCURS 25.
* will process 10 rows

EXEC SQL FOR :NUMBER-OF-ROWS
INSERT INTO prod (name, id) VALUES (:NAME, :ID1) END-EXEC
* will process 25 rows

EXEC SQL FOR ARRAY_SIZE
INSERT INTO prod (name, id) VALUES (:NAME, :ID1) END-EXEC
```

expr

Specify the value to be inserted in the table. Array host variables, host variable literals, strings, and pointer variables can be specified. Structure type arrays and pointer variable arrays cannot be specified.

Do not use pointer variables and ARRAY_SIZE at the same time. The reason for this is that the number of elements in the area represented by the pointer variable cannot be determined.

query

A query (SELECT statement) that supplies the rows to be inserted. The number of rows returned by *query* must be 1. If two or more rows are returned, an error will occur. This cannot be used at the same time as ARRAY_SIZE.

outputHostVar and indicatorVal

These must be array host variables or pointer variables.

Error Messages

The messages below are output if an error occurs when the bulk INSERT is used.

Message

The value for the FOR clause must be a positive integer.

Cause

The value given for *numOfRows* is less than or equal to 0.

Solution

Specify a value that is more than or equal to 1 for numOfRows.

Message

Array host variable is needed when using FOR ARRAY_SIZE.

Cause

An array host variable is not specified in the VALUES clause.

Solution

Specify more than one array host variable in the VALUES clause.

Message

The SELECT..INTO query returned too many rows in row number %d.

Cause

The "SELECT ... INTO" query in the INSERT statement returned more than one row.

Solution

If *numOfRows* is more than one, the maximum number of rows that can be returned in the "SELECT ... INTO" query in the INSERT statement is one.

Limitations

The limitations when using bulk INSERT are given below.

- Array of structures should not be used as an input in the 'VALUES' clause.
- Array of pointers should not be used as an input in the 'VALUES' clause.
- ECOBPG supports the use of 'WITH' clause in single INSERT statements. 'WITH' clause cannot be used in bulk INSERT statements.
- If an error occurs, all bulk INSERT actions will be rolled back, therefore, no rows are inserted. However, if the RETURNING clause was used, and the error occurred while obtaining the rows after the insertion was successful, the insertion processing will not be rolled back.

Samples

Given below are some sample usages of the bulk INSERT functionality.

Basic Bulk INSERT

```
01 GROUP-ITEM.
05 IN-F1 PIC S9(9) OCCURS 4.

MOVE 1 TO IN-F1(1)

MOVE 2 TO IN-F1(2)

MOVE 3 TO IN-F1(3)

MOVE 4 TO IN-F1(4)

...

EXEC SQL FOR 3 INSERT INTO target (f1) VALUES (:IN-F1) END-EXEC
```

The number of rows to insert indicated by the FOR clause is 3, so the data in the first 3 elements of the host array variable are inserted into the table. The contents of the target table will be:

```
f1
----
1
2
3
(3 rows)
```

Also a host integer variable can be used to indicate the number of rows that will be inserted in FOR clause, which will produce the same result as above:

```
01 NUM PIC S9(9) COMP VALUE 3.
01 GROUP-ITEM.
05 IN-F1 PIC S9(9) OCCURS 4.
MOVE 1 TO IN-F1(1)
MOVE 2 TO IN-F1(2)
MOVE 3 TO IN-F1(3)
MOVE 4 TO IN-F1(4)
...
EXEC SQL FOR :NUM INSERT INTO target (f1) VALUES (:IN-F1) END-EXEC
```

Inserting constant values

Constant values can also be bulk INSERTed into the table as follows:

```
EXEC SQL FOR 3 INSERT INTO target (f1,f2) VALUES (DEFAULT, 'hello') END-EXEC
```

Assuming the 'DEFAULT' value for the 'f1' column is '0', the contents of the target table will be:

```
f1 | f2
---+------
0 | hello
0 | hello
0 | hello
(3 rows)
```

Using ARRAY_SIZE

'FOR ARRAY_SIZE' can be used to insert the entire contents of a host array variable, without explicitly specifying the size, into the table.

```
01 GROUP-ITEM.
05 IN-F1 PIC S9(9) OCCURS 4.

MOVE 1 TO IN-F1(1)

MOVE 2 TO IN-F1(2)

MOVE 3 TO IN-F1(3)

MOVE 4 TO IN-F1(4)

...

EXEC SQL FOR ARRAY_SIZE INSERT INTO target (f1) VALUES (:IN-F1) END-EXEC
```



If there are multiple host array variables specified as input values, then the number of rows inserted is same as the smallest array size. The example given below demonstrates this usage.

```
01 GROUP-ITEM.
05 IN-F1 PIC S9(9) OCCURS 4.
05 IN-F3 PIC X(10) OCCURS 3.
MOVE 1 TO IN-F1(1)
MOVE 2 TO IN-F1(2)
MOVE 3 TO IN-F1(3)
```

```
MOVE 4 TO IN-F1(4)

MOVE "one" TO IN-F3(1)

MOVE "two" TO IN-F3(2)

MOVE "three" TO IN-F3(3)

...

EXEC SQL FOR ARRAY_SIZE INSERT INTO target (f1,f3) VALUES (:IN-F1,:IN-F3) END-EXEC
```

In the above example, the array sizes are 3 and 4. Given that the smallest array size is 3, only three rows are inserted into the table. The table contents are given below.

```
f1 | f3
---+----
1 | one
2 | two
3 | three
(3 rows)
```

Using SELECT query

The result of a SELECT query can be used to insert values.

```
EXEC SQL FOR 4 INSERT INTO target(f1) SELECT age FROM source WHERE name LIKE 'foo' END-EXEC
```

In the example above, assuming that the SELECT query returns one row, the same row will be inserted into the table four times



If "2" or more is specified for the FOR clause, the INSERT statement returns an error when two or more rows of query results are returned.

If "1" is specified for the FOR clause, all rows returned by the SELECT query will be inserted into the table.

```
EXEC SQL FOR 1 INSERT INTO target(f1) SELECT age FROM source END-EXEC
```

In the example above, "1" specified for the FOR clause indicates all returned rows.

Using RETURNING clause

Bulk INSERT supports the same RETURNING clause syntax as normal INSERT. An example is given below.

```
01 GROUP-ITEM.
05 IN-F1 PIC S9(9) OCCURS 4.
05 OUT-F1 PIC S9(9) OCCURS 4.
MOVE 1 TO IN-F1(1)
MOVE 2 TO IN-F1(2)
MOVE 3 TO IN-F1(3)
MOVE 4 TO IN-F1(4)
...
EXEC SQL FOR 3 INSERT INTO target (f1) VALUES (:IN-F1) RETURNING f1 INTO :OUT-F1 END-EXEC
```

After the execution of the above INSERT statement, the 'out_f1' array will have 3 elements with the values of '1', '2' and '3'.

7.4.4 DECLARE STATEMENT

Refer to "DECLARE STATEMENT" in the PostgreSQL Documentation.

7.4.5 Creating Applications while in Database Multiplexing Mode

This section explains points to consider when creating applications while in database multiplexing mode.



- Refer to the Cluster Operation Guide (Database Multiplexing) for information on database multiplexing mode.
- Refer to "Application Development" in the Cluster Operation Guide (PRIMECLUSTER) for points to consider when creating applications using the failover feature integrated with the cluster software.

......

7.4.5.1 Errors when an Application Connection Switch Occurs and Corresponding Actions

If an application connection switch occurs while in database multiplexing mode, explicitly close the connection and then reestablish the connection or reexecute the application.

The table below shows errors that may occur during a switch, and the corresponding action to take.

State		Error information (*1)	Action
Server failure or Fujitsu Enterprise Postgres system failure	Failure occurs during access	57P01 57P02 YE000 26000 40001	After the switch is complete, reestablish the connection, or reexecute the application.
	Accessed during system failure	08001	
Switch to the standby server	Switched during access	57P01 57P02 YE000 26000 40001	
	Accessed during switch	08001	

^{*1:} Return value of SQLSTATE.

Chapter 8 SQL References

This chapter explains the SQL statement features expanded by Fujitsu Enterprise Postgres.

8.1 Expanded Trigger Definition Feature

This section explains the expanded trigger definition feature.

8.1.1 CREATE TRIGGER

In addition to features of PostgreSQL, triggers can be created with DO option.

Synopsis

Description

Refer to the PostgreSQL Documentation for information about CREATE TRIGGER. This section describes DO option.

A trigger which is created with DO option will be associated with the specified table or view and will execute the specified code by the specified procedural language of DO (unnamed code block) when certain events occur.

Parameters

lang_name

The name of the language that the function is implemented in.

plpgsql is supported in CREATE TRIGGER.

code

When the certain events occur, it executes the code in a specified procedural language. The unnamed code block does not require a prior definition like a function. Syntax is same as procedural language.



- A trigger defined with DO option cannot be replaced by a trigger defined with EXECUTE PROCEDURE option.
- A trigger defined with EXECUTE PROCEDURE option cannot be replaced by a trigger defined with DO option.

Examples

It executes the code block that is specified by DO before the table is updated. (Example that LANGUAGE is plpgsql)

```
CREATE TRIGGER check_update
BEFORE UPDATE ON accounts
```

FOR EACH ROW

DO \$\$BEGIN RETURN NEW; END;\$\$;



When a trigger created with DO option, a new function is created internally. The name of function is "schema name". "on table name"_"trigger name"_TRIGPROC(serial number).

8.1.2 How to Define Triggers in pgAdmin

The expanded features of the trigger definition can also be used in pgAdmin.



See

Refer to "pgAdmin Help" for information on how to define triggers using pgAdmin.

Chapter 9 Compatibility with Oracle Databases

This chapter describes the environment settings and functionality offered for features that are compatible with Oracle databases.

9.1 Overview

Features compatible with Oracle databases are provided. These features enable you to easily migrate to Fujitsu Enterprise Postgres and reduce the costs of reconfiguring applications.

The table below lists features compatible with Oracle databases.

Table 9.1 Features compatible with Oracle databases

Category		Feature		
		Item	Overview	
SQL	Queries	Outer join operator (+)	Operator for outer joins	
		DUAL table	Table provided by the system	
	Functions	DECODE	Compares values, and if they match, returns a corresponding value	
		SUBSTR	Extracts part of a string using characters to specify position and length	
		NVL	Returns a substitute value when a value is NULL	
Package		DBMS_ALERT	Sends alerts	
		DBMS_ASSERT	Perform assertions on input values	
		DBMS_OUTPUT	Sends messages to clients	
		DBMS_PIPE	Execution of inter-session communication	
		DBMS_RANDOM	Random number generation	
		DBMS_UTILITY	Addition of various functions	
		UTL_FILE	Enables text file operations	
		DBMS_SQL	Enables dynamic SQL execution	



In addition to the above, refer to the file below for information on the features compatible with Oracle databases.

- Linu
 - fujitsuEnterprisePostgresInstallDir/share/doc/extension/README.asciidoc
- Windows(R):

 $\textit{fujitsuEnterprisePostgresInstallDir} \\ \texttt{doc} \\ \texttt{extension} \\ \texttt{README}. \\ \texttt{asciidoc} \\ \\ \texttt{extension} \\ \texttt{readme}. \\ \texttt{extension} \\ \texttt{e$

9.2 Precautions when Using the Features Compatible with Oracle Databases

This section provides notes on using the features compatible with oracle databases.

9.2.1 Notes on search_path

Objects defined with the features compatible with oracle databases are defined in the oracle schema. Therefore, when using this function, it is necessary to add "oracle" to the "search_path" parameter in postgresql.conf.

```
search_path = '"$user", public, oracle'
```



- The search_path parameter specifies the order in which schemas are searched.
- Refer to "Statement Behavior" in "Client Connection Defaults" in "Server Administration" in the PostgreSQL Documentation for information on search_path.

9.2.2 Notes on SUBSTR

SUBSTR is implemented in Fujitsu Enterprise Postgres and Oracle databases using different external specifications.

For this reason, when using SUBSTR, define which specification is to take precedence. In the default configuration of Fujitsu Enterprise Postgres, the specifications of Fujitsu Enterprise Postgres take precedence.

When using the SUBSTR function compatible with Oracle databases, set "oracle" and "pg_catalog" in the "search_path" parameter of postgresql.conf. You must specify "oracle" before "pg_catalog" when doing this.

```
search_path = '"$user", public, oracle, pg_catalog'
```

9.2.3 Notes when Integrating with the Interface for Application Development

The SQL noted in "Table 9.1 Features compatible with Oracle databases" can be used in the interface for application development.

Note that both "public" and the schema name in the SQL statement must be specified as the SearchPath parameter before "oracle" and "pg_catalog" when using the Oracle database-compatible feature SUBSTR.

9.3 Queries

The following queries are supported:

- Outer Join Operator (+)
- DUAL Table

9.3.1 Outer Join Operator (+)

In the WHERE clause conditional expression, by adding the plus sign (+), which is the outer join operator, to the column of the table you want to add as a table join, it is possible to achieve an outer join that is the same as a joined table (OUTER JOIN).

Syntax

SELECT statement

```
SELECT ... [WHERE [NOT] joinCond ...] ...
SELECT ... [WHERE srchCond ]... ] ...

Join condition
{ colSpec(+) = colSpec | colSpec = colSpec(+) }
```



Here we are dealing only with the WHERE clause of the SELECT statement. Refer to "SQL Commands" in "Reference" in the PostgreSQL Documentation for information on the overall syntax of the SELECT statement.

General rules

WHERE clause

- The WHERE clause specifies search condition or join conditions for the tables that are derived.
- Search conditions are any expressions that return BOOLEAN types as the results of evaluation. Any rows that do not
 meet these conditions are excluded from the output. When the values of the actual rows are assigned to variables and
 if the expression returns TRUE, those rows are considered to have met the conditions.
- Join conditions are comparison conditions that specify outer join operators. Join conditions in a WHERE clause return a table that includes all the rows that meet the join conditions, including rows that do not meet all the join conditions.
- Join conditions take precedence over search conditions. For this reason, all rows returned by the join conditions are subject to the search conditions.
- The following rules and restrictions apply to queries that use outer join operators. It is therefore recommended to use FROM clause joined tables (OUTER JOIN) rather than outer join operators:
 - Outer join operators can only be specified in the WHERE clause.
 - Outer join operators can only be specified for base tables or views.
 - To perform outer joins using multiple join conditions, it is necessary to specify outer join operators for all join conditions.
 - When combining join conditions with constants, specify outer join operators in the corresponding column specification. When not specified, they will be treated as search conditions.
 - The results column of the outer join of table t1 is not returned if table t1 is joined with table t2 by specifying an outer join operator in the column of t1, then table t1 is joined with table t3 by using search conditions.
 - It is not possible to specify columns in the same table as the left/right column specification of a join condition.
 - It is not possible to specify an expression other than a column specification for outer join operators, but they may be specified for the columns that compose the expression.

There are the following limitations on the functionality of outer join operators when compared with joined tables (OUTER JOIN). To use functionality that is not available with outer join operators, use joined tables (OUTER JOIN).

Table 9.2 Range of functionality with outer join operators

Functionality available with joined tables (OUTER JOIN)	Outer join operator
Outer joins of two tables	Y
Outer joins of three or more tables	Y (*1)
Used together with joined tables within the same query	N
Use of the OR logical operator to a join condition	N
Use of an IN predicate to a join condition	N
Use of subqueries to a join condition	N

Y: Available N: Not available *1: The outer joins by outer join operators can return outer join results only for one other table. For this reason, to combine outer joins of table t1 and table t2 or table t2 and table t3, it is not possible to specify outer join operators simultaneously for table t2.



Table configuration

t1

col1	col2	col3
1001	AAAAA	1000
1002	BBBBB	2000
1003	CCCCC	3000

t2

col1	col2
1001	aaaaa
1002	bbbbb
1004	ddddd

Example 1: Return all rows in table t2, including those that do not exist in table t1.

This is the same syntax as the joined table (OUTER JOIN) of the FROM clause shown next.

```
SELECT *
FROM t1 RIGHT OUTER JOIN t2
ON t1.col1 = t2.col1;
```

Example 2: In the following example, the results are filtered to records above 2000 in t1.col3 by search conditions, and the records are those in table t2 that include ones that do not exist in table t1. After filtering with the join conditions, there is further filtering with the search conditions, so there will only be one record returned.

This is the same syntax as the joined table (OUTER JOIN) of the FROM clause shown next.

```
SELECT *
FROM t1 RIGHT OUTER JOIN t2
ON t1.col1 = t2.col1
WHERE t1.col3 >= 2000;
```

9.3.2 DUAL Table

DUAL table is a virtual table provided by the system. Use when executing SQL where access to a base table is not required, such as when performing tests to get result expressions such as functions and operators.

.....



In the following example, the current system date is returned.

```
SELECT CURRENT_DATE "date" FROM DUAL;
date
------
2013-05-14
(1 row)
```

9.4 SQL Function Reference

The following SQL functions are supported:

- DECODE
- SUBSTR
- NVL

9.4.1 DECODE

Description

Compares values and if they match, returns a corresponding value.

Syntax

```
DECODE(expr, srch, result [, srch, result ]... [, default ])
```

General rules

- DECODE compares values of the value expression to be converted and the search values one by one. If the values match, a corresponding result value is returned. If no values match, the default value is returned if it has been specified. A NULL value is returned if a default value has not been specified.
- If the same search value is specified more than once, then the result value returned is the one listed for the first occurrence of the search value.
- The following data types can be used in result values and in the default value:
 - CHAR
 - VARCHAR
 - NCHAR
 - NCHAR VARYING
 - TEXT
 - INTEGER
 - BIGINT

- NUMERIC
- DATE
- TIME WITHOUT TIME ZONE
- TIMESTAMP WITHOUT TIME ZONE
- TIMESTAMP WITH TIME ZONE
- The same data type must be specified for the values to be converted and the search values. However, note that different data types may also be specified if a literal is specified in the search value, and the value expression to be converted contains data types that can be converted. When specifying literals, refer to "Table A.1 Data type combinations that contain literals and can be converted implicitly" in "A.3 Implicit Data Type Conversions" for information on the data types that can be specified.
- If the result values and default value are all literals, the data types for these values will be as shown below:
 - If all values are string literals, all will become character types.
 - If there is one or more numeric literal, all will become numeric types.
 - If there is one or more literal cast to the datetime/time types, all will become datetime/time types.
- If the result values and default value contain a mixture of literals and non-literals, the literals will be converted to the data types of the non-literals. When specifying literals, refer to "Table A.1 Data type combinations that contain literals and can be converted implicitly" in "A.3 Implicit Data Type Conversions" for information on the data types that can be converted.
- The same data type must be specified for all result values and for the default value. However, different data types can be specified if the data type of any of the result values or default value can be converted these data types are listed below:

Table 9.3 Data type combinations that can be converted by DECODE (summary)

		Other result values or default value		
		Numeric type	Characte r type	Date/time type
Result value	Numeric type	Y	N	N
(any)	Character type	N	Y	N
	Date/time type	N	N	S (*1)

Y: Can be converted

Table 9.4 Result value and default value date/time data types that can be converted by DECODE

		Other result values or default value			
		DATE	TIME WITHOUT TIME ZONE	TIMESTAMP WITHOUT TIME ZONE	TIMESTAM P WITH TIME ZONE
Result	DATE	Y	N	Y	Y
value (any)	TIME WITHOUT TIME ZONE	N	Y	N	N
	TIMESTAMP WITHOUT TIME ZONE	Y	N	Y	Y
	TIMESTAMP WITH TIME ZONE	Y	N	Y	Y

S: Some data types can be converted

N: Cannot be converted

^{*1:} The data types that can be converted for date/time types are listed below:

- Y: Can be converted
- N: Cannot be converted
- The data type of the return value will be the data type within the result or default value that is longest and has the highest precision.



In the following example, the value of col3 in table t1 is compared and converted to a different value. If the col3 value matches search value 1, the result value returned is "one". If the col3 value does not match any of search values 1, 2, or 3, the default value "other number" is returned.

9.4.2 SUBSTR

Description

Extracts part of a string using characters to specify position and length.

Syntax

```
SUBSTR(str, startPos [, len ])
```

General rules

- SUBSTR extracts and returns a substring of string str, beginning at position startPos, for number of characters len.
- When startPos is positive, it will be the number of characters from the beginning of the string.
- When *startPos* is 0, it will be treated as 1.
- When startPos is negative, it will be the number of characters from the end of the string.
- When len is not specified, all characters to the end of the string are returned. NULL is returned when len is less than 1.
- For *startPos* and *len*, specify a SMALLINT or INTEGER type. When specifying literals, refer to "Table A.1 Data type combinations that contain literals and can be converted implicitly" in "A.3 Implicit Data Type Conversions" for information on the data types that can be specified.
- The data type of the return value is TEXT.



- There are two types of SUBSTR. One that behaves as described above, and one that behaves the same as SUBSTRING. The search_path parameter must be modified for it to behave the same as the specification described above.
- It is recommended to set search_path in postgresql.conf. In this case, it will be effective for each instance. Refer to "9.2.2 Notes on SUBSTR" for information on how to configure postgresql.conf.
- The configuration of search_path can be done at the user level or at the database level. Setting examples are shown below.
 - Example of setting at the user level

This can be set by executing an SQL command. In this example, user1 is used as the username.

```
ALTER USER user1 SET search_path = "$user",public,oracle,pg_catalog;
```

- Example of setting at the database level

This can be set by executing an SQL command. In this example, db1 will be used as the database name.

```
ALTER DATABASE db1 SET search_path = "$user",public,oracle,pg_catalog;
```

You must specify "oracle" before "pg_catalog".

- If the change has not been implemented, SUBSTR is the same as SUBSTRING.



See

Refer to "SQL Commands" in "Reference" in the PostgreSQL Documentation for information on ALTER USER and ALTER DATABASE.



Information

The general rules for SUBSTRING are as follows:

- The start position will be from the beginning of the string, whether positive, 0, or negative.
- When *len* is not specified, all characters to the end of the string are returned.
- An empty string is returned if no string is extracted or *len* is less than 1.



See

Refer to "String Functions and Operators" under "The SQL Language" in the PostgreSQL Documentation for information on SUBSTRING.



Example

In the following example, part of the string "ABCDEFG" is extracted:

```
SELECT SUBSTR('ABCDEFG',3,4) "Substring" FROM DUAL;

Substring
-----
CDEF
(1 row)

SELECT SUBSTR('ABCDEFG',-5,4) "Substring" FROM DUAL;

Substring
-----
(1 row)
```

9.4.3 NVL

Description

Returns a substitute value when a value is NULL.

Syntax

```
NVL(expr1, expr2)
```

General rules

- NVL returns a substitute value when the specified value is NULL. When *expr1* is NULL, *expr2* is returned. When *expr1* is not NULL, *expr1* is returned.
- Specify the same data types for *expr1* and *expr2*. However, if a constant is specified in *expr2*, and the data type can also be converted by *expr1*, different data types can be specified. When this happens, the conversion by *expr2* is done to suit the data type in *expr1*, so the value of *expr2* returned when *expr1* is a NULL value will be the value converted in the data type of *expr1*.
- When specifying literals, refer to "Table A.1 Data type combinations that contain literals and can be converted implicitly" in "A.3 Implicit Data Type Conversions" for information on the data types that can be converted.

......



In the following example, "IS NULL" is returned if the value of col1 in table t1 is a NULL value.

```
SELECT col2, NVL(col1,'IS NULL') "nvl" FROM t1;

col2 | nvl

----+----

aaa | IS NULL

(1 row)
```

9.5 Package Reference

A "package" is a group of features, brought together by schemas, that have a single functionality, and are used by calling from PL/pgSQL.

The following packages are supported:

- DBMS_ALERT
- DBMS_ASSERTION
- DBMS_OUTPUT
- DBMS_PIPE
- DBMS_RANDOM
- DBMS_UTILITUY
- UTL_FILE
- DBMS_SQL

To call each feature from PL/pgSQL, use the PERFORM or SELECT statement and qualify the feature name with the package name. For more information on the calling format, refer to the feature-specific description for each package.

In the following, explain DBMS_SQL among the supported packages. For other packages, refer to the README stored in the installation location.



See

For packages other than DBMS_SQL, refer to the following file.

- Linux:

fujitsuEnterprisePostgresInstallDir/share/doc/extension/README.asciidoc

9.5.1 DBMS_SQL

Overview

Dynamic SQL can be executed from PL/pgSQL.

Features

Feature	Description	
BIND_VARIABLE	Sets values in the host variable within the SQL statement.	
CLOSE_CURSOR	Closes the cursor.	
COLUMN_VALUE	Retrieves the value of the column in the select list extracted with FETCH_ROWS.	
DEFINE_COLUMN	Defines the column from which values are extracted and the storage destination.	
EXECUTE	Executes SQL statements.	
FETCH_ROWS	Positions the specified cursor at the next row and extracts values from the row.	
OPEN_CURSOR	Opens a new cursor.	
PARSE	Parses SQL statements.	



- In DBMS_SQL, the data types supported in dynamic SQL are limited, and therefore the user must consider this. The supported data types are:
 - INTEGER
 - DECIMAL
 - NUMERIC
 - REAL
 - DOUBLE PRECISION
 - CHAR(*1)
 - VARCHAR(*1)
 - NCHAR(*1)
 - NCHAR VARYING(*1)
 - TEXT
 - DATE
 - TIMESTAMP WITHOUT TIME ZONE
 - TIMESTAMP WITH TIME ZONE
 - INTERVAL(*2)
 - SMALLINT

- BIGINT

*1:

The host variables with CHAR, VARCHAR, NCHAR, and NCHAR VARYING data types are treated as TEXT, to match the string function arguments and return values. Refer to "String Functions and Operators" in "Functions and Operators" in "The SQL Language" in the PostgreSQL Documentation for information on string functions. When specifying the arguments of the features compatible with Oracle databases NVL and/or DECODE, use CAST to convert the data types of the host variables to ensure that data types between arguments are the same.

*2:

When using COLUMN_VALUE to obtain an INTERVAL type value specified in the select list, use an INTERVAL type variable with a wide range such as when no interval qualifier is specified, or with a range that matches that of the variable in the select list. If an interval qualifier variable with a narrow range is specified, then the value within the interval qualifier range will be obtained, but an error that the values outside the range have been truncated will not occur.



This example illustrates where a value expression that returns an INTERVAL value is set in the select list and the result is received with COLUMN_VALUE. Note that the SQL statement operation result returns a value within the INTERVAL DAY TO SECOND range.

[Bad example]

Values of MINUTE, and those after MINUTE, are truncated, because the variable(v_interval) is INTERVAL DAY TO HOUR.

```
v_interval INTERVAL DAY TO HOUR;
...
    PERFORM DBMS_SQL.PARSE(cursor, 'SELECT CURRENT_TIMESTAMP - ''2010-01-01'' FROM
DUAL', 1);
...
    SELECT value INTO v_interval FROM DBMS_SQL.COLUMN_VALUE(cursor, 1, v_interval);
    result:1324 days 09:00:00
```

[Good example]

By ensuring that the variable(v_interval) is INTERVAL, the values are received correctly.

```
v_interval INTERVAL;
...

PERFORM DBMS_SQL.PARSE(cursor, 'SELECT CURRENT_TIMESTAMP - ''2010-01-01'' FROM
DUAL', 1);
...

SELECT value INTO v_interval FROM DBMS_SQL.COLUMN_VALUE(cursor, 1, v_interval);
    result:1324 days 09:04:37.530623
```

Syntax

```
{ BIND_VARIABLE(cursor, varName, val [, len ]) | CLOSE_CURSOR(cursor) | COLUMN_VALUE(cursor, colPos, varName) | DEFINE_COLUMN(cursor, colPos, varName [, len ]) | EXECUTE(cursor) | FETCH_ROWS(cursor) | OPEN_CURSOR([parm1 ]) | PARSE(cursor, sqlStmt, parm1 [, parm2, parm3, parm4 ])
```

9.5.1.1 Description

This section explains each feature of DBMS_SQL.

BIND_VARIABLE

- BIND_VARIABLE sets values in the host variable within the SQL statement.
- Specify the cursor number to be processed.
- Specify the name of the host variable within the SQL statement using a string for the host variable name.
- Specify the value set in the host variable. The data type of the host variable is the same as that of the value expression it is implicitly converted in accordance with its position within the SQL statement. Refer to "A.3 Implicit Data Type Conversions" for information on implicit conversions.
- If the value is a character type, the string length is the number of characters. If the string length is not specified, the size is the total length of the string.
- It is necessary to place a colon at the beginning of the host variable in SQL statements to identify the host variable. The colon does not have to be added to the host variable names specified at BIND_VARIABLE. The following shows examples of host variable names specified with SQL statements and host variable names specified with BIND_VARIABLE:

```
PERFORM DBMS_SQL.PARSE(cursor, 'SELECT emp_name FROM emp WHERE sal > :x', 1);
```

In this example, BIND_VARIABLE will be as follows:

```
PERFORM DBMS_SQL.BIND_VARIABLE(cursor, ':x', 3500);
```

Or.

```
PERFORM DBMS_SQL.BIND_VARIABLE(cursor, 'x', 3500);
```

- The length of the host variable name can be up to 30 bytes (excluding colons).
- If the data type of the set value is string, specify the effective size of the column value as the fourth argument.

<section-header> Example

If the data type of the value to be set is not a string:

```
PERFORM DBMS_SQL.BIND_VARIABLE(cursor, ':NO', 1);
```

If the data type of the value to be set is a string:

```
PERFORM DBMS_SQL.BIND_VARIABLE(cursor, ':NAME', h_memid, 5);
```

CLOSE_CURSOR

- CLOSE_CURSOR closes the cursor.
- Specify the cursor number to be processed.
- The value returned is a NULL value.



```
cursor := DBMS_SQL.CLOSE_CURSOR(cursor);
```

COLUMN_VALUE

- COLUMN_VALUE retrieves the value of the column in the select list extracted with FETCH_ROWS.

- Specify the cursor number to be processed.
- Specify the position of the column of the select list in the SELECT statement. The position of the first column is 1.
- Specify the destination variable name.
- Use a SELECT statement to obtain the values of the value, column_error, and actual_length columns.
- The value column returns the value of the column specified at the column position. The data type of the variable name must match that of the column. If the data type of the column in the SELECT statement specified in PARSE is not compatible with DBMS_SQL, use CAST to convert to a compatible data type.
- The data type of the column_error column is NUMERIC. If the column value could not be set correctly in the value column, a value other than 0 will be returned:

22001: The extracted string has been truncated

22002: The extracted value contains a NULL value

- The data type of the actual_length column is INTEGER. If the extracted value is a character type, the number of characters will be returned (if the value was truncated, the number of characters prior to the truncation will be returned), otherwise, the number of bytes will be returned.

•••••



When retrieving the value of the column, the error code, and the actual length of the column value:

```
SELECT value, column_error, actual_length INTO v_memid, v_col_err, v_act_len FROM DBMS_SQL.COLUMN_VALUE(cursor, 1, v_memid);
```

When retrieving just the value of the column:

```
SELECT value INTO v_memid FROM DBMS_SQL.COLUMN_VALUE(cursor, 1, v_memid);
```

DEFINE_COLUMN

- DEFINE_COLUMN defines the column from which values are extracted and the storage destination.
- Specify the cursor number to be processed.
- Specify the position of the column in the select list in the SELECT statement. The position of the first column is 1.
- Specify the destination variable name. The data type should be match with the data type of the column from which the value is to be extracted. If the data type of the column in the SELECT statement specified in PARSE is not compatible with DBMS_SQL, use CAST to convert to a compatible data type.
- Specify the maximum number of characters of character type column values.
- If the data type of the column value is string, specify the effective size of the column value as the fourth argument.



When the data type of the column value is not a string:

```
PERFORM DBMS_SQL.DEFINE_COLUMN(cursor, 1, v_memid);
```

When the data type of the column value is a string:

```
PERFORM DBMS_SQL.DEFINE_COLUMN(cursor, 1, v_memid, 10);
```

EXECUTE

- EXECUTE executes SQL statements.

- Specify the cursor number to be processed.
- The return value is an INTEGER type, is valid only with INSERT statement, UPDATE statement, and DELETE statement, and is the number of rows processed. Anything else is invalid.



```
ret := DBMS_SQL.EXECUTE(cursor);
```

FETCH_ROWS

- FETCH_ROWS positions at the next row and extracts values from the row.
- Specify the cursor number to be processed.
- The return value is an INTEGER type and is the number of rows extracted. 0 is returned if all are extracted.
- The extracted information is retrieved with COLUMN_VALUE.



```
LOOP

IF DBMS_SQL.FETCH_ROWS(cursor) = 0 THEN

EXIT;

END IF;

...

END LOOP;
```

OPEN_CURSOR

- OPEN_CURSOR opens a new cursor.
- The parameter is used for compatibility with Oracle databases only, and is ignored by Fujitsu Enterprise Postgres. An INTEGER type can be specified, but it will be ignored. If migrating from an Oracle database, specify 1.
- Close unnecessary cursors by executing CLOSE_CURSOR.
- The return value is an INTEGER type and is the cursor number.



```
cursor := DBMS_SQL.OPEN_CURSOR();
```

PARSE

- PARSE analyzes dynamic SQL statements.
- Specify the cursor number to be processed.
- Specify the SQL statement to be parsed.
- Parameters 1, 2, 3, and 4 are used for compatibility with Oracle databases only, and are ignored by Fujitsu Enterprise Postgres. If you are specifying values anyway, specify the following:
 - Parameter 1 is an INTEGER type. Specify 1.

- Parameters 2 and 3 are TEXT types. Specify NULL.
- Parameter 4 is a BOOLEAN type. Specify TRUE.

If migrating from an Oracle database, the specified values for parameters 2, 3, and 4 do not need to be changed.

- Add a colon to the beginning of host variables in SQL statements.
- The DDL statement is executed when PARSE is issued. EXECUTE is not required for the DDL statement.
- If PARSE is called again for opened cursors, the content in the data regions within the cursors is reset, and the SQL statement is parsed anew.




PERFORM DBMS_SQL.PARSE(cursor, 'SELECT memid, memnm FROM member WHERE memid = :NO', 1);

9.5.1.2 Example

This section explains the flow of DBMS_SQL and provides an example.

Flow of DBMS_SQL

Flow of DBMS_SQL



Example

```
CREATE FUNCTION smp_00()
RETURNS INTEGER
AS $$
DECLARE

str_sql VARCHAR(255);
cursor INTEGER;
h_smpid INTEGER;
v_smpid INTEGER;
v_smpid INTEGER;
v_smpid VARCHAR(20);
v_smpage INTEGER;
errcd INTEGER;
```

```
length
             INTEGER;
   ret
              INTEGER;
BEGIN
   str_sql := 'SELECT smpid, smpnm, smpage FROM smp_tbl WHERE smpid < :H_SMPID ORDER BY
smpid';
   h_smpid
             := 3;
   v_smpid
              := 0;
              := '';
   v_smpnm
   v\_smpage := 0;
   cursor := DBMS_SQL.OPEN_CURSOR();
   PERFORM DBMS_SQL.PARSE(cursor, str_sql, 1);
   PERFORM DBMS_SQL.BIND_VARIABLE(cursor, ':H_SMPID', h_smpid);
   PERFORM DBMS_SQL.DEFINE_COLUMN(cursor, 1, v_smpid);
   PERFORM DBMS_SQL.DEFINE_COLUMN(cursor, 2, v_smpnm, 10);
   PERFORM DBMS_SQL.DEFINE_COLUMN(cursor, 3, v_smpage);
   ret := DBMS_SQL.EXECUTE(cursor);
   1000
       if DBMS_SQL.FETCH_ROWS(cursor) = 0 then
           EXIT;
       end if;
       SELECT value,column_error,actual_length INTO v_smpid,errcd,length FROM
DBMS_SQL.COLUMN_VALUE(cursor, 1, v_smpid);
       RAISE NOTICE '----';
       RAISE NOTICE '----';
       RAISE NOTICE 'smpid = %', v_smpid;
RAISE NOTICE 'errcd = %', errcd;
       RAISE NOTICE 'length
                               = %', length;
       SELECT value,column_error,actual_length INTO v_smpnm,errcd,length FROM
DBMS_SQL.COLUMN_VALUE(cursor, 2, v_smpnm);
       RAISE NOTICE '-----
       RAISE NOTICE 'smpnm = %', v_smpnm;
RAISE NOTICE 'errcd = %', errcd;
RAISE NOTICE 'length = %', length;
       select value,column_error,actual_length INTO v_smpage,errcd,length FROM
DBMS_SQL.COLUMN_VALUE(cursor, 3, v_smpage);
       RAISE NOTICE '-----';
       RAISE NOTICE 'smpage = %', v_smpage;
RAISE NOTICE 'errcd = %', errcd;
RAISE NOTICE 'length = %', length;
       RAISE NOTICE '';
   end loop;
   cursor := DBMS_SQL.CLOSE_CURSOR(cursor);
$$ LANGUAGE plpgsql;
```

Chapter 10 Application Connection Switch Feature

The application connection switch feature enables automatic connection to the target server when there are multiple servers with redundant configurations.

When using this feature, specify the primary server and secondary server as the connected servers in the application connection information. A standby server can optionally be prioritized over the primary server as the target server.

If an application connection switch occurs, explicitly close the connection and then reestablish the connection or reexecute the application. Refer to "Errors when an Application Connection Switch Occurs and Corresponding Actions" of the relevant client interface for information on how to confirm the switch.

10.1 Connection Information for the Application Connection Switch Feature

To use the application connection switch feature, set the information shown below when connecting the database.

IP address or host name

Specify the IP address or host name that will be used to configure the database multiplexing system.

Port number

A port number used by each database server to listen for connections from applications.

In each client interface, multiple port numbers can be specified, however in the format shown below, for example:

host1,host2:port2

JDBC and .NET

If only one port number is specified, it will be assumed that host1: 27500 (the default value) and host2:port2 were specified.

Omit all port numbers, or specify only one per server.

Others

If only one port number is specified, it will be assumed that the same port is used for all the hosts.

Target server

From the specified connection destination server information, specify the selection sequence of the servers to which the application will connect. The values specified for the target server have the meanings shown below. If a value is omitted, "any" will be assumed.

Primary server

The primary server is selected as the connection target from the specified "IP addresses or host names". Specify this to perform tasks that can be performed only on the primary server, such as applications in line with updates, or management tasks such as REINDEX and VACUUM.

Standby server

The standby server is selected as the connection target from the specified "IP addresses or host names". On standby server, the update will always fail. If the target server is not standby, the JDBC driver will throw an error stating that it is unable to find a server with the specified targetServerType.

Priority given to a primary server

The primary server is selected preferentially as the connection target from the specified "IP addresses or host names". If there is no primary server, the application will connect to the standby server.

Priority given to a standby server

The standby server is selected preferentially as the connection target from the specified "IP addresses or host names". If there is no standby server, the application will connect to the primary server.

Any

This method is not recommended in database multiplexing systems. This is because, although the connection destination server is selected in the specified sequence from the specified "IP addresses or host names", if the server that was successfully connected to first is the standby server, the write operations will always fail.

The table below shows the server selection order values to set for each driver:

Server selection order	JDBC drivers	.NET Data Provider	Other drivers
Drimory coryor	"primary"(*1)	"read-write"(*1)	"read-write"(*1)
Primary server	"primary"(*1)	"primary"(*2)	"primary"(*2)
Standby corror	"secondary"(*2)	"standby"	"standby"
Standby server	secondary (*2)	"read-only"(*2)	"read-only"(*2)
Priority given to a primary server	"preferPrimary"(*3	"prefer-primary"	-
Priority given to a standby server	"preferSecondary"(*2)	"prefer-standby"	"prefer- standby"(*4)
Any	"any"	"any"	"any"

^{*1:} The primary server whose default transaction mode is read-only is not selected.

SSL server certificate Common Name (CN)

To perform SSL authentication by creating the same server certificate for each server in a multiplexing system, specify the SSL server certificate Common Name (CN) in this parameter. Accordingly, SSL authentication using the CN can be performed without having to consider the names of the multiple servers contained in the multiplexing system.

10.2 Using the Application Connection Switch Feature

This section explains how to set the connection destination server using the application connection switch feature.

Of the parameters used as connection information for each client interface, only the parameters specific to the application connection switch feature are explained here. Refer to "Setup" and "Connecting to the Database" for information on the other parameters of each client interface.

10.2.1 Using the JDBC Driver

Set the following information in the connection string of the DriverManager class, or in the data source.

Table 10.1 Information to be set

Argument	Explanation	
host1 host2	Specify the IP address or host name. The IP address or host name can be omitted. If omitted, the default is localhost.	
port1 port2	Specify the port number for the connection. The port number can be omitted. If omitted, the default is 27500.	
database_name	Specify the database name.	
targetServerType	Specify the selection sequence of the servers to which the application will connect. Refer to "Target server" for details.	

^{*2:} The primary server whose default transaction mode is read-only is also selected.

^{*3:} Prefer primary servers whose default transaction mode is read-write.

^{*4:} While it is possible to specify "prefer-read" instead of "prefer-standby" for compatibility with Fujitsu Enterprise Postgres 13 and earlier, this is not recommended.

Argument	Explanation
sslmode	Specify this to encrypt communications. By default, this is disabled. The setting values for sslmode are as follows:
	disable: Connect without SSL
	require: Connect always with SSL
	verify-ca: Connect with SSL, using a certificate issued by a trusted CA (*1)
	verify-full: Connect with SSL, using a certificate issued by a trusted CA to verify if the server host name matches the certificate (*1)
sslservercertcn	This parameter is enabled only to perform SSL authentication (sslmode=verify-full).
	Specify the server certificate CN. If this is omitted, the value will be null, and the server certificate CN will be authenticated using the host name specified in host.

^{*1:} If specifying either "verify-ca" or "verify-full", the CA certificate file can be specified using connection string sslrootcert.

When using Driver Manager

Specify the following URL in the API of the DriverManager class:

```
jdbc:postgresql://[host1][:port1],[host2][:port2]/dbName[?targetServerType={primary |
secondary| preferPrimary | preferSecondary | any}][&sslmode=verify-
full&sslrootcert=cACertificateFile&sslservercertcn=targetServerCertificateCN]
```

- If the target server is omitted, the default value "any" is used.
- When using IPV6, specify the host in the "[host]" (with square brackets) format.

[Example]

```
jdbc:postgresql://[2001:Db8::1234]:27500,192.168.1.1:27500/dbName
```

When using the data source

Specify the properties of the data source in the following format:

```
source.setServerName("[host1][:port1],[host2][:port2]");
source.setTargetServerType("primary");
source.setSslmode("verify-full");
source.setSslrootcert("cACertificateFile");
source.setSslservercertcn("targetServerCertificateCN");
```

- If the IP address or host name are omitted, localhost will be used.
- If the port number is omitted, the value specified in the portNumber property will be used. Also, if the portNumber property is omitted, the default is 27500.
- If the target server is omitted, the value will be "any".
- When using IPV6, specify the host in the "[host]" (with square brackets) format.

[Example]

```
source.setServerName("[2001:Db8::1234]:27500,192.168.1.1:27500");
```



If using the connection parameter loginTimeout, the value will be applied for the time taken attempting to connect to all of the specified hosts.

10.2.2 Using the ODBC Driver

Set the following information in the connection string or data source.

Table 10.2 Information to be set

Parameter	Explanation	
Servername	Specify IP address 1 and IP address 2, or the host name, using a comma as the delimiter. Based on ODBC rules, it is recommended to enclose the whole string containing comma delimiters with {}.	
	Format: {host1,host2}	
Port	Specify the connection destination port numbers, using a comma as the delimiter. Based on ODBC rules, it is recommended to enclose the whole string containing comma delimiters with {}.	
	Format: {port1,port2}	
	Specify the port number corresponding to the IP address or host specified for the nth Servername as the nth Port.	
	The port number can be omitted. If omitted, the default is 27500.	
	If <i>n</i> server names are specified, and m ports are specified then there will be error reported. The only exceptions are where m=n or m=1. In case only one port is specified, then the same is applied for all the hosts.	
target_session_attrs	Specify the selection sequence of the servers to which the application will connect. Refer to "Target server" for details.	
SSLMode	Specify this to encrypt communications. By default, this is disabled. The setting values for SSLMode are as follows:	
	disable: Connect without SSL	
	allow: Connect without SSL, and if it fails, connect with SSL	
	prefer: Connect with SSL, and if it fails, connect without SSL	
	require: Connect always with SSL	
	verify-ca: Connect with SSL, using a certificate issued by a trusted CA (*1)	
	verify-full: Connect with SSL, using a certificate issued by a trusted CA to verify if the server host name matches the certificate (*1)	
SSLServerCertCN	This parameter is enabled only to perform SSL authentication (SSLMode=verifyfull).	
	Specify the server certificate CN. If this is omitted, the value will be null, and the server certificate CN will be authenticated using the host name specified in Servername.	

^{*1:} If specifying either "verify-ca" or "verify-full", use the system environment variable PGSSLROOTCERT of your operating system to specify the CA certificate file as shown below.

Example)

Variable name: PGSSLROOTCERT Variable value: *cACertificateFile*

When specifying a connection string

Specify the following connection string:

```
...;Servername={host1,host2};Port={port1,port2};[target_session_attrs={any | read-write | read-only | primary | standby | prefer-standby}];[ SSLMode=verify-full;SSLServerCertCN=targetServerCertificateCN]...
```

- When using IPV6, specify the host in the "host" format.

[Example]

```
Servername={2001:Db8::1234,192.168.1.1};Port={27500,27500};
```

When using the data source

Specify the properties of the data source in the following format:

```
Servername={host1,host2}

Port={port1,port2}

target_session_attrs={any | read-write | read-only | primary | standby | prefer-standby}

SSLMode=verify-full

SSLServerCertCN=targetServerCertificateCN
```


- When using IPV6, specify the host in the "host" format.

[Example]

```
Servername={2001:Db8::1234,192.168.1.1}
```

Registering the data source using the ODBC Data Source Administrator

Using the ODBC Data Source Administrator, specify the items within the red border below:





If using the connection parameter login_timeout, this value is applied for connections to each of the specified hosts. If both multiplexed database servers have failed, the connection will time out when a time equal to double the login_timeout value elapses.

10.2.3 Using a .NET Data Provider

Set the following information in the connection string of NpgsqlConnection, or in the data source.

Table 10.3 Information to be set

Argument	Explanation				
host1 host2	Specify the IP address or host name.				
port1 port2	Specify the port number for the connection.				
TargetSessionAttributes	Specify the selection sequence of the servers to which the application will connect. Refer to "Target server" for details.				

When specifying a connection string

Specify the following connection string:

```
host1[:port1],host2[:port2];[Target Session Attributes={any | primary | standby | prefer-primary | prefer-standby | read-write | read-only}];
```

- If the port number is omitted from the host string, the value specified for the Port keyword of the connection string will be used. Refer to "4.3.3 Connection String" for information on the Port keyword.
- When using IPV6, specify the host in the "[host]" (with square brackets) format.
- If the target server type is omitted, the value will be any.

[Example]

```
host=[2001:Db8::1234]:27500,192.168.1.1:27500;
```

When specifying the NpgsqlConnectionStringBuilder property

Specify the Host property of the data source in the following format:

```
host1[:port1],host2[:port2]
```

- If the port number is omitted from the host string, the value specified in the Port property will be used. Also, if the Port property is omitted, the default is 27500.

Specify the TargetSessionAttributes property of the data source in the following format:

any	primary	standby	prefer-primary	prefer-standby	read-write	read-only
				· · · · · · · · · · · · · · · · · · ·		•

- If the target server type is omitted, the value will be any.



If using the connection parameter Timeout, this value is applied for connections to each of the specified hosts. If both multiplexed database servers have failed, the connection will time out when a time equal to double the Timeout value elapses.

10.2.4 Using a Connection Service File

Set the connection parameters as follows.

Table 10.4 Information to be set

Parameter	Explanation		
host	Specify the host names, using a comma as the delimiter.		
hostaddr Specify IP address 1 and IP address 2, using a comma as the delin			

Parameter	Explanation		
port	Specify the connection destination port numbers, using a comma as the delimiter. Specify the port number for the server specified for the nth host or hostaddr as the nth port. The port number can be omitted. If omitted, the default is 27500.		
	If n server names are specified, and m ports are specified then there will be error reported. The only exceptions are where $m=n$ or $m=1$. In case only one port is specified, then the same is applied for all the hosts.		
target_session_attrs	Specify the selection sequence of the servers to which the application will connect. Refer to "Target server" for details.		
sslmode	Specify this to encrypt communications. By default, this is disabled. The setting values for sslmode are as follows:		
	disable: Connect without SSL		
	allow: Connect without SSL, and if it fails, connect with SSL		
	prefer: Connect with SSL, and if it fails, connect without SSL		
	require: Connect always with SSL		
	verify-ca: Connect with SSL, using a certificate issued by a trusted CA (*1)		
	verify-full: Connect with SSL, using a certificate issued by a trusted CA to verify if the server host name matches the certificate (*1)		
sslservercertcn	This parameter is enabled only to perform SSL authentication (sslmode=verify-full).		
	Specify the server certificate CN. If this is omitted, the value will be null, and the server certificate CN will be authenticated using the host name specified in host.		

*1: If specifying either "verify-ca" or "verify-full", use the system environment variable PGSSLROOTCERT (connection parameter sslrootcert) of your operating system to specify the CA certificate file as shown below.

Example)

Variable name: PGSSLROOTCERT Variable value: *cACertificateFile*



If using the connection parameter connect_timeout, this value is applied for connections to each of the specified hosts. If both multiplexed database servers have failed, the connection will time out when a time equal to double the connect_timeout value elapses.



If using the C Library, embedded SQL or psql commands (including other client commands that specify connection destinations), it is recommended to use a connection service file to specify connection destinations.

In the connection service file, a name (service name) is defined as a set, comprising information such as connection destination information and various types of tuning information set for connections. By using the service name defined in the connection service file when connecting to databases, it is no longer necessary to modify applications when the connection information changes.

10.2.5 Using the C Library (libpq)

It is recommended that you use a connection service file. Refer to "10.2.4 Using a Connection Service File" for details.

If a connection service file will not be used, set the following information for the database connection control functions (PQconnectdbParams, PQconnectdb, and so on) or environment variables.

Table 10.5 Information to be set

Parameter (environment variable name)	t Explanation		
host(PGHOST)	Specify the host names, using a comma as the delimiter.		
hostaddr(PGHOSTADDR)	Specify IP address 1 and IP address 2, using a comma as the delimiter.		
port(PGPORT)	Specify the connection destination port numbers, using a comma as the delimiter. Specify the port number for the server specified for the nth host or hostaddr as the nth port.		
	The port number can be omitted. If omitted, the default is 27500.		
	If n server names are specified, and m ports are specified then there will be error reported. The only exceptions are where m=n or m=1. In case only one port is specified, then the same is applied for all the hosts.		
target_session_attrs(PGTA RGETSESSIONATTRS)	Specify the selection sequence of the servers to which the application will connect. Refer to "Target server" for details.		
sslmode(PGSSLMODE)	Specify this to encrypt communications. By default, this is disabled. The setting values for sslmode are as follows:		
	disable: Connect without SSL		
	allow: Connect without SSL, and if it fails, connect with SSL		
	prefer: Connect with SSL, and if it fails, connect without SSL		
	require: Connect always with SSL		
	verify-ca: Connect with SSL, using a certificate issued by a trusted CA (*1)		
	verify-full: Connect with SSL, using a certificate issued by a trusted CA to verify if the server host name matches the certificate (*1)		
sslservercertcn(PGXSSLS ERVERCERTCN)	This parameter is enabled only to perform SSL authentication (sslmode=verifyfull).		
	Specify the server certificate CN. If this is omitted, the value will be null, and the server certificate CN will be authenticated using the host name specified in host.		

^{*1:} If specifying either "verify-ca" or "verify-full", use the system environment variable PGSSLROOTCERT (connection parameter sslrootcert) of your operating system to specify the CA certificate file as shown below.

Example)

Variable name: PGSSLROOTCERT Variable value: *cACertificateFile*

When using URI

```
postgresql://host1[:port1],host2[:port2][,...]/database_name
[?target_session_attrs={read-write | read-only | primary | standby | prefer-standby |
any }]
```

- When using IPV6, specify the host in the "[host]" (with square brackets) format.

[Example]

```
postgresq1://postgres@[2001:Db8::1234]:27500,192.168.1.1:27500/database_name
```

When using key-value

```
host=host1[,host2] port=port1[,port2] user=user1 password=pwd1 dbname=mydb
[target_session_attrs={read-write | read-only | primary | standby | prefer-standby |
any }]
```

- When using IPV6, specify the host in the "host" format.

[Example]

```
host=2001:Db8::1234,192.168.1.1 port=27500,27500
```



If using the connection parameter connect_timeout, this value is applied for connections to each of the specified hosts. If both multiplexed database servers have failed, the connection will time out when a time equal to double the connect_timeout value elapses.

Information

If using a password file (.pgpass), describe the entries matching each server.

- Example 1:

```
host1:port1:dbname:user:password
host2:port2:dbname:user:password
```

- Example 2:

*:port:dbname:user:password

10.2.6 Using Embedded SQL

It is recommended that you use a connection service file. Refer to "10.2.4 Using a Connection Service File" for details.



If using a connection service file, either of the following methods is available:

- Set the service name as a string literal or host variable, as follows: tcp:postgresql://?service=my_service
- Set the service name in the environment variable PGSERVICE, and use CONNECT TO DEFAULT

If a connection service file will not be used, use a literal or variable to specify the connection destination server information for target in the SQL statement below:

```
EXEC SQL CONNECT TO target [AS connection-name] [USER user-name];
```

Method used

```
dbname@host1,host2[:[port1][,port2]]
tcp:postgresql://host1,host2[:[port1][,port2]] [/dbname] [?target_session_attrs={read-write | read-only | primary | standby | prefer-standby | any}][&sslmode=verify-full&sslservercertcn=targetServerCertificateCN]
```

- The above format cannot be specified directly without using a literal or variable.

Table 10.6 Information to be set

Argument	Explanation			
host1 host2	Specify the IP address or host name. IPv6 format addresses cannot be specified.			
port1 port2	Specify the connection destination port numbers, using a comma as the delimiter. The port number can be omitted. If omitted, the default is 27500.			
dbname	Specify the database name.			
target_session_attrs	Specify the selection sequence of the servers to which the application will connect. Refer to "Target server" for details.			
sslmode	Specify this to encrypt communications. By default, this is disabled. The setting values for sslmode are as follows:			
	disable: Connect without SSL			
	allow: Connect without SSL, and if it fails, connect with SSL			
	prefer: Connect with SSL, and if it fails, connect without SSL			
	require: Connect always with SSL			
	verify-ca: Connect with SSL, using a certificate issued by a trusted CA (*1)			
	verify-full: Connect with SSL, using a certificate issued by a trusted CA to verify if the server host name matches the certificate (*1)			
sslservercertcn	This parameter is enabled only to perform SSL authentication (sslmode=verify-full).			
	Specify the server certificate CN. If this is omitted, the value will be null, and the server certificate CN will be authenticated using the host name specified in host.			

^{*1:} If specifying either "verify-ca" or "verify-full", use the system environment variable PGSSLROOTCERT (connection parameter sslrootcert) of your operating system to specify the CA certificate file as shown below.

Example)

Variable name: PGSSLROOTCERT Variable value: cACertificateFile



Environment variables can also be used. Refer to "10.2.5 Using the C Library (libpq)" for information on environment variables.



If using the connection parameter connect_timeout, this value is applied for connections to each of the specified hosts. If both multiplexed database servers have failed, the connection will time out when a time equal to double the connect_timeout value elapses.

10.2.7 Using the psql Command

It is recommended that you use a connection service file. Refer to "10.2.4 Using a Connection Service File" for details.

If a connection service file will not be used, specify the following information in the psql command option/environment variable.

Table 10.7 Information to be set

Option (environment variable)	Explanation		
-h/host(PGHOST/ PGHOSTADDR)	Specify IP address 1 and IP address 2, or the host name, using a comma as the delimiter. This can also be specified for the environment variable PGHOST or PGHOSTADDR.		
-p/port(PGPORT)	Specify the connection destination port numbers, using a comma as the delimiter. This can also be specified for the environment variable PGPORT. Specify the port number corresponding to the IP address specified for the nth -h option as the nth -p option.		
	The port number can be omitted. If omitted, the default is 27500.		
	If n -h options are specified, and m -p options are specified then there will be error reported. The only exception is where m=n or m=1. In case only one port is specified, then the same is applied for all the hosts.		
(PGTARGETSESSIONA TTRS)	Specify the selection sequence of the servers to which the application will connect. Refer to "Target server" for details.		
(PGSSLMODE)	Specify this to encrypt communications. By default, this is disabled. The setting values for PGSSLMODE are as follows:		
	disable: Connect without SSL		
	allow: Connect without SSL, and if it fails, connect with SSL		
	prefer: Connect with SSL, and if it fails, connect without SSL		
	require: Connect always with SSL		
	verify-ca: Connect with SSL, using a certificate issued by a trusted CA (*1)		
	verify-full: Connect with SSL, using a certificate issued by a trusted CA to verify if the server host name matches the certificate (*1)		
(PGXSSLSERVERCERT CN)	This environment variable is enabled only to perform SSL authentication (PGSSLMODE=verify-full).		
	Specify the server certificate CN. If this is omitted, the value will be null, and the server certificate CN will be authenticated using the host name specified in host.		

^{*1:} If specifying either "verify-ca" or "verify-full", use the system environment variable PGSSLROOTCERT (connection parameter sslrootcert) of your operating system to specify the CA certificate file as shown below.

Example)

Variable name: PGSSLROOTCERT Variable value: *cACertificateFile*



If using the connection parameter connect_timeout, this value is applied for connections to each of the specified hosts. If both multiplexed database servers have failed, the connection will time out when a time equal to double the connect_timeout value elapses.



Use the same method as for psql commands to specify connection destination server information for other client commands used to specify connection destinations.

••••••••••

Chapter 11 Scan Using a Vertical Clustered Index (VCI)

This chapter describes scanning using a VCI.



This feature can only be used in Advanced Edition.

11.1 Operating Conditions

Faster aggregation can be achieved by using a VCI defined for all columns to be referenced.

This section describes the conditions under which a scan can use a VCI.

Whether to use VCI is determined based on cost estimation in the same way as normal indexes. Therefore, another execution plan will be selected if it is cheaper than a VCI even if a VCI is available.

SQL statements that can use VCIs

In addition to general SELECT statements, VCIs can be used for the SQL statements below (as long as they do not specify any of the elements listed in "SQL statements that cannot use VCIs" below):

- SELECT INTO
- CREATE TABLE AS SELECT
- CREATE MATERIALIZED VIEW ... AS SELECT
- CREATE VIEW ... AS SELECT
- COPY (SELECT ...) TO

SQL statements that cannot use VCIs

VCIs cannot be used for SQL statements that specify any of the following:

- Subquery to reference the column in which the parent query is referencing is specified
- Lock clause (such as FOR UPDATE)
- Cursor declared with WITH HOLD or scrollable
- SERIALIZABLE transaction isolation level
- Function or operator listed in "Functions and operators that do not use a VCI"
- User-defined function


Table 11.1 Functions and operators that cannot use VCIs

Classification		Function/operator	
Mathematical functions and operators	Random functions	random and setseed	
String functions and operators	String functions	format (if the <i>format</i> argument is specified), regexp_matches, regexp_split_to_array and regexp_split_to_table	
Date/time functions and operators	Date/time functions	age(timestamp), current_date, current_time, current_timestamp, localtime, localtimestamp, statement_timestamp and transaction_timestamp	
	Delaying execution functions	pg_sleep, pg_sleep_for, and pg_sleep_until	

Classification		Function/operator	
Enum support functions		All functions and operators	
Geometric functions and	operators	All functions and operators	
Network address function	ns and operators	All functions and operators	
Text search functions and	l operators	All functions and operators	
XML functions		All functions	
JSON functions and oper	ators	All functions and operators	
Sequence manipulation for	unctions	All functions	
Array functions and oper	ators	All functions and operators	
Range functions and open	rators	All functions and operators	
Aggregate functions	General-purpose aggregate functions	array_agg, json_agg, json_object_agg, string_agg and xmlagg	
	Aggregate functions for statistics	corr, covar_pop, covar_samp, regr_avgx, regr_avgy, regr_count, regr_intercept, regr_r2, regr_slope, regr_sxx, regr_sxy and regr_syy	
	Ordered-set aggregate functions	All functions	
	Hypothetical-set aggregate functions	All functions	
Window functions	•	All functions	
Subquery expressions		Subquery expressions with its row constructor specified on the left side	
Row and array compariso	ons	Row constructor and composite type comparisons	
Set returning functions		All functions	
System information functions		All functions	
System administration functions		All functions	
Trigger functions		All functions	
Session information func	tions	current_role and current_user	

11.2 **Usage**

This section describes how to use a VCI in line with the following steps:



11.2.1 Designing

Design as follows before using a VCI.

- Execution multiplicity and number of parallel processes
- Parameters

Execution multiplicity and number of parallel processes

Determine the maximum number of SQL statements that can be executed simultaneously and the number of parallel processes based on the number of CPU cores that can be allocated for scans that use VCI to perform aggregate processing. Design in advance the multiplicity of SQL statements for executing scans that use VCI and the number of parallel processes for scans that use VCI.

For example, if the number of CPUs that can be allocated is 32 cores, then the maximum number of SQL statements that can be executed simultaneously is 8 and the number of parallel processes is 4.



A temporary file is created in /dev/shm or in a directory specified for the vci.smc_directory parameter as the dynamic shared memory for each SQL statement during a scan using a VCI.

A temporary file is created in a directory under the data storage directory or in a directory specified for the vci.smc_directory parameter as the dynamic shared memory for each SQL statement during a scan using a VCI.

Ensure that this directory has sufficient space to meet the memory requirements estimated for the execution multiplicity and number of parallel processes of SQL statements (refer to "Memory used per scanning" in "VCI Memory Requirements" in the Installation and Setup Guide for Server for details). If it does not have sufficient space when a scan is performed, SQL statements will return errors due to the insufficient memory.

Parameters

The VCI parallel scan feature cannot be used for setting parameters immediately after creating an instance.

Therefore, set the parameters below based on the values determined in "Execution multiplicity and number of parallel processes of SQL statements" above.

Parameter name	Description	Default	Value index
vci.max_parallel_degre e	Maximum number of VCI parallel	0	Specify the number of parallel processes.



Parameter name	Description	Default	Value index
	(background processes) to be used per SQL statement.		
vci.smc_directory	Directory name in which a temporary file is created as the dynamic shared memory during a scan using a VCI.	L /dev/shm W A directory (dataStorageDir\ base\pgsql_tmp) under the data storage directory	Specify a directory that has enough free space for the memory used for each query during the scan.
max_worker_processes	Maximum number of background processes that the system supports.	8	Add the value of the maximum number of SQL statements that can be executed simultaneously for scans that use VCI multiplied by vci.max_parallel_degree.



Refer to "Parameters" in the Operation Guide for information on the details of and how to set the parameters.

••••••

11.2.2 Checking

Execute the SQL statement with "EXPLAIN ANALYZE" to check the following:

- If a VCI was used "Custom Scan (VCI...)" is displayed in the plan if a VCI was used.
- Number of parallel processes
 The number of parallel processes when the SQL statement is executed is displayed in "Allocated Workers". Check that
 it is running the designed number of parallel processes.
- Response
 Check if the execution time displayed in "Execution time" is as estimated.

The following shows an example of the output result of EXPLAIN ANALYZE:



A cost output by the execution plan that uses a VCI may be inaccurate. A VCI works if all or part of the best execution plan when the SQL statement was executed is replaced with an execution plan that uses a VCI. If the cost of the execution plan to be replaced is lower than a certain value (vci.cost_threshold parameter), it will not be replaced or recalculated. Therefore, the cost of the original execution plan is output as is.

11.2.3 Evaluating

If the results in "11.2.2 Checking" is any of the following, tune accordingly:

If a VCI is not used

- Check if the "11.1 Operating Conditions" are met.
- Check if vci.enable is set to "on".
- A VCI may not be appropriately used when statistics are outdated, such as immediately after inserting a large amount of data. In such cases, execute the VACUUM ANALYZE statement or the ANALYZE statement.
- A VCI is not used if there is insufficient memory for VCI scan. This may occur during time-consuming transactions involving tables for which VCIs were defined. Set vci.log_query to "on", and check if either "could not use VCI: local ROS size (%zu) exceeds limit (%zu)" or "out of memory during local ROS generation" is output. If it is, then increase the value of the vci.max_local_ros.

Response is not as expected

Tuning may improve response. Check the following:

- If vci.max_parallel_degree is not set or is set to 0, set an appropriate value according to "11.2.1 Designing".
- If there is a margin in the CPU usage, increase the value of vci.max_parallel_degree and check again. In addition, if the value that of max_worker_processes is lower than the maximum number of SQL statements that can be executed simultaneously for parallel scan multiplied by vci.max_parallel_degree, increase it and check again.

11.3 Usage Notes

This section provides notes on using VCI.

- Regardless of whether VCI is used, the content of the result does not change. However, records may be returned in a different order if the ORDER BY clause is not specified.
- To reduce resource consumption, edit postgresql.conf or use the SET statement to enable/disable vci.enable when you use this feature only for specific times or jobs (SQL applications).
- The optimizer hint (pg hint plan) cannot be specified for a VCI. The hint clause is ignored if it is specified.
- If a plan other than VCI is specified for the optimizer hint (pg_hint_plan), a VCI may be used. Therefore, if you specify a query plan with the hint clause, use the SET statement to set vci.enable to "off".
- The message below may be output when a scan that uses VCI is performed on the streaming replication standby server:

```
"LOG: recovery has paused"
"HINT: Execute pg_wal_replay_resume() to continue."
```

This message is output because application of the WAL to the VCI temporarily pauses due to the scan being performed.

- Even if a scan is performed using a VCI, information in the idx_scan, idx_tup_read, and idx_tup_fetch columns of the collected statistics views, pg_stat_all_indexes and pg_stat_user_indexes, will not be updated.

- Currently, it is not possible to replace the query plan for parallel aggregation with the query plan using VCI. Therefore, if you create a VCI on a column of a partition table and aggregate (sum () etc.) on that column, one of the following plans will be selected. Use different setting parameters according to the situation of the target table.
 - Plan of the parallel aggregations using scan methods other than VCI scan

It is selected when max_parallel_workers_per_gather is 1 or more.

```
explain select sum(value) from test;

QUERY PLAN

Finalize Aggregate (cost=99906.30..99906.31 rows=1 width=8)

-> Gather (cost=99906.08..99906.29 rows=2 width=8)

Workers Planned: 2

-> Partial Aggregate (cost=98906.08..98906.09 rows=1 width=8)

-> Parallel Append (cost=0.00..94739.83 rows=1666500 width=4)

-> Parallel Seq Scan on test_1 (cost=0.00..43203.67 rows=833250 width=4)

-> Parallel Seq Scan on test_2 (cost=0.00..43203.67 rows=833250 width=4)
```

This plan is fast when the number of records to be aggregated (number of records that hit the search conditions) is very large. This is because the benefit of parallelizing aggregation is important, not the performance of scanning. For example, each parallel worker will perform a sequential scan and aggregate most of the scanned records.

- Plan that aggregates VCI scan results by a single aggregator node

It is selected by setting max_parallel_workers_per_gather to 0 and not creating a query plan of parallel aggregate.

This plan is fast when the number of aggregated items is not large or when the size of the aggregated column is smaller than the record size. This is because the scan performance is more important, so it is faster to aggregate the results of VCI scans of each partition.

- Originally, if there is only one partition to be accessed, the following VCI aggregation plan can be used. Below is an example of scanning only one partition with partition pruning.

```
explain select sum(value) from test where id < 1000001;

QUERY PLAN

Custom Scan (VCI Aggregate) (cost=62786.50..62786.51 rows=1 width=8)

Allocated Workers: 2

-> Custom Scan (VCI Scan) using test_l_id_value_idx on test_1 (cost=0.00..57787.00 rows=1999800 width=4)

Filter: (id < 1000001)
```

However, the current planner does not try to choose VCI aggregation because it creates a plan for parallel aggregation if the table is partitioned. So in this case, set max_parallel_workers_per_gather to 0 to force the planner to choose VCI aggregation.

Appendix A Precautions when Developing Applications

This appendix describes precautions when developing applications with Fujitsu Enterprise Postgres.

A.1 Precautions when Using Functions and Operators

This section describes notes for using functions and operators.

A.1.1 General rules of Functions and Operators

This section describes general rules for using functions and operators. Ensure the general rules are followed when using functions and operators to develop applications.

General rules

- Specify the stated numbers for arguments when specifying numbers for arguments in functions.
- Specify the stated data types when specifying data types for functions. If you use a data type other than the stated data types, use CAST to explicitly convert the data type.
- Specify data types that can be compared when specifying data types for operators. If you use a data type that cannot be compared, use CAST to explicitly convert the data type.



Refer to "Functions and Operators" under "The SQL Language" in the PostgreSQL Documentation for information on the functions and operators available with Fujitsu Enterprise Postgres.

A.1.2 Errors when Developing Applications that Use Functions and/ or Operators

This section provides examples of problems that may occur when developing applications that use functions and/or operators, and describes how to deal with them.

The error "Function ***** does not exist" occurs when executing SQL

The following error will occur when executing an SQL statement that does not abide by the general rules for functions:

```
ERROR: Function ***** does not exist
```

Note: "*****" denotes the function for which the error occurred, and the data type of its arguments.

The cause of the error will be one of the following:

- The specified function does not exist.
- The wrong number of arguments or wrong argument data type was specified

Corrective action

Check the following points and correct any errors:

- Check if there are any errors in the specified function name, number of arguments, or argument data type, and revise accordingly.
- Check the argument data type of the function displayed in the message. If an unintended data type is displayed, use a function such as CAST to convert it.

The error "Operator does not exist" occurs when executing SQL

The following error will occur when executing an SQL statement that specifies a data type in the operator that cannot be compared:

```
ERROR: Operator does not exist: *****
```

Note: "*****" denotes the operator for which the error occurred, and the data type of the specified value.

Corrective action

Ensure the data type of the expressions specified on the left and right sides of the operator can be compared. If required, revise to ensure these data types can be compared by using a function such as CAST to explicitly convert them.

A.2 Notes when Using Temporary Tables

In standard SQL, a temporary table can be defined in advance to enable an empty temporary table to be created automatically when the application connects to the database. However, in Fujitsu Enterprise Postgres, a temporary table must be created when the application connects to the database by explicitly using the CREATE TABLE statement.

If the same temporary table is repeatedly created and deleted during the same session, the system table might expand, and memory usage might increase. To prevent this, specify the CREATE TABLE statement to ensure the temporary table is reused.

For example, in cases where a temporary table would be created and deleted for repeatedly executed transactions, specify the CREATE TABLE statement as shown below:

- Specify "IF NOT EXISTS" to create a temporary table only if none exists when the transaction starts.
- Specify "ON COMMIT DELETE ROWS" to ensure all rows are deleted when the transaction ends.



Refer to "SQL Commands" under "Reference" in the PostgreSQL Documentation for information on the CREATE TABLE statement.

Examples of SQL using a temporary table are shown below:

Example of bad use (creating and deleting a temporary table)

```
BEGIN;
CREATE TEMPORARY TABLE mytable(coll CHAR(4), col2 INTEGER) ON COMMIT DROP;
(mytable processes)

COMMIT;
```

Example of good use (reusing a temporary table)

```
BEGIN;
CREATE TEMPORARY TABLE IF NOT EXISTS mytable(coll CHAR(4), col2 INTEGER) ON COMMIT
DELETE ROWS;
(mytable processes)

COMMIT;
```

A.3 Implicit Data Type Conversions

An implicit data type conversion refers to a data type conversion performed automatically by Fujitsu Enterprise Postgres, without the need to explicitly specify the data type to convert to.

The combination of possible data type conversions differs, depending on whether the expression in the conversion source is a literal.

For non-literals, data types can only be converted to other types within the same range.

For literals, character string literal types can be converted to the target data type. Numeric literals are implicitly converted to specific numeric types. These implicitly converted numeric literals can then have their types converted to match the conversion target data type within the numeric type range. For bit character string literals, only the bit column data type can be specified. The following shows the range of type conversions for literals.

Table A.1 Data type combinations that contain literals and can be converted implicitly

Conversion target		Conversion source			
		Characte r literal (*1)	Numeric literal(*2)	Bit character string literal	
Numeric type	SMALLINT	Y	N	N	
	INTEGER	Y	Y (*3)	N	
	BIGINT	Y	Y (*4)	N	
	DECIMAL	Y	Y (*5)	N	
	NUMERIC	Y	Y (*5)	N	
	REAL	Y	N	N	
	DOUBLE PRECISION	Y	N	N	
	SMALLSERIAL	Y	N	N	
	SERIAL	Y	Y (*3)	N	
	BIGSERIAL	Y	Y (*4)	N	
Currency type	MONEY	Y	N	N	
Character type	CHAR	Y	N	N	
	VARCHAR	Y	N	N	
	NCHAR	Y	N	N	
	NCHAR VARYING	Y	N	N	
	TEXT	Y	N	N	
Binary data type	BYTEA	Y	N	N	
Date/time type	TIMESTAMP WITHOUT TIME ZONE	Y	N	N	
	TIMESTAMP WITH TIME ZONE	Y	N	N	
	DATE	Y	N	N	
	TIME WITHOUT TIME ZONE	Y	N	N	
	TIME WITH TIME ZONE	Y	N	N	
	INTERVAL	Y	N	N	
Boolean type	BOOLEAN	Y	N	N	
Geometric type	POINT	Y	N	N	
	LSEG	Y	N	N	
	BOX	Y	N	N	

Conversion target		Co	Conversion source		
		Characte r literal (*1)	Numeric literal(*2)	Bit character string literal	
	PATH	Y	N	N	
	POLYGON	Y	N	N	
	CIRCLE	Y	N	N	
Network address type	CIDR	Y	N	N	
	INET	Y	N	N	
	MACADDR	Y	N	N	
	MACADDR8	Y	N	N	
Bit string type	BIT	Y	N	Y	
	BIT VARYING	Y	N	Y	
Text search type	TSVECTOR	Y	N	N	
	TSQUERY	Y	N	N	
UUID type	UUID	Y	N	N	
XML type	XML	Y	N	N	
JSON type	JSON	Y	N	N	

Y: Can be converted

Implicit data type conversions can be used when comparing or storing data.

The conversion rules differ, depending on the reason for converting. Purpose-specific explanations are provided below.

A.3.1 Function Argument

Value expressions specified in a function argument will be converted to the data type of that function argument.



Refer to "Functions and Operators" under "The SQL Language" in the PostgreSQL Documentation for information on data types that can be specified in function arguments.

.............

A.3.2 Operators

Comparison operators, BETWEEN, IN

Combinations of data types that can be compared using comparison operators, BETWEEN, or IN are shown below.

N: Cannot be converted

^{*1:} Only strings that can be converted to the data type of the conversion target can be specified (such as "1" if the conversion target is a numeric type)

^{*2: &}quot;Y" indicates specific numeric types that are converted first.

^{*3:} Integers that can be expressed as INTEGER types can be specified

^{*4:} Integers that cannot be expressed as INTEGER types, but can be expressed as BIGINT types, can be specified

^{*5:} Integers that cannot be expressed as INTEGER or BIGINT types, but that can be expressed as NUMERIC types, or numeric literals that contain a decimal point or the exponent symbol (e), can be specified

Table A.2 Combinations of comparable data type

Left side	Right side		
	Numeric type	Charact er string type	Date/time type
Numeric type	Y	N	N
Character type	N	Y	N
Date/time type	N	N	Y

Y: Can be compared

N: Cannot be compared

When strings with different lengths are compared, the shorter one is padded with spaces to make the lengths match.

When numeric values with different precisions are compared, data will be converted to the type with the higher precision.

Set operation and CASE also follow the same rules.

Other operators

Value expressions specified in operators will be converted to data types that are valid for that operator.



Refer to "Functions and Operators" under "The SQL Language" in the PostgreSQL Documentation for information on data types that can be specified in operators.

A.3.3 Storing Values

Value expressions specified in the VALUES clause of the INSERT statement or the SET clause of the UPDATE statement will be converted to the data type of the column in which they will be stored.

A.4 Notes on Using Index

This section explains the notes on using the following indexes:

- SP-GiST index

A.4.1 SP-GiST Index

If more than 2 concurrent updates are performed on a table in which the SP-GiST index is defined, applications may stop responding. When this occur, all system processes including the Check Pointer process will also be in the state of no response. For these reasons, use of the SP-GiST index is not recommended.

A.5 Notes on Using Multibyte Characters in Definition Names

Do not use multibyte characters in database names or user names if using a Windows database server.

Multibyte characters must not be used in database names or user names on non-Windows database servers, because certain conditions may apply or it may not be possible to connect to some clients.

Related notes and constraints are described below.

1) Configuring the client encoding system

The client encoding system must be configured when the names are created.



Refer to "Character Set Support" in "Server Administration" in the PostgreSQL Documentation for information on how to configure the client encoding system.

.....

2) Encoding system of names used for connection

Ensure that the encoding system of names used for connection is the same as that of the database that was connected when these names were created.

The reasons for this are as follows:

- Storage system for names in Fujitsu Enterprise Postgres

The system catalog saves encoded names by using the encoding system of the database at the time the names were created.

- Encoding conversion policy when connected

When connected, names sent from the client are matched with names in the system catalog without performing encoding conversion.

Accordingly, if the database that was connected when the names were defined uses the EUC_JP encoding system, but the database name is specified using UTF-8 encoding, then the database will be considered to be non-existent.

3) Connection constraints

The table below shows the connection constraints for each client type, based on the following assumptions:

- The conditions described in 1) and 2) above are satisfied.
- The database name and user names use the same encoding system.

Client type	Client operating system		
Client type	Windows(R)	Linux	
JDBC driver	Cannot be connected	Cannot be connected	
ODBC driver	Cannot be connected	No connection constraints	
.NET Data Provider	Can only connect when the encoding system used for definitions is UTF-8	-	
SQLEmbedded SQL in C	Can only connect when the connection service file (pg_service.conf) is used	No connection constraints	
psql command	Can only connect when the connection service file (pg_service.conf) is used	No connection constraints	

A.6 How to Build and Run an Application that Uses Shared Libraries

This section describes the following supplementary items regarding the use of shared libraries when developing applications with Fujitsu Enterprise Postgres.

- Setting DT_RUNPATH for the application
- Direct linking of indirectly used libraries to applications

A.6.1 Setting DT_RUNPATH for Applications

Searching for libraries used by applications

If your application uses Fujitsu Enterprise Postgres shared libraries such as libpq and ecpg, you need to load those libraries when you run your application.

When loading a library, it searches the machine for the library.

Library Search allows you to specify the path to search.

To specify the path, you can use the DT_RUNPATH attribute recorded in the application or library, or the environment variable LD_LIBRARY_PATH.

In general, use of the DT_RUNPATH attribute is recommended. This is because the environment variable LD_LIBRARY_PATH may affect the execution of applications other than the corresponding application.

The following explains when to use DT_RUNPATH.

Build so that the path in the operating environment that stores the library to be used is set in the DT_RUNPATH attribute of the application.

For example, when using libpq or ecpg, set "<*Fujitsu Enterprise Postgres client feature installation directory in operating environment*>/lib". (Note 1)

For how to set the DT_RUNPATH attribute, refer to the documentation of your compiler or linker.

There are the following three types of paths to be set and how to operate them. Please select the one suitable for each operation, and set and operate.

(1) Specify an absolute path


Set "<absolute path of the directory where the library is stored>/lib".

(2) Specify a relative path (Note 1)

Set "<relative path from the application of the directory that stores the library>/lib".

(3) Specify an arbitrary path and using a symbolic link

Set "(any path)". Also, create a symbolic link to the directory that stores the library at any path location.



Note 1:

For example, if you are building an application with gcc, add "-Wl,-rpath,< Fujitsu Enterprise Postgres installation directory in the operating environment/lib>,--enable-new-dtags" as gcc options, DT_RUNPATH can be set.



See

For setting DT_RUNPATH, refer to the linker ld's rpath and enable-new-dtags options, for example.

Also, refer to \$ORIGIN in ld.so for specifying the relative position from the application.

When DT_RUNPATH cannot be set

If none of the above settings work, you can find the shared libraries your application needs by setting the environment variable LD_LIBRARY_PATH to the path to the directory that stores the libraries when you run your application.

However, if you set LD_LIBRARY_PATH and execute commands or programs other than the applicable application, please be aware that they may result in run-time errors or unexpected behavior.



See

Check the ld.so man page for more information on searching for shared libraries.

A.6.2 Direct Linking of Indirectly Used Libraries to Applications



This content is complicated, and the possibility that it corresponds to this content in many applications is considered to be low.

Use this as a reference only when the application cannot be executed unexpectedly during a test, etc., when creating the application (as explained below, libF3 fails to load, or a runtime error occurs due to conflicts with other libraries).

apprentiation (as originated sector), not a time to rotat, or a running error sectors due to comment with state in time to rotate and the sector of the sect

For applications that use the libraries bundled with Fujitsu Enterprise Postgres, for example, if libF1, libF2, and libF3 have the following dependencies with other applications and libraries as shown below, libF3 can be directly Please specify in the build option to link.

If you do not link directly, specify the path where libF3 is stored in LD_LIBRARY_PATH when running the application.

Otherwise, an error may occur during application execution.

Example

The assumptions for this example are:

- There are the following dependencies between the application and the library.
 - apl1->libF1->libF2 (apl1 requires libF1 and libF1 requires libF2)
 - apl1->libA->libF3->libF2 (apl1 requires libA, libA requires libF3, libF3 requires libF2)
- libF1, libF2, and libF3 are libraries bundled with Fujitsu Enterprise Postgres.

When libF3 is not linked directly to apl1

/opt/apl/apl1

Required libraries: libF1

libA

DT_RPATH: /opt/fsepv<x>client64/lib

/opt/fsepv<x>client64/lib/libF1

Required libraries: libF2

DT RPATH: /opt/fsepv<x>client64/lib

/lib/libA

Required libraries: libF3

DT_RPATH: None *Search under /lib by default

/opt/fsepv<x>client64/lib/libF3


Required libraries: libF2

DT_RPATH: /opt/fsepv<x>client64/lib

/lib/libF3

Required libraries: libF2


DT_RPATH: None *Search under /lib by default



/lib/libF3 does not exist, or /lib/libF3 and /opt/fsepv<x>client64/lib/libF2 conflict (A symbol referenced by libF3 is not defined in libF2)

When libF3 is directly linked to apl1

/opt/apl/apl1 Required libraries: libF1 libF3 libA DT RPATH: /opt/fsepv<x>client64/lib /opt/fsepv<x>client64/lib/libF1 Required libraries: libF2 DT RPATH: /opt/fsepv<x>client64/lib Required libraries: libF3 DT RPATH: None *Search under /lib by default /opt/fsepv<x>client64/lib/libF3 Required libraries: libF2 DT RPATH: /opt/fsepv<x>client64/lib /lib/libF3 Required libraries: libF2 None *Search under /lib by default DT RPATH:



/opt/fsepv<x>client64/lib/libF3 and /opt/fsepv<x>client64/lib/libF2 does not conflict

Generally, when loading libraries, if some libA requires libB and libB requires libC, only the DT_RUNPATH value set in libA is used to search libB, and only the DT_RUNPATH value set in libB is used to search libC.

In the figure above "When libF3 is not linked directly to apl1", the search for libF1 uses the DT_RUNPATH value set in apl1. The search for libF2 uses the DT_RUNPATH value set for libF1. Since "Fujitsu Enterprise Postgres installation directory/lib" is set in each DT_RUNPATH, libF1 and libF2 bundled with Fujitsu Enterprise Postgres can be found and loaded when searching for libF1 and libF2.

However, searching for libF3 fails to load libF3 shipped with Fujitsu Enterprise Postgres. This is because the search for libF3 uses the DT_RUNPATH value set for libA, but libA does not have a DT_RUNPATH value.

Therefore, either libF3 cannot be found and loading of libF3 fails, or even if an unexpected libF3 is found in the machine and loaded, there is a conflict between libF2 bundled with Fujitsu Enterprise Postgres conflicts and can result in run-time errors. (The conflict here is that a symbol referenced in libF3 is not defined in libF2.)

By directly linking libF3 to the application, as shown in the figure above "When libF3 is directly linked to apl1", it is possible to use the DT_RUNPATH of the application and load the libF3 bundled with Fujitsu Enterprise Postgres. can. Alternatively, you can load libF3 shipped with Fujitsu Enterprise Postgres by setting LD_LIBRARY_PATH.

Appendix B Conversion Procedures Required due to Differences from Oracle Database

This appendix explains how to convert from an Oracle database to Fujitsu Enterprise Postgres, within the scope noted in "Chapter 9 Compatibility with Oracle Databases" from the following perspectives:

- Feature differences
- Specification differences

This document assumes that the version of the Oracle database to be converted is 7-10.2g.

B.1 Outer Join Operator (Perform Outer Join)

Features

In the WHERE clause conditional expression, by adding the plus sign (+), which is the outer join operator, to the column of the table you want to add as a table join, it is possible to achieve an outer join that is the same as a joined table (OUTER JOIN).

B.1.1 Comparing with the ^= Comparison Operator

Oracle database

```
SELECT *
FROM t1, t2
WHERE t1.col1(+) ^= t2.col1;
```

Note: col1 is assumed to be CHAR(4) type

Fujitsu Enterprise Postgres

```
SELECT *
FROM t1, t2
WHERE t1.col1(+) != t2.col1;
```

Note: col1 is assumed to be CHAR(4) type

Feature differences

Oracle database

The ^= comparison operator can be specified.

Fujitsu Enterprise Postgres

The ^= comparison operator cannot be specified.

Conversion procedure

Convert using the following procedure:

- 1. Locate the places where the keyword "^=" is used.
- 2. Ensure that the keyword, "(+)", is either on the right or left-hand side.
- 3. Change "^=" to "!=".

B.2 DECODE (Compare Values and Return Corresponding Results)

Features

DECODE compares values of the conversion target value expression and the search values one by one, and if the values of the conversion target value expression and the search values match, a corresponding result value is returned.

B.2.1 Comparing Numeric Data of Character String Types and Numeric Characters

Oracle database

Note: col1 is assumed to be CHAR(4) type

Fujitsu Enterprise Postgres

```
SELECT DECODE( CAST(col1 AS INTEGER),

1000, 'ITEM-A',

2000, 'ITEM-B',

'ITEM-C')

FROM t1;
```

Note: col1 is assumed to be CHAR(4) type

Feature differences

Oracle database

When the value expression is a string and the search value is a numeric, the string value will be converted to the data type of the comparison target numeric, so that they can be compared.

Fujitsu Enterprise Postgres

If the conversion target value expression is a string value, then no search value can be specified with numbers.

Conversion procedure

Since the data type that can be specified for the conversion target value expression is unknown, use CAST to explicitly convert the conversion target value expression (col1 in the example) to a numeric (INTEGER type in the example).

B.2.2 Obtaining Comparison Result from more than 50 Conditional Expressions

Oracle database

```
SELECT DECODE(col1,

1,'A',

2,'B',

...

78,'BZ',

NULL,'UNKNOWN',

'OTHER')

FROM t1;
```

Note: col1 is assumed to be INTEGER type

Fujitsu Enterprise Postgres

```
SELECT CASE

WHEN col1 = 1 THEN 'A'
WHEN col1 = 2 THEN 'B'

...
WHEN col1 = 78 THEN 'BZ'
WHEN col1 IS NULL THEN 'UNKNOWN'
ELSE 'OTHER'
END
FROM t1;
```

Note: col1 is assumed to be INTEGER type

Feature differences

Oracle database

Search value with a maximum of 127 items (up to 255 arguments in total) can be specified.

Fujitsu Enterprise Postgres

Search value with a maximum of 49 items (up to 100 arguments in total) only can be specified.

Conversion procedure

Convert to the CASE expression using the following procedure:

- 1. Specify the DECODE conversion target value expression (col1 in the first argument, in the example) and the search value (1 in the second argument, in the example) for the CASE expression search condition. Specify the DECODE result value ('A' in the third argument, in the example) for the CASE expression THEN (WHEN col1 = 1 THEN 'A', in the example). Note that if the search value is NULL, specify "IS NULL" for the search condition for the CASE expression.
- 2. If the DECODE default value ('OTHER' in the last argument, in the example) is specified, specify the default value for the CASE expression ELSE (ELSE 'OTHER', in the example).

B.2.3 Obtaining Comparison Result from Values with Different Data Types

Oracle database

Note: col1 is assumed to be CHAR(4) type

Fujitsu Enterprise Postgres

Note: col1 is assumed to be CHAR(4) type

Feature differences

Oracle database

The data types of all result values are converted to the data type of the first result value.

Fujitsu Enterprise Postgres

Results in an error.

Conversion procedure

Convert using the following procedure:

- 1. Check the literal data type for the first result value specified.
- 2. Change the literals specified for each result value to the literal data type checked in the step 1.

B.3 SUBSTR (Extract a String of the Specified Length from Another String)

Features

SUBSTR returns the number of characters specified in the third argument (starting from the position specified in the second argument) from the string specified in the first argument.

Refer to "9.2.2 Notes on SUBSTR" for details on precautions when using SUBSTR.

B.3.1 Specifying a Value Expression with a Data Type Different from the One that can be Specified for Function Arguments

Oracle database

Note: col1 and col2 are assumed to be CHAR type

Fujitsu Enterprise Postgres

Note: col1 and col2 are assumed to be CHAR type

Feature differences

Oracle database

If the type can be converted to a data type that can be specified for function arguments, conversion is performed implicitly.

Fujitsu Enterprise Postgres

If the data types are different from each other, or if loss of significance occurs, implicit conversion is not performed.

Conversion procedure

Since the data type of the string length is clear, first execute the following CREATE CAST only once so that the CHAR type value (col2 in the example) specified for the string length is implicitly converted to INTEGER type.

```
CREATE CAST (CHAR AS INTEGER) WITH INOUT AS IMPLICIT;
```

B.3.2 Extracting a String with the Specified Format from a Datetime Type Value

Oracle database

```
SELECT SUBSTR( CURRENT_TIMESTAMP,

1,
8)
FROM DUAL;
```

Fujitsu Enterprise Postgres

Feature differences

Oracle database

A datetime value such as CURRENT_TIMESTAMP can be specified for character value expressions.

Fujitsu Enterprise Postgres

A datetime value such as CURRENT_TIMESTAMP cannot be specified for character value expressions.

Conversion procedure

First, specify TO_CHAR for the SUBSTR character value expression.

Specify datetime type (CURRENT_TIMESTAMP, in the example) in firstArg of TO_CHAR, and specify the format template pattern ('DD-MON-YY HH.MI.SS.US PM', in the example) for secondArg to match with the result of SUBSTR before conversion.

TO_CHAR specification format: TO_CHAR(firstArg, secondArg)



Refer to "Data Type Formatting Functions" in the PostgreSQL Documentation for information on format template patterns that can be specified for TO_CHAR in Fujitsu Enterprise Postgres.

B.3.3 Concatenating a String Value with a NULL value

Oracle database

Note: col1 and col2 are assumed to be character string type, and col2 may contain NULL

Fujitsu Enterprise Postgres

```
SELECT SUBSTR( col1 || NVL(col2, '')
2,
5)
FROM t1;
```

Note: col1 and col2 are assumed to be character string type, and col2 may contain NULL

Feature differences

Oracle database

NULL is handled as an empty string, and strings are joined.

Fujitsu Enterprise Postgres

NULL is not handled as an empty string, and the result of joining the strings becomes NULL.

Conversion procedure

Convert using the following procedure:

- 1. Locate the places where the keyword "||" is used.
- 2. Check if any of the value expressions can contain NULL if they can, then execute step 3.
- 3. Modify to NVL(valExpr,").

B.4 NVL (Replace NULL)

Features

NVL converts NULL values.

B.4.1 Obtaining Result from Arguments with Different Data Types

Oracle database

```
SELECT NVL( col1, col2)
FROM t1;
```

Note: col1 is assumed to be VARCHAR(100) type, and col2 is assumed to be CHAR(100) type

Fujitsu Enterprise Postgres

```
SELECT NVL( col1,

CAST(col2 AS VARCHAR(100)))

FROM t1;
```

Note: col1 is assumed to be VARCHAR(100) type, and col2 is assumed to be CHAR(100) type

Feature differences

Oracle database

Value expressions with different data types can be specified. If the first argument is a string value, then VARCHAR2 is returned, and if it is a numeric, then a numeric type with greater range is returned.

Fujitsu Enterprise Postgres

Value expressions with different data types cannot be specified.

Conversion procedure

Since the data types that can be specified for the expressions in the two arguments are unknown, use the following steps to convert:

- 1. Check the data types specified for each of the two expressions.
- 2. Using the data type that is to be received as a result, explicitly convert the other argument with CAST.

B.4.2 Operating on Datetime/Numeric, Including Adding Number of Days to a Particular Day

Oracle database

```
SELECT NVL( col1 + 10, CURRENT_DATE)
FROM t1;
```

Note: col1 is assumed to be TIMESTAMP WITHOUT TIME ZONE type or TIMESTAMP WITH TIME ZONE type

Fujitsu Enterprise Postgres

```
SELECT NVL( CAST(coll AS DATE) + 10, CURRENT_DATE)
FROM t1;
```

Note: col1 is assumed to be TIMESTAMP WITHOUT TIME ZONE type or TIMESTAMP WITH TIME ZONE type

Feature differences

Oracle database

Numbers can be operated (added to or subtracted from) with either TIMESTAMP WITHOUT TIME ZONE type or TIMESTAMP WITH TIME ZONE type. Operation result will be DATE type.

Fujitsu Enterprise Postgres

Numbers cannot be operated (added to or subtracted from) with neither TIMESTAMP WITHOUT TIME ZONE type nor TIMESTAMP WITH TIME ZONE type. However, numbers can be operated (added to or subtracted from) with DATE type.

Conversion procedure

Convert using the following procedure:

- 1. Search locations where the keyword "+" or "-" is used in addition or subtraction, and check if these operations are between numbers and TIMESTAMP WITHOUT TIME ZONE type or TIMESTAMP WITH TIME ZONE type.
- 2. If they are, use CAST to explicitly convert TIMESTAMP WITHOUT TIME ZONE type or TIMESTAMP WITH TIME ZONE type to DATE type.

B.4.3 Calculating INTERVAL Values, Including Adding Periods to a Date

Oracle database

```
SELECT NVL( CURRENT_DATE + (col1 * 1.5), col2)
FROM tl;
```

Note: col1 and col2 are assumed to be INTERVAL YEAR TO MONTH types

Fujitsu Enterprise Postgres

```
SELECT NVL( CURRENT_DATE +
CAST(col1 * 1.5 AS
```

```
INTERVAL YEAR TO MONTH), col2)
FROM t1;
```

Note: col1 and col2 are assumed to be INTERVAL YEAR TO MONTH types

Feature differences

Oracle database

INTERVAL YEAR TO MONTH type multiplication and division result in INTERVAL YEAR TO MONTH type and any fraction (number of days) will be truncated.

Fujitsu Enterprise Postgres

INTERVAL YEAR TO MONTH type multiplication and division result in INTERVAL type and fractions (number of days) will not be truncated.

Conversion procedure

Convert using the following procedure:

- 1. Search locations where the keywords "*" or "/" are used in multiplication or division, and check if the specified value is INTERVAL YEAR TO MONTH type.
- 2. If the value is INTERVAL YEAR TO MONTH type, use CAST to explicitly convert the operation result to INTERVAL YEAR TO MONTH type.

B.5 DBMS_OUTPUT (Output Messages)

Features

DBMS_OUTPUT sends messages to clients such as psql from PL/pgSQL.

B.5.1 Outputting Messages Such As Process Progress Status

Oracle database

```
set serveroutput on; ... (1)
DECLARE
 v_col1
               CHAR(20);
               INTEGER;
 v_col2
 CURSOR c1 IS
   SELECT col1, col2 FROM t1;
BEGIN
 DBMS_OUTPUT.PUT_LINE('-- BATCH_001 Start --');
 OPEN c1;
 DBMS_OUTPUT.PUT_LINE('-- LOOP Start --');
   FETCH c1 INTO v_col1, v_col2;
   EXIT WHEN c1%NOTFOUND;
   DBMS OUTPUT.PUT('.');
 END LOOP;
 DBMS_OUTPUT.NEW_LINE; ...(2)
 DBMS_OUTPUT.PUT_LINE('-- LOOP End --');
 CLOSE c1;
 DBMS_OUTPUT.PUT_LINE('-- BATCH_001 End --');
EXCEPTION
 WHEN OTHERS THEN
    DBMS_OUTPUT.PUT_LINE('-- SQL Error --');
    DBMS_OUTPUT.PUT_LINE('ERROR : ' || SQLERRM );
```

```
END;
/
```

Fujitsu Enterprise Postgres

```
DO $$
DECLARE
   v_col1
               CHAR(20);
   v_col2
               INTEGER;
   c1 CURSOR FOR
       SELECT col1, col2 FROM t1;
BEGIN
   PERFORM DBMS_OUTPUT.SERVEROUTPUT(TRUE); ...(1)
   PERFORM DBMS_OUTPUT.ENABLE(NULL); ...(1)
   PERFORM DBMS_OUTPUT.PUT_LINE('-- BATCH_001 Start --');
   OPEN cl;
   PERFORM DBMS_OUTPUT.PUT_LINE('-- LOOP Start --');
   LOOP
       FETCH c1 INTO v_col1, v_col2;
       EXIT WHEN FOUND = false;
       PERFORM DBMS_OUTPUT.PUT('.');
   END LOOP;
   PERFORM DBMS_OUTPUT.NEW_LINE(); ...(2)
   PERFORM DBMS_OUTPUT.PUT_LINE('-- LOOP End --');
   CLOSE c1;
   PERFORM DBMS_OUTPUT.PUT_LINE('-- BATCH_001 End --');
EXCEPTION
   WHEN OTHERS THEN
       PERFORM DBMS_OUTPUT.PUT_LINE('-- SQL Error --');
       PERFORM DBMS_OUTPUT.PUT_LINE('ERROR : ' | SQLERRM );
END;
$$
```

(1) SERVEROUTPUT/ENABLE

Specification differences

Oracle database

Use SET statement and specify SERVEROUTPUT ON.

Fujitsu Enterprise Postgres

Specify DBMS_SQL.SERVEROUTPUT(TRUE).

Conversion procedure

Convert using the following procedure:

- 1. Check if a SET SERVEROUTPUT statement is specified before the PL/SQL block of a stored procedure.
- If a SET SERVEROUTPUT straight after BEGIN
 of PL/pgSQL. If ON is specified to have messages output to a window, then specify TRUE. If OFF is specified, then
 specify FALSE.
- 3. Specify DBMS_SQL.ENABLE only if SET SERVEROUTPUT is ON. The values to be specified for the argument are as follows:
 - If SIZE is specified for the SET SERVEROUTPUT statement, specify this size for the argument.

- If SIZE is not specified for the SET SERVEROUTPUT statement, then specify 2000 for Oracle10.1g or earlier, NULL for Oracle10.2g or later.

If DBMS_SQL.ENABLE is specified for the PL/SQL block of the stored procedure, specify the same value as that argument.

(2) NEW_LINE

Specification differences

Oracle database

If there is no argument for packageName.featureName, parenthesis can be omitted.

Fujitsu Enterprise Postgres

Even if there is no argument for packageName.featureName, parenthesis cannot be omitted.

Conversion procedure

Convert using the following procedure:

- 1. Locate the places where the keyword "DBMS_OUTPUT.NEW_LINE" is used in the stored procedure.
- 2. If there is no parenthesis after packageName.featureName, add the parenthesis.

B.5.2 Receiving a Return Value from a Procedure (PL/SQL) Block (For GET_LINES)

Oracle database

```
set serveroutput off;
DECLARE
                INTEGER;
   v_num
BEGIN
   DBMS_OUTPUT.DISABLE; ...(3)
   DBMS_OUTPUT.ENABLE(20000); ...(4)
   DBMS_OUTPUT.PUT_LINE('-- ITEM CHECK --');
   SELECT count(*) INTO v_num FROM t1;
   IF v_num = 0 THEN
       DBMS_OUTPUT.PUT_LINE('-- NO ITEM --');
       DBMS_OUTPUT.PUT_LINE('-- IN ITEM(' | v_num |  ') --');
    END IF;
END;
set serveroutput on;
DECLARE
   v_buffs
               DBMSOUTPUT_LINESARRAY; ...(5)
               INTEGER := 10;
   v_num
BEGIN
   DBMS_OUTPUT.GET_LINES(v_buffs, v_num); ...(5)
    FOR i IN 1..v_num LOOP
        DBMS_OUTPUT.PUT_LINE('LOG : ' || v_buffs(i)); ...(5)
```

```
END LOOP;
END;
```

Fujitsu Enterprise Postgres

```
DO $$
DECLARE
   v_num
                 INTEGER;
BEGIN
   PERFORM DBMS_OUTPUT.SERVEROUTPUT(FALSE);
   PERFORM DBMS_OUTPUT.DISABLE(); ...(3)
   PERFORM DBMS_OUTPUT.ENABLE(20000); ...(4)
PERFORM DBMS_OUTPUT.PUT_LINE('-- ITEM CHECK --');
   SELECT count(*) INTO v_num FROM t1;
   IF v_num = 0 THEN
       PERFORM DBMS_OUTPUT.PUT_LINE('-- NO ITEM --');
   ELSE
        PERFORM DBMS_OUTPUT.PUT_LINE('-- IN ITEM(' | v_num |  ') --');
   END IF;
END;
$$
DO $$
DECLARE
   v_buffs
                VARCHAR[]; ...(5)
                INTEGER := 10;
   v_num
BEGIN
   PERFORM DBMS_OUTPUT.SERVEROUTPUT(TRUE);
   SELECT lines, numlines INTO v_buffs, v_num FROM DBMS_OUTPUT.GET_LINES(v_num); ...(5)
   FOR i IN 1..v_num LOOP
       PERFORM DBMS_OUTPUT.PUT_LINE('LOG : ' | v_buffs[i]); ...(5)
   END LOOP;
END;
$$
```

(3) DISABLE

Same as the NEW_LINE in the DBMS_OUTPUT package. Refer to NEW_LINE for information on specification differences and conversion procedures associated with specification differences.

(4) ENABLE

Same as NEW_LINE in the DBMS_OUTPUT package. Refer to NEW_LINE for information on specification differences and conversion procedures associated with specification differences.

(5) GET_LINES

Specification format for Oracle database

DBMS_OUTPUT.GET_LINES(firstArg, secondArg)

Specification differences

Oracle database

Obtained values are received with variables specified for arguments.

Fujitsu Enterprise Postgres

Since obtained values are the search results for DBMS_OUTPUT.GET_LINES, they are received with variables specified for the INTO clause of the SELECT statement.

Conversion procedure

Convert using the following procedure:

- 1. Locate the places where the keyword "DBMS_OUTPUT.GET_LINES" is used in the stored procedure.
- 2. Change the data type (DBMSOUTPUT_LINESARRAY in the example) of the variable (v_buffs in the example) specified as *firstArg* of DBMS_OUTPUT.GET_LINES into a VARCHAR type array (VARCHAR[] in the example).
- 3. Replace the DBMS_OUTPUT.GET_LINES location called with a SELECT INTO statement.
 - Use the literal "lines, numlines" in the select list.
 - Specify *firstArg* (v_buffs in the example) and *secondArg* (v_num in the example) configured in DBMS_OUTPUT.GET_LINES, in the INTO clause.
 - Use DBMS_OUTPUT.GET_LINES in the FROM clause. Specify only *secondArg* (v_num in the example) before modification.
- 4. Identify the location that references *firstArg* (v_buffs in the example), and change it to the PL/pgSQL array reference format (v_buffs[i] in the example).

B.5.3 Receiving a Return Value from a Procedure (PL/SQL) Block (For GET_LINE)

Oracle database

Note: Only the process to obtain a value is stated

Fujitsu Enterprise Postgres

```
SELECT line, status INTO v_buff1, v_num FROM DBMS_OUTPUT.GET_LINE(); ...(6)
    EXIT WHEN v_num = 1;
    v_buff2 := v_buff2 || v_buff1;
    END LOOP;

PERFORM DBMS_OUTPUT.PUT_LINE(v_buff2);
END;
$$
;
```

Note: Only the process to obtain a value is stated

(6) GET_LINE

Specification format for Oracle database

DBMS_OUTPUT.GET_LINE(firstArg, secondArg)

Specification differences

Oracle database

Obtained values are received with variables specified for arguments.

Fujitsu Enterprise Postgres

Since obtained values are the search results for DBMS_OUTPUT.GET_LINES, they are received with variables specified for the INTO clause of the SELECT statement.

Conversion procedure

Convert using the following procedure:

- 1. Locate the places where the keyword "DBMS_OUTPUT.GET_LINE" is used in the stored procedure.
- 2. Replace the DBMS_OUTPUT.GET_LINE location called with a SELECT INTO statement.
 - Use the literal "line, status" in the select list.
 - Specify *firstArg* (v_buff1 in the example) and *secondArg* (v_num in the example) configured in DBMS_OUTPUT.GET_LINE, in the INTO clause.
 - Use DBMS_OUTPUT.GET_LINE in the FROM clause. Although arguments are not specified, parenthesis must be specified.

B.6 UTL_FILE (Perform File Operation)

Features

UTL_FILE reads and writes text files from PL/pgSQL.

B.6.1 Registering a Directory to Load and Write Text Files

Oracle database

```
[Oracle9i or earlier]
Configure the following with initialization parameter
UTL_FILE_DIR='/home/fsep' ...(1)

[Oracle9.2i or later]
Configure the following with CREATE DIRECTORY statement
CREATE DIRECTORY DIR AS '/home/fsep'; ...(1)
```

Fujitsu Enterprise Postgres

```
INSERT INTO UTL_FILE.UTL_FILE_DIR(dir)

VALUES('/home/fsep'); ...(1)
```

(1) UTL_FILE_DIR/CREATE DIRECTORY

Feature differences

Oracle database

Configure the directory to be operated, using the CREATE DIRECTORY statement or the initialization parameter UTL_FILE_DIR.

Fujitsu Enterprise Postgres

The directory to be operated cannot be configured using the CREATE DIRECTORY statement or the initialization parameter UTL_FILE_DIR.

Conversion procedure

Configure the target directory information in the UTL_FILE.UTL_FILE_DIR table using the INSERT statement. Note that this conversion procedure should be performed only once before executing the PL/pgSQL function.

- When using the initialization parameter UTL_FILE_DIR:
 - 1. Check the initialization parameter UTL_FILE_DIR value ('/home/fsep' in the example).
 - 2. Using the INSERT statement, specify and execute the directory name checked in step 1.
 - Specify UTL_FILE.UTL_FILE_DIR(dir) for the INTO clause.
 - Using the character string literal ('/home/fsep' in the example), specify the target directory name for the VALUES clause.
 - If multiple directories are specified, execute the INSERT statement for each directory.
- When using the CREATE DIRECTORY statement:
 - 1. Check the directory name ('/home/fsep' in the example) registered with the CREATE DIRECTORY statement. To check, log in SQL*Plus as a user with DBA privileges, and execute "show ALL_DIRECTORIES;".
 - 2. Using the INSERT statement, specify and execute the directory name checked in step 1. Same steps are used to specify the INSERT statement as when using the initialization parameter UTL_FILE_DIR.

B.6.2 Checking File Information

Oracle database

```
CREATE PROCEDURE read_file(fname VARCHAR2) AS

v_file    UTL_FILE.FILE_TYPE;
v_exists    BOOLEAN;
v_length    NUMBER;
v_bsize    INTEGER;
v_rbuff    VARCHAR2(1024);

BEGIN

UTL_FILE.FGETATTR('DIR', fname, v_exists, v_length, v_bsize); ...(2)

IF v_exists <> true THEN
    DBMS_OUTPUT.PUT_LINE('-- FILE NOT FOUND --');
    RETURN;
END IF;
```

```
DBMS_OUTPUT.PUT_LINE('-- FILE DATA --');
   v_file := UTL_FILE.FOPEN('DIR', fname, 'r', 1024); ...(3)
   FOR i IN 1..3 LOOP
       UTL_FILE.GET_LINE(v_file, v_rbuff, 1024); ...(4)
       DBMS_OUTPUT.PUT_LINE(v_rbuff);
   END LOOP;
   DBMS_OUTPUT.PUT_LINE('... more');
   DBMS_OUTPUT.PUT_LINE('-- READ END --');
   UTL_FILE.FCLOSE(v_file); ...(5)
   RETURN;
EXCEPTION
   WHEN NO_DATA_FOUND THEN
       DBMS_OUTPUT.PUT_LINE('-- FILE END --');
       UTL_FILE.FCLOSE(v_file);
       RETURN;
   WHEN OTHERS THEN
       DBMS_OUTPUT.PUT_LINE('-- SQL Error --');
       DBMS_OUTPUT.PUT_LINE('ERROR : ' | | SQLERRM );
       UTL_FILE.FCLOSE_ALL; ...(6)
       RETURN;
END;
set serveroutput on
call read_file('file01.txt');
```

Fujitsu Enterprise Postgres

```
CREATE FUNCTION read_file(fname VARCHAR) RETURNS void AS $$
DECLARE
   v_file
              UTL_FILE.FILE_TYPE;
   v_exists BOOLEAN;
   v_length NUMERIC;
   v_bsize INTEGER;
   v_rbuff
              VARCHAR(1024);
BEGIN
   PERFORM DBMS OUTPUT.SERVEROUTPUT(TRUE);
   SELECT fexists, file_length, blocksize
      INTO v_exists, v_length, v_bsize
      FROM UTL_FILE.FGETATTR('/home/fsep', fname); ...(2)
   IF v_exists <> true THEN
       PERFORM DBMS_OUTPUT.PUT_LINE('-- FILE NOT FOUND --');
       RETURN;
   END IF;
   PERFORM DBMS_OUTPUT.PUT_LINE('-- FILE DATA --');
   v_file := UTL_FILE.FOPEN('/home/fsep', fname, 'w', 1024); ...(3)
   FOR i IN 1..3 LOOP
       v_rbuff := UTL_FILE.GET_LINE(v_file, 1024); ...(4)
       PERFORM DBMS_OUTPUT.PUT_LINE(v_rbuff);
   END LOOP;
   PERFORM DBMS_OUTPUT.PUT_LINE('... more');
   PERFORM DBMS_OUTPUT.PUT_LINE('-- READ END --');
   v_file := UTL_FILE.FCLOSE(v_file); ...(5)
```

(2) FGETATTR

Specification format for Oracle database

UTL_FILE.FGETATTR(firstArg, secondArg, thirdArg, fourthArg, fifthArg)

Feature differences

Oracle database

If using a CREATE DIRECTORY statement (Oracle9.2i or later), specify a directory object name for the directory name.

Fujitsu Enterprise Postgres

A directory object name cannot be specified for the directory name.

Specification differences

Oracle database

Obtained values are received with variables specified for arguments.

Fujitsu Enterprise Postgres

Since obtained values are the search results for UTL_FILE.FGETATTR, they are received with variables specified for the INTO clause of the SELECT statement.

Conversion procedure

Convert using the following procedure. Refer to UTL_FILE_DIR/CREATE DIRECTORY for information on how to check if the directory object name corresponds to the actual directory name.

- $1. \ \ Locate \ the \ places \ where \ the \ keyword \ "UTL_FILE.FOPEN" \ is \ used \ in \ the \ stored \ procedure.$
- 2. Check the actual directory name ('/home/fsep' in the example) that corresponds to the directory object name ('DIR' in the example).
- 3. Replace the directory object name ('DIR' in the example) in *firstArg* with the actual directory name ('/home/fsep' in the example) verified in step 2.
- 4. Replace the UTL_FILE.FGETATTR location called with a SELECT INTO statement.
 - Use the literal "fexists, file_length, blocksize" in the select list.
 - Specify *thirdArg*, *fourthArg*, and *fifthArg* (v_exists, v_length, v_bsize, in the example) specified for UTL_FILE.FGETATTR to the INTO clause in the same order as that of the arguments.
 - Use UTL_FILE.FGETATTR in the FROM clause. Specify only the actual directory name for *firstArg* ('home/fsep' in the example) and *secondArg* (fname in the example) before modification for the arguments.

(3) FOPEN

Specification format for Oracle

UTL_FILE.FOPEN(firstArg, secondArg, thirdArg, fourthArg, fifthArg)

Feature differences

Oracle database

If using a CREATE DIRECTORY statement (Oracle9.2i or later), specify a directory object name for the directory name.

Fujitsu Enterprise Postgres

A directory object name cannot be specified for the directory name.

Conversion procedure

Convert using the following procedure. Refer to UTL_FILE_DIR/CREATE DIRECTORY for information on how to check if the directory object name corresponds to the actual directory name.

- 1. Locate the places where the keyword "UTL_FILE.FOPEN" is used in the stored procedure.
- Check the actual directory name ('/home/fsep' in the example) that corresponds to the directory object name ('DIR' in the example).
- 3. Replace the directory object name ('DIR' in the example) in *firstArg* with the actual directory name ('/home/fsep' in the example) checked in step 1.

(4) GET_LINE

Specification format for Oracle database

UTL_FILE.GET_LINE(firstArg, secondArg, thirdArg, fourthArg)

Specification differences

Oracle database

Obtained values are received with variables specified for arguments.

Fujitsu Enterprise Postgres

Since obtained values are the returned value of UTL_FILE.GET_LINE, they are received with variables specified for substitution statement.

Conversion procedure

Convert using the following procedure:

- 1. Locate the places where the keyword "UTL_FILE.GET_LINE" is used in the stored procedure.
- 2. Replace the UTL_FILE.GET_LINE location called with a value assignment (:=).
 - On the left-hand side, specify secondArg (v_rbuff in the example) specified for UTL_FILE.GET_LINE.
 - Use UTL_FILE.GET_LINE in the right-hand side. Specify only *firstArg* (v_file in the example) and *thirdArg* (1024 in the example) before modification.

(5) FCLOSE

Specification format for Oracle database

UTL_FILE.FCLOSE(firstArg)

Specification differences

Oracle database

After closing, the file handler specified for the argument becomes NULL.

Fujitsu Enterprise Postgres

After closing, set the file handler to NULL by assigning the return value of UTL_FILE.FCLOSE to it.

Conversion procedure

Convert using the following procedure:

- 1. Locate the places where the keyword "UTL_FILE.FCLOSE" is used in the stored procedure.
- 2. Replace the UTL_FILE.FCLOSE location called with a value assignment (:=) so that the file handler (v_file in the example) becomes NULL.
 - On the left-hand side, specify the argument (v_file in the example) specified for UTL_FILE.FCLOSE.
 - Use UTL_FILE.FCLOSE in the right-hand side. For the argument, specify the same value (v_file in the example) as before modification.

(6) FCLOSE_ALL

Same as NEW_LINE in the DBMS_OUTPUT package. Refer to NEW_LINE in the DBMS_OUTPUT for information on specification differences and conversion procedures associated with specification differences.

B.6.3 Copying Files

Oracle database

```
CREATE PROCEDURE copy_file(fromname VARCHAR2, toname VARCHAR2) AS
BEGIN

UTL_FILE.FCOPY('DIR1', fromname, 'DIR2', toname, 1, NULL); ...(7)

RETURN;

EXCEPTION
WHEN OTHERS THEN
DBMS_OUTPUT.PUT_LINE('-- SQL Error --');

DBMS_OUTPUT.PUT_LINE('ERROR : ' || SQLERRM );
RETURN;

END;
/
set serveroutput on
call copy_file('file01.txt','file01_bk.txt');
```

Fujitsu Enterprise Postgres

```
CREATE FUNCTION copy_file(fromname VARCHAR, toname VARCHAR) RETURNS void AS $$
BEGIN

PERFORM DBMS_OUTPUT.SERVEROUTPUT(TRUE);

PERFORM UTL_FILE.FCOPY('/home/fsep', fromname, '/home/backup', toname, 1, NULL); ...(7)

RETURN;

EXCEPTION

WHEN OTHERS THEN

PERFORM DBMS_OUTPUT.PUT_LINE('-- SQL Error --');

PERFORM DBMS_OUTPUT.PUT_LINE('ERROR : ' || SQLERRM );

RETURN;

END;
```

```
$$
LANGUAGE plpgsql;

SELECT copy_file('file01.txt','file01_bk.txt');
```

(7) FCOPY

Specification format for Oracle database

UTL_FILE.FCOPY(firstArg, secondArg, thirdArg, fourthArg, fifthArg, sixthArg)

Feature differences

Oracle database

If using a CREATE DIRECTORY statement (Oracle9.2i or later), specify a directory object name for the directory name.

Fujitsu Enterprise Postgres

A directory object name cannot be specified for the directory name.

Conversion procedure

Convert using the following procedure. Refer to UTL_FILE_DIR/CREATE DIRECTORY for information on how to check if the directory object name corresponds to the actual directory name.

- 1. Locate the places where the keyword "UTL_FILE.FCOPY" is used in the stored procedure.
- 2. Check the actual directory names ('/home/fsep' and '/home/backup', in the example) that correspond to the directory object names ('DIR1' and 'DIR2', in the example) of *firstArg* and *thirdArg* argument.
- 3. Replace the directory object name ('DIR1' and 'DIR2', in the example) with the actual directory names ('/home/fsep' in the example) checked in step 1.

B.6.4 Moving/Renaming Files

Oracle database

```
CREATE PROCEDURE move_file(fromname VARCHAR2, toname VARCHAR2) AS
BEGIN

UTL_FILE.FRENAME('DIR1', fromname, 'DIR2', toname, FALSE); ...(8)
RETURN;

EXCEPTION
WHEN OTHERS THEN
DBMS_OUTPUT.PUT_LINE('-- SQL Error --');
DBMS_OUTPUT.PUT_LINE('ERROR : ' || SQLERRM );
RETURN;

END;
/

set serveroutput on
call move_file('file01.txt','file02.txt');
```

Fujitsu Enterprise Postgres

```
CREATE FUNCTION move_file(fromname VARCHAR, toname VARCHAR) RETURNS void AS $$
BEGIN
PERFORM DBMS_OUTPUT.SERVEROUTPUT(TRUE);
```

```
PERFORM UTL_FILE.FRENAME('/home/fsep', fromname, '/home/backup', toname, FALSE); ...(8)
    RETURN;

EXCEPTION
    WHEN OTHERS THEN
        PERFORM DBMS_OUTPUT.PUT_LINE('-- SQL Error --');
        PERFORM DBMS_OUTPUT.PUT_LINE('ERROR : ' || SQLERRM );
        RETURN;

END;
$$
LANGUAGE plpgsql;

SELECT move_file('file01.txt','file02.txt');
```

(8) FRENAME

Same as FCOPY for the UTL_FILE package. Refer to FCOPY in the UTL_FILE package for information on specification differences and conversion procedures associated with specification differences.

B.7 DBMS_SQL (Execute Dynamic SQL)

Features

For DBMS_SQL, dynamic SQL can be executed from PL/pgSQL.

B.7.1 Searching Using a Cursor

Oracle database

```
CREATE PROCEDURE search_test(h_where CLOB) AS
   str_sql
             CLOB;
              INTEGER;
   v_cnt
              DBMS_SQL.VARCHAR2A;
   v_array
              INTEGER;
   v cur
   v_smpid
              INTEGER;
              VARCHAR2(20);
   v smpnm
   v_addbuff VARCHAR2(20);
   v_smpage INTEGER;
   errcd
              INTEGER;
              INTEGER;
   length
              INTEGER;
BEGIN
             := 'SELECT smpid, smpnm FROM smp_tbl WHERE ' || h_where || ' ORDER BY smpid';
   str_sql
             := 0;
   v_smpid
              := '';
   v_smpnm
   v_smpage
               := 0;
   v_cur := DBMS_SQL.OPEN_CURSOR; ...(1)
   v cnt :=
     CEIL(DBMS LOB.GETLENGTH(str sql)/1000);
   FOR idx IN 1 .. v_cnt LOOP
       v_array(idx) :=
           DBMS_LOB.SUBSTR(str_sql,
                           (idx-1)*1000+1);
```

```
END LOOP;
   {\tt DBMS\_SQL.PARSE}(v\_cur,\ v\_array,\ 1,\ v\_cnt,\ {\tt FALSE},\ {\tt DBMS\_SQL.NATIVE});\ \dots (2)
   DBMS_SQL.DEFINE_COLUMN(v_cur, 1, v_smpid);
   DBMS_SQL.DEFINE_COLUMN(v_cur, 2, v_smpnm, 10);
   ret := DBMS_SQL.EXECUTE(v_cur);
   LOOP
      v_addbuff := '';
       IF DBMS_SQL.FETCH_ROWS(v_cur) = 0 THEN
          EXIT;
       END IF;
       DBMS_OUTPUT.PUT_LINE('-----');
       DBMS_SQL.COLUMN_VALUE(v_cur, 1, v_smpid, errcd, length); ...(3)
       IF errcd = 1405 THEN ...(3)
        DBMS_OUTPUT.PUT_LINE('smpid = (NULL)');
       ELSE
        DBMS_OUTPUT.PUT_LINE('smpid = ' || v_smpid);
       END IF;
       DBMS_SQL.COLUMN_VALUE(v_cur, 2, v_smpnm, errcd, length);
       IF errcd = 1406 THEN
        v_addbuff := '... [len=' | length | | ']';
       IF errcd = 1405 THEN
        DBMS_OUTPUT.PUT_LINE('v_smpnm
                                       = (NULL)');
        DBMS_OUTPUT.PUT_LINE('v_smpnm = ' || v_smpnm || v_addbuff );
       END IF;
DBMS_OUTPUT.PUT_LINE('-----');
      DBMS_OUTPUT.NEW_LINE;
   END LOOP;
   DBMS_SQL.CLOSE_CURSOR(v_cur); ...(4)
   RETURN;
END;
Set serveroutput on
call search_test('smpid < 100');</pre>
```

Fujitsu Enterprise Postgres

```
CREATE FUNCTION search_test(h_where text) RETURNS void AS $$

DECLARE

str_sql text;

v_cur INTEGER;
```

```
v_smpid INTEGER;
v_smpnm VARCHAR(20);
   v_addbuff VARCHAR(20);
   v_smpage INTEGER;
            INTEGER;
   errcd
   length
             INTEGER;
   ret
             INTEGER;
BEGIN
   PERFORM DBMS_OUTPUT.SERVEROUTPUT(TRUE);
   str_sql := 'SELECT smpid, smpnm FROM smp_tbl WHERE ' | h_where | | ' ORDER BY smpid';
   v_smpid
   v_smpnm
              := '';
   v_smpage
              := 0;
   v_cur := DBMS_SQL.OPEN_CURSOR(); ...(1)
   PERFORM DBMS_SQL.PARSE(v_cur, str_sql, 1); ...(2)
   PERFORM DBMS_SQL.DEFINE_COLUMN(v_cur, 1, v_smpid);
   PERFORM DBMS_SQL.DEFINE_COLUMN(v_cur, 2, v_smpnm, 10);
   ret := DBMS_SQL.EXECUTE(v_cur);
   LOOP
       v_addbuff := '';
       IF DBMS_SQL.FETCH_ROWS(v_{cur}) = 0 THEN
          EXIT;
       END IF;
       PERFORM
DBMS_OUTPUT.PUT_LINE('-----');
       SELECT value,column_error,actual_length
        INTO v_smpid, errcd, length
        FROM DBMS_SQL.COLUMN_VALUE(v_cur,
                                 v_smpid); ...(3)
       IF errcd = 22002 THEN ...(3)
        PERFORM DBMS_OUTPUT.PUT_LINE('smpid
                                           = (NULL)');
        PERFORM DBMS_OUTPUT.PUT_LINE('smpid = ' | v_smpid);
       END IF;
       SELECT value,column_error,actual_length INTO v_smpnm, errcd, length FROM
DBMS_SQL.COLUMN_VALUE(v_cur, 2, v_smpnm);
       IF errcd = 22001 THEN
        v_addbuff := '... [len=' || length || ']';
       END IF;
       IF errcd = 22002 THEN
        PERFORM DBMS OUTPUT.PUT LINE('v smpnm = (NULL)');
        PERFORM DBMS_OUTPUT.PUT_LINE('v_smpnm = ' | v_smpnm | v_addbuff );
       END IF;
       PERFORM
DBMS_OUTPUT.PUT_LINE('-----');
      PERFORM DBMS_OUTPUT.NEW_LINE();
   END LOOP;
   v_cur := DBMS_SQL.CLOSE_CURSOR(v_cur); ...(4)
   RETURN;
END:
$$
LANGUAGE plpgsql;
```

(1) OPEN_CURSOR

Same as NEW_LINE in the DBMS_OUTPUT package. Refer to NEW_LINE in the DBMS_OUTPUT package for information on specification differences and conversion procedures associated with specification differences.

(2) PARSE

Specification format for Oracle database

DBMS_SQL.PARSE(firstArg, secondArg, thirdArg, fourthArg, fifthArg)

Feature differences

Oracle database

SQL statements can be specified with string table types (VARCHAR2A type, VARCHAR2S type). Specify this for *secondArg*.

DBMS_SQL.NATIVE, DBMS_SQL.V6, DBMS_SQL.V7 can be specified for processing SQL statements.

Fujitsu Enterprise Postgres

SQL statements cannot be specified with string table types.

DBMS_SQL.NATIVE, DBMS_SQL.V6, DBMS_SQL.V7 cannot be specified for processing SQL statements.

Conversion procedure

Convert using the following procedure:

- 1. Locate the places where the keyword "DBMS_SQL.PARSE" is used in the stored procedure.
- 2. Check the data type of the SQL statement specified for secondArg (v_array in the example).
 - If the data type is either DBMS_SQL.VARCHAR2A type or DBMS_SQL.VARCHAR2S type, then it is a table type specification. Execute step 3 and continue the conversion process.
 - If the data type is neither DBMS_SQL.VARCHAR2A type nor DBMS_SQL.VARCHAR2S type, then it is a string specification. Execute step 7 and continue the conversion process.
- Check the SQL statement (str_sql in the example) before it was divided into DBMS_SQL.VARCHAR2A type and DBMS_SQL.VARCHAR2S type.
- 4. Delete the sequence of the processes (processes near FOR idx in the example) where SQL is divided into DBMS_SQL.VARCHAR2A type and DBMS_SQL.VARCHAR2S type.
- 5. Replace *secondArg* with the SQL statement (str_sql in the example) before it is divided, that was checked in step
- 6. Delete thirdArg, fourthArg, and fifthArg (v_cnt, FALSE, DBMS_SQL.NATIVE, in the example).
- 7. If DBMS_SQL.NATIVE, DBMS_SQL.V6, and DBMS_SQL.V7 are specified, then replace *thirdArg* with a numeric literal 1.
 - If either DBMS_SQL.VARCHAR2A type or DBMS_SQL.VARCHAR2S type is used, then sixthArg becomes relevant.
 - If neither DBMS_SQL.VARCHAR2A type nor DBMS_SQL.VARCHAR2S type is used, then *thirdArg* becomes relevant.

(3) COLUMN_VALUE

Specification format for Oracle database

DBMS_SQL.COLUMN_VALUE(firstArg, secondArg, thirdArg, fourthArg, fifthArg)

Feature differences

Oracle database

The following error codes are returned for column_error.

- 1406: fetched column value was truncated
- 1405: fetched column value is NULL

Fujitsu Enterprise Postgres

The following error codes are returned for column_error.

- 22001: string_data_right_truncation
- 22002: null_value_no_indicator_parameter

Specification differences

Oracle database

Obtained values are received with variables specified for arguments.

Fujitsu Enterprise Postgres

Since obtained values are the search results for DBMS_SQL.COLUMN_VALUE, they are received with variables specified for the INTO clause of the SELECT statement.

Conversion procedure

Convert using the following procedure:

- 1. Locate the places where the keyword "DBMS_SQL.COLUMN_VALUE" is used in the stored procedure.
- 2. Replace the DBMS_SQL.COLUMN_VALUE location called with a SELECT INTO statement.
 - Check the number of arguments (v_smpid, errcd, and length in the example) specified after *secondArg* (1 in the example) of DBMS_SQL.COLUMN_VALUE.
 - Specify "value", "column_error", and "actual_length" in the select list, according to the number of arguments checked in the previous step (for example, if only *thirdArg* is specified, then specify "value" only.)
 - Specify *thirdArg*, *fourthArg*, and *fifthArg* (v_smpid, errcd, length in the example) configured for DBMS_SQL.COLUMN_VALUE, for the INTO clause.
 - Use DBMS_SQL.COLUMN_VALUE in the FROM clause. Specify *firstArg*, *secondArg*, and *thirdArg* (v_cur, 1, v_smpid, in the example) before modification.
- 3. If the *fourthArg* (column_error value in the example) is used, then check the location of the target variable (errcd in the example).
- 4. If a decision process is performed in the location checked, then modify the values used in the decision process as below:
 - 1406 to 22001
 - 1405 to 22002

(4) CLOSE_CURSOR

Specification format for Oracle database

 $DBMS_SQL.CLOSE_CURSOR(\textit{firstArg})$

Specification differences

Oracle database

After closing, the cursor specified in firstArg becomes NULL.

Fujitsu Enterprise Postgres

After closing, set the cursor to NULL by assigning the return value of DBMS_SQL.CLOSE_CURSOR to it.

Conversion procedure

Convert using the following procedure:

- $1. \ Locate the places where the keyword "DBMS_SQL.CLOSE_CURSOR" is used in the stored procedure.$
- 2. Set the cursor to NULL by assigning (:=) the return value of DBMS_SQL.CLOSE_CURSOR to it.
 - On the left-hand side, specify the argument (v_cur in the example) specified for DBMS_SQL.CLOSE_CURSOR.
 - Use DBMS_SQL.CLOSE_CURSOR in the right-hand side. For the argument, specify the same value (v_cur in the example) as before modification.

Appendix C Tables Used by the Features Compatible with Oracle Databases

This chapter describes the tables used by the features compatible with Oracle databases.

C.1 UTL_FILE.UTL_FILE_DIR

Register the directory handled by the UTL_FILE package in the UTL_FILE.UTL_FILE_DIR table.

Name	Туре	Description
dir	text	Name of the directory handled by the UTL_FILE package

Appendix D ECOBPG - Embedded SQL in COBOL

This appendix describes application development using embedded SQL in COBOL.

D.1 Precautions when Using Functions and Operators

An embedded SQL program consists of code written in an ordinary programming language, in this case COBOL, mixed with SQL commands in specially marked sections. To build the program, the source code (*.pco) is first passed through the embedded SQL preprocessor, which converts it to an ordinary COBOL program (*.cob), and afterwards it can be processed by a COBOL compiler. (For details about the compiling and linking see "D.9 Processing Embedded SQL Programs".) Converted ECOBPG applications call functions in the libpq library through the embedded SQL library (ecpglib), and communicate with the PostgreSQL server using the normal frontend-backend protocol.

Embedded SQL has advantages over other methods for handling SQL commands from COBOL code. First, it takes care of the tedious passing of information to and from variables in your C program. Second, the SQL code in the program is checked at build time for syntactical correctness. Third, embedded SQL in COBOL is specified in the SQL standard and supported by many other SQL database systems. The PostgreSQL implementation is designed to match this standard as much as possible, and it is usually possible to port embedded SQL programs written for other SQL databases to PostgreSQL with relative ease.

As already stated, programs written for the embedded SQL interface are normal COBOL programs with special code inserted to perform database-related actions. This special code always has the form:

EXEC SQL ... END-EXEC

These statements syntactically take the place of a COBOL statement. Depending on the particular statement, they can appear at the data division or at the procedure division. Actual executable SQLs need to be placed at the procedure division, and host variable declarations need to be placed at data division. However, the precompiler does not validate their placements. Embedded SQL statements follow the case-sensitivity rules of normal SQL code, and not those of COBOL.

For COBOL code notation, "fixed' or "variable" can be used. In each line, columns 1 to 6 constitute the line number area, and column 7 is the indicator area. Embedded SQL programs also should be placed in area B (column 12 and beyond).

Note that sample code in this document omits indents for each area.

ECOBPG processes or outputs programs according to the COBOL code notation. COBOL code notation is specified using the ecobpg command. Note, however, that the following restrictions apply:

- For "fixed" notation, area B is from columns 12 to 72. Characters in column 73 and beyond are deleted in the precompiled source.
- For "variable" notation, area B is from column 12 to the last column of that record (up to column 251). Characters in column 252 and beyond are deleted in the precompiled source.

ECOBPG accepts as many COBOL statements as possible. Note, however, that the following restrictions apply:

- In declaring host variable section, you can't use debug line.
- Outside of declaring host variable section, you can use debug line, but you can't contain any SQL in debug lines.
- In declaring host variable section, you can't use commas or semicolons as separator. Use space instead.
- EXEC SQL VAR command, it can be used in ECPG, is not available in ECOBPG. Use REDEFINE clause of COBOL instead.

The following sections explain all the embedded SQL statements.

D.2 Managing Database Connections

This section describes how to open, close, and switch database connections.

D.2.1 Connecting to the Database Server

One connects to a database using the following statement:

```
EXEC SQL CONNECT TO target [AS connection-name] [USER user-name] END-EXEC.
```

The target can be specified in the following ways:

- dbname[@hostname][:port]
- tcp:postgresql://hostname[:port][/dbname][?options]
- unix:postgresql://hostname[:port][/dbname][?options]
- an SQL string literal containing one of the above forms
- a reference to a character variable containing one of the above forms (see examples)
- DEFAULT

If you specify the connection target literally (that is, not through a variable reference) and you don't quote the value, then the case-insensitivity rules of normal SQL are applied. In that case you can also double-quote the individual parameters separately as needed. In practice, it is probably less error-prone to use a (single-quoted) string literal or a variable reference. The connection target DEFAULT initiates a connection to the default database under the default user name. No separate user name or connection name can be specified in that case.

There are also different ways to specify the user name:

- username
- username/password
- username IDENTIFIED BY password
- username USING password

As above, the parameters username and password can be an SQL identifier, an SQL string literal, or a reference to a character variable

The connection-name is used to handle multiple connections in one program. It can be omitted if a program uses only one connection. The most recently opened connection becomes the current connection, which is used by default when an SQL statement is to be executed (see later in this chapter).

Here are some examples of CONNECT statements:

```
EXEC SQL CONNECT TO mydb@sql.mydomain.com END-EXEC.
```

EXEC SQL CONNECT TO tcp:postgresql://sql.mydomain.com/mydb AS myconnection USER john END-EXEC.

```
EXEC SQL BEGIN DECLARE SECTION END-EXEC.

01 TARGET PIC X(25).

01 USER PIC X(5).

EXEC SQL END DECLARE SECTION END-EXEC.

...

MOVE "mydb@sql.mydomain.com" TO TARGET.

MOVE "john" TO USER.

EXEC SQL CONNECT TO :TARGET USER :USER END-EXEC.
```

The last form makes use of the variant referred to above as character variable reference. For this purpose, only fixed-length string(no VARYING) variable can be used. Trailing spaces are ignored. You will see in later sections how COBOL variables can be used in SQL statements when you prefix them with a colon.

Be advised that the format of the connection target is not specified in the SQL standard. So if you want to develop portable applications, you might want to use something based on the last example above to encapsulate the connection target string somewhere.

D.2.2 Choosing a Connection

SQL statements in embedded SQL programs are by default executed on the current connection, that is, the most recently opened one. If an application needs to manage multiple connections, then there are two ways to handle this.

The first option is to explicitly choose a connection for each SQL statement, for example:

```
EXEC SQL AT connection-name SELECT ... END-EXEC.
```

This option is particularly suitable if the application needs to use several connections in mixed order.

The second option is to execute a statement to switch the current connection. That statement is:

```
EXEC SQL SET CONNECTION connection-name END-EXEC.
```

This option is particularly convenient if many statements are to be executed on the same connection.

Here is an example program managing multiple database connections:

```
EXEC SQL BEGIN DECLARE SECTION END-EXEC.
   01 DBNAME PIC X(7).
EXEC SQL END DECLARE SECTION END-EXEC.
   EXEC SQL CONNECT TO testdb1 AS con1 USER testuser END-EXEC.
   EXEC SQL CONNECT TO testdb2 AS con2 USER testuser END-EXEC.
   EXEC SQL CONNECT TO testdb3 AS con3 USER testuser END-EXEC.
   This query would be executed in the last opened database "testdb3".
   EXEC SQL SELECT current_database() INTO :DBNAME END-EXEC.
   DISPLAY "current=" DBNAME " (should be testdb3)".
   Using "AT" to run a query in "testdb2"
   EXEC SQL AT con2 SELECT current_database() INTO :DBNAME END-EXEC.
   DISPLAY "current=" DBNAME " (should be testdb2)".
   Switch the current connection to "testdb1".
   EXEC SQL SET CONNECTION con1 END-EXEC.
   EXEC SQL SELECT current_database() INTO :DBNAME END-EXEC.
   DISPLAY "current=" DBNAME " (should be testdb1)".
    EXEC SQL DISCONNECT ALL END-EXEC.
```

This example would produce this output:

```
current=testdb3 (should be testdb3)
current=testdb2 (should be testdb2)
current=testdb1 (should be testdb1)
```

D.2.3 Closing a Connection

To close a connection, use the following statement:

```
EXEC SQL DISCONNECT [connection] END-EXEC.
```

The connection can be specified in the following ways:

- connection-name
- DEFAULT
- CURRENT
- ALL

If no connection name is specified, the current connection is closed.

It is good style that an application always explicitly disconnect from every connection it opened.

D.3 Running SQL Commands

Any SQL command can be run from within an embedded SQL application. Below are some examples of how to do that.

D.3.1 Executing SQL Statements

Creating a table:

```
EXEC SQL CREATE TABLE foo (number integer, ascii char(16)) END-EXEC.

EXEC SQL CREATE UNIQUE INDEX numl ON foo(number) END-EXEC.

EXEC SQL COMMIT END-EXEC.
```

Inserting rows:

```
EXEC SQL INSERT INTO foo (number, ascii) VALUES (9999, 'doodad') END-EXEC.
EXEC SQL COMMIT END-EXEC.
```

Deleting rows:

```
EXEC SQL DELETE FROM foo WHERE number = 9999 END-EXEC.

EXEC SQL COMMIT END-EXEC.
```

Updates:

```
EXEC SQL UPDATE foo

SET ascii = 'foobar'

WHERE number = 9999 END-EXEC.

EXEC SQL COMMIT END-EXEC.
```

SELECT statements that return a single result row can also be executed using EXEC SQL directly. To handle result sets with multiple rows, an application has to use a cursor; see "D.3.2 Using Cursors" below. (As a special case, an application can fetch multiple rows at once into an array host variable; see "Arrays".)

Single-row select:

```
EXEC SQL SELECT foo INTO :FooBar FROM table1 WHERE ascii = 'doodad' END-EXEC.
```

Also, a configuration parameter can be retrieved with the SHOW command:

```
EXEC SQL SHOW search_path INTO :var END-EXEC.
```

The tokens of the form :something are *host variables*, that is, they refer to variables in the COBOL program. They are explained in "D.4 Using Host Variables".

D.3.2 Using Cursors

To retrieve a result set holding multiple rows, an application has to declare a cursor and fetch each row from the cursor. The steps to use a cursor are the following: declare a cursor, open it, fetch a row from the cursor, repeat, and finally close it.

Select using cursors:

```
EXEC SQL DECLARE foo_bar CURSOR FOR

SELECT number, ascii FROM foo

ORDER BY ascii END-EXEC.

EXEC SQL OPEN foo_bar END-EXEC.

EXEC SQL FETCH foo_bar INTO :FooBar, :DooDad END-EXEC.

...

EXEC SQL CLOSE foo_bar END-EXEC.

EXEC SQL CLOSE foo_bar END-EXEC.
```

For more details about declaration of the cursor, see "D.11.4 DECLARE", and refer to "SQL Commands" in "Reference" in the PostgreSQL Documentation for information on FETCH command.

Note: The ECOBPG DECLARE command does not actually cause a statement to be sent to the PostgreSQL backend. The cursor is opened in the backend (using the backend's DECLARE command) at the point when the OPEN command is executed.

D.3.3 Managing Transactions

In the default mode, statements are committed only when EXEC SQL COMMIT is issued. The embedded SQL interface also supports autocommit of transactions (similar to libpq behavior) via the -t command-line option to ecobpg or via the EXEC SQL SET AUTOCOMMIT TO ON statement. In autocommit mode, each command is automatically committed unless it is inside an explicit transaction block. This mode can be explicitly turned off using EXEC SQL SET AUTOCOMMIT TO OFF.



Refer to "ecpg" in "PostgreSQL Client Applications" in the PostgreSQL Documentation for information on -t command-line option to ecobpg.

The following transaction management commands are available:

EXEC SQL COMMIT END-EXEC

Commit an in-progress transaction.

EXEC SQL ROLLBACK END-EXEC

Roll back an in-progress transaction.

EXEC SQL SET AUTOCOMMIT TO ON END-EXEC

Enable autocommit mode.

EXEC SQL SET AUTOCOMMIT TO OFF END-EXEC

Disable autocommit mode. This is the default.

D.3.4 Prepared Statements

When the values to be passed to an SQL statement are not known at compile time, or the same statement is going to be used many times, then prepared statements can be useful.

The statement is prepared using the command PREPARE. For the values that are not known yet, use the placeholder "?":

```
EXEC SQL PREPARE stmt1 FROM "SELECT oid, datname FROM pg_database WHERE oid = ?" END-EXEC.
```

If a statement returns a single row, the application can call EXECUTE after PREPARE to execute the statement, supplying the actual values for the placeholders with a USING clause:

```
EXEC SQL EXECUTE stmt1 INTO :dboid, :dbname USING 1 END-EXEC.
```

If a statement returns multiple rows, the application can use a cursor declared based on the prepared statement. To bind input parameters, the cursor must be opened with a USING clause:

```
EXEC SQL PREPARE stmt1 FROM "SELECT oid,datname FROM pg_database WHERE oid > ?" END-EXEC.

EXEC SQL DECLARE foo_bar CURSOR FOR stmt1 END-EXEC.

* when end of result set reached, break out of while loop

EXEC SQL WHENEVER NOT FOUND GOTO FETCH-END END-EXEC.

EXEC SQL OPEN foo_bar USING 100 END-EXEC.

...

PERFORM NO LIMIT

EXEC SQL FETCH NEXT FROM foo_bar INTO :dboid, :dbname END-EXEC

END-PERFORM.

FETCH-END.

EXEC SQL CLOSE foo_bar END-EXEC.
```

When you don't need the prepared statement anymore, you should deallocate it:

```
EXEC SQL DEALLOCATE PREPARE name END-EXEC.
```

For more details about PREPARE, see "D.11.10 PREPARE". Also see "D.5 Dynamic SQL" for more details about using placeholders and input parameters.

D.4 Using Host Variables

In "D.3 Running SQL Commands" you saw how you can execute SQL statements from an embedded SQL program. Some of those statements only used fixed values and did not provide a way to insert user-supplied values into statements or have the program process the values returned by the query. Those kinds of statements are not really useful in real applications. This section explains in detail how you can pass data between your COBOL program and the embedded SQL statements using a simple mechanism called host variables. In an embedded SQL program we consider the SQL statements to be guests in the COBOL program code which is the host language. Therefore the variables of the COBOL program are called host variables.

Another way to exchange values between PostgreSQL backends and ECOBPG applications is the use of SQL descriptors, described in "D.6 Using Descriptor Areas".

D.4.1 Overview

Passing data between the COBOL program and the SQL statements is particularly simple in embedded SQL. Instead of having the program paste the data into the statement, which entails various complications, such as properly quoting the value, you can simply write the name of a COBOL variable into the SQL statement, prefixed by a colon. For example:

```
EXEC SQL INSERT INTO sometable VALUES (:v1, 'foo', :v2) END-EXEC.
```

This statements refers to two COBOL variables named v1 and v2 and also uses a regular SQL string literal, to illustrate that you are not restricted to use one kind of data or the other.

This style of inserting COBOL variables in SQL statements works anywhere a value expression is expected in an SQL statement.

D.4.2 Declare Sections

To pass data from the program to the database, for example as parameters in a query, or to pass data from the database back to the program, the COBOL variables that are intended to contain this data need to be declared in specially marked sections, so the embedded SQL preprocessor is made aware of them.

This section starts with:

```
EXEC SQL BEGIN DECLARE SECTION END-EXEC.
```

and ends with:

```
EXEC SQL END DECLARE SECTION END-EXEC.
```

Between those lines, there must be normal COBOL variable declarations, such as:

```
01 INTX PIC S9(9) COMP VALUE 4.
01 FOO PIC X(15).
01 BAR PIC X(15).
```

As you can see, you can optionally assign an initial value to the variable. The variable's scope is determined by the location of its declaring section within the program.

You can have as many declare sections in a program as you like.

The declarations are also echoed to the output file as normal COBOL variables, so there's no need to declare them again. Variables that are not intended to be used in SQL commands can be declared normally outside these special sections.

The definition of a group item also must be listed inside a DECLARE section. Otherwise the preprocessor cannot handle these types since it does not know the definition.

D.4.3 Retrieving Query Results

Now you should be able to pass data generated by your program into an SQL command. But how do you retrieve the results of a query? For that purpose, embedded SQL provides special variants of the usual commands SELECT and FETCH. These commands have a special INTO clause that specifies which host variables the retrieved values are to be stored in. SELECT is used for a query that returns only single row, and FETCH is used for a query that returns multiple rows, using a cursor.

Here is an example:

```
* assume this table:
* CREATE TABLE test (a int, b varchar(50));

*

EXEC SQL BEGIN DECLARE SECTION END-EXEC.
01 V1 PIC S9(9).
01 V2 PIC X(50) VARYING.
EXEC SQL END DECLARE SECTION END-EXEC.
...

EXEC SQL SELECT a, b INTO :V1, :V2 FROM test END-EXEC.
```

So the INTO clause appears between the select list and the FROM clause. The number of elements in the select list and the list after INTO (also called the target list) must be equal.

Here is an example using the command FETCH:

```
EXEC SQL BEGIN DECLARE SECTION END-EXEC.

01 V1 PIC S9(9).

01 V2 PIC X(50) VARYING.

EXEC SQL END DECLARE SECTION END-EXEC.

...

EXEC SQL DECLARE foo CURSOR FOR SELECT a, b FROM test END-EXEC.

...

PERFORM WITH

...

EXEC SQL FETCH NEXT FROM foo INTO :V1, :V2 END-EXEC
```

END-PERFORM.

Here the INTO clause appears after all the normal clauses.

D.4.4 Type Mapping

When ECOBPG applications exchange values between the PostgreSQL server and the COBOL application, such as when retrieving query results from the server or executing SQL statements with input parameters, the values need to be converted between PostgreSQL data types and host language variable types (specifically COBOL language data types). One of the main points of ECOBPG is that it takes care of this automatically in most cases.

In this respect, there are two kinds of data types: Some simple PostgreSQL data types, such as integer and text, can be read and written by the application directly. Other PostgreSQL data types, such as timestamp and date can only be accessed through character strings. special library functions does not exist in ecobpg. (pgtypes, exists in ECPG, for COBOL is not implemented yet)

"Table D.1 Mapping Between PostgreSQL Data Types and COBOL Variable Types" shows which PostgreSQL data types correspond to which COBOL data types. When you wish to send or receive a value of a given PostgreSQL data type, you should declare a COBOL variable of the corresponding COBOL data type in the declare section.

Table D.1 Mapping Between PostgreSQL Data Types and COBOL Variable Types

PostgreSQL data type	COBOL Host variable type
smallint	PIC S9([1-4]) {BINARY COMP COMP-5}
integer	PIC S9([5-9]) {BINARY COMP COMP-5}
bigint	PIC S9([10-18]) {BINARY COMP COMP-5}
decimal	PIC S9(m)V9(n) PACKED-DECIMAL PIC 9(m)V9(n) DISPLAY (*1) PIC S9(m)V9(n) DISPLAY PIC S9(m)V9(n) DISPLAY SIGN TRAILING [SEPARATE] PIC S9(m)V9(n) DISPLAY SIGN LEADING [SEPARATE]
numeric	(same with decimal)
real	COMP-1
double precision	COMP-2
small serial	PIC S9([1-4]) {BINARY COMP COMP-5}
serial	PIC S9([1-9]) {BINARY COMP COMP-5}
bigserial	PIC S9([10-18]) {BINARY COMP COMP-5}
oid	PIC 9(9) {BINARY COMP COMP-5}
character(n), varchar(n), text	PIC X(n), PIC X(n) VARYING
name	PIC X(NAMEDATALEN)
boolean	BOOL(*2)
bytea	BYTEA(n)
other types(e.g. timestamp)	PIC X(n), PIC X(n) VARYING

^{*1:} If no USAGE is specified, host variable is regarded as DISPLAY.

Body of BOOL is PIC X(1). '1' for true and '0' for false.

You can use some pattern of digits for integer(see table), but if database sends big number with more digits than specified, behavior is undefined.

^{*2:} Type definition is added automatically on pre-compiling.

PostgreSQL data type

COBOL Host variable type

VALUE clause can't be used with VARYING. (Can be used with other types)

REDEFINE clause can be used, but it won't be validated on pre-compilation (The COBOL compiler will do this).

Handling Character Strings

To handle SQL character string data types, such as varchar and text, there is a possible way to declare the host variables.

The way is using the PIC X(n) VARYING type (we call it VARCHAR type from now on), which is a special type provided by ECOBPG. The definition on type VARCHAR is converted into a group item consists of named variables. A declaration like:

```
01 VAR PIC X(180) VARYING.
```

is converted into:

```
01 VAR.
49 LEN PIC S9(4) COMP-5.
49 ARR PIC X(180).
```

if --varchar-with-named-member option is used, it is converted into:

```
01 VAR.
49 VAR-LEN PIC S9(4) COMP-5.
49 VAR-ARR PIC X(180).
```

You can use level 1 to 48 for VARCHAR. Don't use level 49 variable right after VARCHAR variable. To use a VARCHAR host variable as an input for SQL statement, LEN must be set the length of the string included in ARR.

To use a VARCHAR host variable as an output of SQL statement, the variable must be declared in a sufficient length. If the length is insufficient, it can cause a buffer overrun.

PIC X(n) and VARCHAR host variables can also hold values of other SQL types, which will be stored in their string forms.

Accessing Special Data Types

ECOBPG doesn't have special support for date, timestamp, and interval types.

(ECPG has pgtypes, but ECOBPG doesn't.)

You can use PIC X(n) or VARCHAR for DB I/O with these types. See "Data Types" section in PostgreSQL's document.

bytea

Handling of bytea types is similar to VARCHAR. The definition of an array of type bytea is converted into a group item consists of named variables.

A declaration like:

```
01 var bytea(100).
```

is converted into:

```
01 var .
49 LEN PIC S9(9) COMP-5 .
49 ARR PIC X(100) .
```

if bytea-with-named-member option is used, it is converted into:

```
01 var .
49 var-LEN PIC S9(9) COMP-5 .
49 var-ARR PIC X(100) .
```

The data item ARR holds binary format data. Unlike VARCHAR, it is not affected by the locale or character encoding when processing data.

Other usage and prohibitions are the same as VARCHAR.



The bytea variable can only be used if byta_output is set to hex.

Host Variables with Nonprimitive Types

As a host variable you can also use arrays, typedefs, and group items.

Arrays

To create and use array variables, OCCURENCE syntax is provided by COBOL.

The typical use case is to retrieve multiple rows from a query result without using a cursor. Without an array, to process a query result consisting of multiple rows, it is required to use a cursor and the FETCH command. But with array host variables, multiple rows can be received at once. The length of the array has to be defined to be able to accommodate all rows, otherwise a buffer overrun will likely occur.

Following example scans the pg_database system table and shows all OIDs and names of the available databases:

```
EXEC SQL BEGIN DECLARE SECTION END-EXEC.

01 GROUP-ITEM.

05 DBID PIC S9(9) COMP OCCURS 8.

05 DBNAME PIC X(16) OCCURS 8.

01 I PIC S9(9) COMP.

EXEC SQL END DECLARE SECTION END-EXEC.

* Retrieve multiple rows into arrays at once.

EXEC SQL SELECT oid,datname INTO :DBID, :DBNAME FROM pg_database END-EXEC.

PERFORM VARYING I FROM 1 BY 1 UNTIL I > 8

DISPLAY "oid=" DBID(I) ", dbname=" DBNAME(I)

END-PERFORM.

EXEC SQL COMMIT END-EXEC.

EXEC SQL COMMIT END-EXEC.

EXEC SQL DISCONNECT ALL END-EXEC.
```

You can use member of array as simple host variable by specifying subscript of array. For specifying subscript, use C-style "[1]", not COBOL-style "(1)". But subscript starts with 1, according to COBOL syntax.

```
EXEC SQL BEGIN DECLARE SECTION END-EXEC.

01 GROUP-ITEM.

05 DBID PIC S9(9) COMP OCCURS 8.

EXEC SQL END DECLARE SECTION END-EXEC.

EXEC SQL CONNECT TO testdb END-EXEC.

EXEC SQL SELECT oid INTO :DBID[1] FROM pg_database where oid=1 END-EXEC.

DISPLAY "oid=" DBID(1)

EXEC SQL COMMIT END-EXEC.

EXEC SQL DISCONNECT ALL END-EXEC.
```

Group Item

A group item whose subordinate item names match the column names of a query result, can be used to retrieve multiple columns at once. The group item enables handling multiple column values in a single host variable.

The following example retrieves OIDs, names, and sizes of the available databases from the pg_database system table by using the pg_database_size() function. In this example, a group item variable dbinfo_t with members whose names match each column in the SELECT result is used to retrieve one result row without putting multiple host variables in the FETCH statement.

```
EXEC SQL BEGIN DECLARE SECTION END-EXEC.
    01 DBINFO-T TYPEDEF.
       02 OID PIC S9(9) COMP.
       02 DATNAME PIC X(65).
       02 DBSIZE PIC S9(18) COMP.
    01 DBVAL TYPE DBINFO-T.
EXEC SQL END DECLARE SECTION END-EXEC.
   EXEC SQL DECLARE curl CURSOR FOR SELECT oid, datname, pg_database_size(oid) AS size
FROM pg_database END-EXEC.
   EXEC SQL OPEN curl END-EXEC.
   when end of result set reached, break out of loop
    EXEC SQL WHENEVER NOT FOUND GOTO END-FETCH END-EXEC.
    PERFORM NO LIMIT
         Fetch multiple columns into one structure.
        EXEC SQL FETCH FROM curl INTO :DBVAL END-EXEC
         Print members of the structure.
        DISPLAY "oid=" OID ", datname=" DATNAME ", size=" DBSIZE
    END-PERFORM.
    END-FETCH.
    EXEC SQL CLOSE curl END-EXEC.
```

group item host variables "absorb" as many columns as the group item as subordinate items. Additional columns can be assigned to other host variables. For example, the above program could also be restructured like this, with the size variable outside the group item:

```
EXEC SQL BEGIN DECLARE SECTION END-EXEC.
   01 DBINFO-T TYPEDEF.
      02 OID PIC S9(9) COMP.
      02 DATNAME PIC X(65).
   01 DBVAL TYPE DBINFO-T.
   01 DBSIZE PIC S9(18) COMP.
EXEC SQL END DECLARE SECTION END-EXEC.
   EXEC SQL DECLARE curl CURSOR FOR SELECT oid, datname, pg_database_size(oid) AS size
FROM pg_database END-EXEC.
   EXEC SQL OPEN curl END-EXEC.
   when end of result set reached, break out of loop
   EXEC SQL WHENEVER NOT FOUND GOTO END-FETCH END-EXEC.
   PERFORM NO LIMIT
       Fetch multiple columns into one structure.
        EXEC SQL FETCH FROM curl INTO :DBVAL, :DBSIZE END-EXEC
        Print members of the structure.
       DISPLAY "oid=" OID ", datname=" DATNAME ", size=" DBSIZE
    END-PERFORM
```

```
FETCH-END.
EXEC SQL CLOSE curl END-EXEC.
```

You can use only non-nested group items for host variable of SQL statement. Declaration of nested group items are OK, but you must specify non-nested part of group items for SQL. (VARCHAR, is translated to group item on pre-compilation, is not considered as offense of this rule.) When using inner item of group item in SQL, use C-struct like period separated syntax(not COBOL's A OF B). Here is example.

```
EXEC SQL BEGIN DECLARE SECTION END-EXEC.

01 NESTED-GROUP.

02 CHILD1.

03 A PIC X(10).

03 B PIC S9(9) COMP.

02 CHILD2.

03 A PIC X(10).

03 B PIC S9(9) COMP.

EXEC SQL END DECLARE SECTION END-EXEC.

* This SQL is valid. CHILD1 has no nested group items.

EXEC SQL SELECT * INTO :NESTED-GROUP.CHILD1 FROM TABLE1 END-EXEC.
```

For specifying basic item of group items, full specification is not needed if the specification is enough for identifying the item. This is from COBOL syntax. For more detail, see resources of COBOL syntax.

TYPEDEF

Use the typedef keyword to map new types to already existing types.

```
EXEC SQL BEGIN DECLARE SECTION END-EXEC.

01 MYCHARTYPE TYPEDEF PIC X(40).

01 SERIAL-T TYPEDEF PIC S9(9) COMP.

EXEC SQL END DECLARE SECTION END-EXEC.
```

Note that you could also use:

```
EXEC SQL TYPE SERIAL-T IS PIC S9(9) COMP-5. END-EXEC.
```

This declaration does not need to be part of a declare section.

D.4.5 Handling Nonprimitive SQL Data Types

This section contains information on how to handle nonscalar and user-defined SQL-level data types in ECOBPG applications. Note that this is distinct from the handling of host variables of nonprimitive types, described in the previous section.

Arrays

SQL-level arrays are not directly supported in ECOBPG. It is not possible to simply map an SQL array into a COBOL array host variable. This will result in undefined behavior. Some workarounds exist, however.

If a query accesses elements of an array separately, then this avoids the use of arrays in ECOBPG. Then, a host variable with a type that can be mapped to the element type should be used. For example, if a column type is array of integer, a host variable of type PIC S9(9) COMP can be used. Also if the element type is varchar or text, a host variable of type VARCHAR can be used.

Here is an example. Assume the following table:

The following example program retrieves the 4th element of the array and stores it into a host variable of type PIC S9(9)

```
EXEC SQL BEGIN DECLARE SECTION END-EXEC.

01 II PIC S9(9) COMP.

EXEC SQL END DECLARE SECTION END-EXEC.

EXEC SQL DECLARE cur1 CURSOR FOR SELECT ii[4] FROM t3 END-EXEC.

EXEC SQL OPEN cur1 END-EXEC.

EXEC SQL WHENEVER NOT FOUND GOTO END-FETCH END-EXEC.

PERFORM NO LIMIT

EXEC SQL FETCH FROM cur1 INTO :II END-EXEC

DISPLAY "ii=" II

END-PERFORM.

END-FETCH.

EXEC SQL CLOSE cur1 END-EXEC.
```

To map multiple array elements to the multiple elements in an array type host variables each element of array column and each element of the host variable array have to be managed separately, for example:

```
EXEC SQL BEGIN DECLARE SECTION END-EXEC.

01 GROUP-ITEM.

05 II_A PIC S9(9) COMP OCCURS 8.

EXEC SQL END DECLARE SECTION END-EXEC.

EXEC SQL DECLARE curl CURSOR FOR SELECT ii[1], ii[2], ii[3], ii[4] FROM t3 END-EXEC.

EXEC SQL OPEN curl END-EXEC.

EXEC SQL WHENEVER NOT FOUND GOTO END-FETCH END-EXEC.

PERFORM NO LIMIT

EXEC SQL FETCH FROM curl INTO :II_A[1], :II_A[2], :II_A[3], :II_A[4] END-EXEC

...

END-PERFORM.
```

Note again that.

```
EXEC SQL BEGIN DECLARE SECTION END-EXEC.

01 GROUP-ITEM.

05 II_A PIC S9(9) COMP OCCURS 8.

EXEC SQL END DECLARE SECTION END-EXEC.

EXEC SQL DECLARE curl CURSOR FOR SELECT ii FROM t3 END-EXEC.

EXEC SQL OPEN curl END-EXEC.
```

```
EXEC SQL WHENEVER NOT FOUND GOTO END-FETCH END-EXEC.

PERFORM NO LIMIT

* WRONG

EXEC SQL FETCH FROM curl INTO :II_A END-EXEC

...

END-PERFORM.
```

would not work correctly in this case, because you cannot map an array type column to an array host variable directly.

Another workaround is to store arrays in their external string representation in host variables of type VARCHAR. For more details about this representation.



Refer to "Arrays" in "Tutorial" in the PostgreSQL Documentation for information more details about this representation.

Note that this means that the array cannot be accessed naturally as an array in the host program (without further processing that parses the text representation).

Composite Types

Composite types are not directly supported in ECOBPG, but an easy workaround is possible. The available workarounds are similar to the ones described for arrays above: Either access each attribute separately or use the external string representation.

For the following examples, assume the following type and table:

```
CREATE TYPE comp_t AS (intval integer, textval varchar(32));
CREATE TABLE t4 (compval comp_t);
INSERT INTO t4 VALUES ( (256, 'PostgreSQL') );
```

The most obvious solution is to access each attribute separately. The following program retrieves data from the example table by selecting each attribute of the type comp_t separately:

```
EXEC SQL BEGIN DECLARE SECTION END-EXEC.
01 INTVAL PIC S9(9) COMP.
01 TEXTVAL PIC X(33) VARYING.
EXEC SQL END DECLARE SECTION END-EXEC.
* Put each element of the composite type column in the SELECT list.
EXEC SQL DECLARE curl CURSOR FOR SELECT (compval).intval, (compval).textval FROM t4 END-
EXEC.
EXEC SOL OPEN curl END-EXEC.
EXEC SOL WHENEVER NOT FOUND GOTO END-FETCH END-EXEC.
PERFORM NO LIMIT
   Fetch each element of the composite type column into host variables.
   EXEC SQL FETCH FROM curl INTO :INTVAL, :TEXTVAL END-EXEC
   DISPLAY "intval=" INTVAL ", textval=" ARR OF TEXTVAL
END-PERFORM.
END-FETCH.
EXEC SQL CLOSE curl END-EXEC.
```

To enhance this example, the host variables to store values in the FETCH command can be gathered into one group item. For more details about the host variable in the group item form, see "Group Item". To switch to the group item, the example can be modified as below. The two host variables, intval and textval, become subordinate items of the comp_t group item, and the group item is specified on the FETCH command.

```
EXEC SQL BEGIN DECLARE SECTION END-EXEC.
01 COMP-T TYPEDEF.
   02 INTVAL PIC S9(9) COMP.
   02 TEXTVAL PIC X(33) VARYING.
01 COMPVAL TYPE COMP-T.
EXEC SQL END DECLARE SECTION END-EXEC.
* Put each element of the composite type column in the SELECT list.
EXEC SQL DECLARE curl CURSOR FOR SELECT (compval).intval, (compval).textval FROM t4 END-
EXEC.
EXEC SQL OPEN curl END-EXEC.
EXEC SQL WHENEVER NOT FOUND GOTO END-FETCH END-EXEC.
PERFORM NO LIMIT
   Put all values in the SELECT list into one structure.
   EXEC SQL FETCH FROM curl INTO : COMPVAL END-EXEC
   DISPLAY "intval=" INTVAL ", textval=" ARR OF TEXTVAL
END-PERFORM.
END-FETCH.
EXEC SQL CLOSE curl END-EXEC.
```

Although a group item is used in the FETCH command, the attribute names in the SELECT clause are specified one by one. This can be enhanced by using a * to ask for all attributes of the composite type value.

```
...

EXEC SQL DECLARE curl CURSOR FOR SELECT (compval).* FROM t4 END-EXEC.

EXEC SQL OPEN curl END-EXEC.

EXEC SQL WHENEVER NOT FOUND GOTO END-FETCH END-EXEC.

PERFORM NO LIMIT

* Put all values in the SELECT list into one structure.

EXEC SQL FETCH FROM curl INTO :COMPVAL END-EXEC

DISPLAY "intval=" INTVAL ", textval=" ARR OF TEXTVAL
END-PERFORM.
```

This way, composite types can be mapped into structures almost seamlessly, even though ECOBPG does not understand the composite type itself.

Finally, it is also possible to store composite type values in their external string representation in host variables of type VARCHAR. But that way, it is not easily possible to access the fields of the value from the host program.

User-defined Base Types

New user-defined base types are not directly supported by ECOBPG. You can use the external string representation and host variables of type VARCHAR, and this solution is indeed appropriate and sufficient for many types.

Here is an example using the data type complex.



See

Refer to "User-defined Types" in "Server Programming" in the PostgreSQL Documentation for information on the data type complex.

The external string representation of that type is (%lf,%lf), which is defined in the functions complex_in() and complex_out() functions. The following example inserts the complex type values (1,1) and (3,3) into the columns a and b, and select them from the table after that.

```
EXEC SQL BEGIN DECLARE SECTION END-EXEC.

01 A PIC X(64) VARYING.

01 B PIC X(64) VARYING.

EXEC SQL END DECLARE SECTION END-EXEC.

EXEC SQL INSERT INTO test_complex VALUES ('(1,1)', '(3,3)') END-EXEC.

EXEC SQL DECLARE curl CURSOR FOR SELECT a, b FROM test_complex END-EXEC.

EXEC SQL OPEN curl END-EXEC.

EXEC SQL WHENEVER NOT FOUND GOTO END-FETCH END-EXEC.

PERFORM NO LIMIT

EXEC SQL FETCH FROM curl INTO :A, :B END-EXEC

DISPLAY "a=" ARR OF A ", b=" ARR OF B

END-PERFORM.

END-FETCH.

EXEC SQL CLOSE curl END-EXEC.
```

Another workaround is avoiding the direct use of the user-defined types in ECOBPG and instead create a function or cast that converts between the user-defined type and a primitive type that ECOBPG can handle. Note, however, that type casts, especially implicit ones, should be introduced into the type system very carefully.

For example:

```
CREATE FUNCTION create_complex(r double precision, i double precision) RETURNS complex LANGUAGE SQL IMMUTABLE
AS $$ SELECT $1 * complex '(1,0)' + $2 * complex '(0,1)' $$;
```

After this definition, the following:

```
EXEC SQL BEGIN DECLARE SECTION END-EXEC.

01 A COMP-2.

01 B COMP-2.

01 C COMP-2.

01 D COMP-2.

EXEC SQL END DECLARE SECTION END-EXEC.

MOVE 1 TO A.

MOVE 2 TO B.

MOVE 3 TO C.

MOVE 4 TO D.

EXEC SQL INSERT INTO test_complex VALUES (create_complex(:A, :B), create_complex(:C, :D))

END-EXEC.
```

has the same effect as

```
EXEC SQL INSERT INTO test_complex VALUES ('(1,2)', '(3,4)') END-EXEC.
```

D.4.6 Indicators

The examples above do not handle null values. In fact, the retrieval examples will raise an error if they fetch a null value from the database. To be able to pass null values to the database or retrieve null values from the database, you need to append a second host variable specification to each host variable that contains data. This second host variable is called the *indicator* and contains a flag that tells whether the datum is null, in which case the value of the real host variable is ignored. Here is an example that handles the retrieval of null values correctly:

```
EXEC SQL BEGIN DECLARE SECTION END-EXEC.
01 VAL PIC X(50) VARYING.
01 VAL_IND PIC S9(9) COMP-5.
```

```
EXEC SQL END DECLARE SECTION END-EXEC.

...

EXEC SQL SELECT b INTO :VAL :VAL_IND FROM test1 END-EXEC.
```

The indicator variable val_ind will be zero if the value was not null, and it will be negative if the value was null.

The indicator has another function: if the indicator value is positive, it means that the value is not null, but it was truncated when it was stored in the host variable.

D.5 Dynamic SQL

In many cases, the particular SQL statements that an application has to execute are known at the time the application is written. In some cases, however, the SQL statements are composed at run time or provided by an external source. In these cases you cannot embed the SQL statements directly into the COBOL source code, but there is a facility that allows you to call arbitrary SQL statements that you provide in a string variable.

D.5.1 Executing Statements without a Result Set

The simplest way to execute an arbitrary SQL statement is to use the command EXECUTE IMMEDIATE. For example:

```
EXEC SQL BEGIN DECLARE SECTION END-EXEC.

01 STMT PIC X(30) VARYING.

EXEC SQL END DECLARE SECTION END-EXEC.

MOVE "CREATE TABLE test1 (...);" TO ARR OF STMT.

COMPUTE LEN OF STMT = FUNCTION STORED-CHAR-LENGTH (ARR OF STMT).

EXEC SQL EXECUTE IMMEDIATE :STMT END-EXEC.
```

EXECUTE IMMEDIATE can be used for SQL statements that do not return a result set (e.g., DDL, INSERT, UPDATE, DELETE). You cannot execute statements that retrieve data (e.g., SELECT) this way. The next section describes how to do that.

D.5.2 Executing a Statement with Input Parameters

A more powerful way to execute arbitrary SQL statements is to prepare them once and execute the prepared statement as often as you like. It is also possible to prepare a generalized version of a statement and then execute specific versions of it by substituting parameters. When preparing the statement, write question marks where you want to substitute parameters later. For example:

```
EXEC SQL BEGIN DECLARE SECTION END-EXEC.

01 STMT PIC X(40) VARYING.

EXEC SQL END DECLARE SECTION END-EXEC.

MOVE "INSERT INTO test1 VALUES(?, ?);" TO ARR OF STMT.

COMPUTE LEN OF STMT = FUNCTION STORED-CHAR-LENGTH (ARR OF STMT).

EXEC SQL PREPARE MYSTMT FROM :STMT END-EXEC.

...

EXEC SQL EXECUTE MYSTMT USING 42, 'foobar' END-EXEC.
```

When you don't need the prepared statement anymore, you should deallocate it:

```
EXEC SQL DEALLOCATE PREPARE name END-EXEC.
```

D.5.3 Executing a Statement with a Result Set

To execute an SQL statement with a single result row, EXECUTE can be used. To save the result, add an INTO clause.

```
EXEC SQL BEGIN DECLARE SECTION END-EXEC.
01 STMT PIC X(50) VARYING.
```

```
01 V1 PIC S9(9) COMP.
01 V2 PIC S9(9) COMP.
01 V3 PIC X(50) VARYING.

EXEC SQL END DECLARE SECTION END-EXEC.

MOVE "SELECT a, b, c FROM test1 WHERE a > ?" TO ARR OF STMT.

COMPUTE LEN OF STMT = FUNCTION STORED-CHAR-LENGTH (ARR OF STMT).

EXEC SQL PREPARE MYSTMT FROM :STMT END-EXEC.

...

EXEC SQL EXECUTE MYSTMT INTO :V1, :V2, :V3 USING 37 END-EXEC.
```

An EXECUTE command can have an INTO clause, a USING clause, both, or neither.

If a query is expected to return more than one result row, a cursor should be used, as in the following example. (See "D.3.2 Using Cursors" for more details about the cursor.)

```
EXEC SQL BEGIN DECLARE SECTION END-EXEC.
01 DBANAME PIC X(128) VARYING.
01 DATNAME PIC X(128) VARYING.
01 STMT PIC X(200) VARYING.
EXEC SQL END DECLARE SECTION END-EXEC.
MOVE "SELECT u.usename as dbaname, d.datname
             " FROM pg_database d, pg_user u
              " WHERE d.datdba = u.usesysid"
TO ARR OF STMT.
COMPUTE LEN OF STMT = FUNCTION STORED-CHAR-LENGTH (ARR OF STMT).
EXEC SQL CONNECT TO testdb AS con1 USER testuser END-EXEC.
EXEC SQL PREPARE STMT1 FROM :STMT END-EXEC.
EXEC SOL DECLARE cursor1 CURSOR FOR STMT1 END-EXEC.
EXEC SQL OPEN cursor1 END-EXEC.
EXEC SQL WHENEVER NOT FOUND GOTO FETCH-END END-EXEC.
PERFORM NO LIMIT
   EXEC SQL FETCH cursor1 INTO :DBANAME,:DATNAME END-EXEC
   DISPLAY "dbaname=" ARR OF DBANAME ", datname=" ARR OF DATNAME
END-PERFORM.
FETCH-END.
EXEC SQL CLOSE cursor1 END-EXEC.
EXEC SOL COMMIT END-EXEC.
EXEC SOL DISCONNECT ALL END-EXEC.
```

D.6 Using Descriptor Areas

An SQL descriptor area is a more sophisticated method for processing the result of a SELECT, FETCH or a DESCRIBE statement. An SQL descriptor area groups the data of one row of data together with metadata items into one data group item. The metadata is particularly useful when executing dynamic SQL statements, where the nature of the result columns might not be known ahead of time. PostgreSQL provides a way to use Descriptor Areas: the named SQL Descriptor Areas.

D.6.1 Named SQL Descriptor Areas

A named SQL descriptor area consists of a header, which contains information concerning the entire descriptor, and one or more item descriptor areas, which basically each describe one column in the result row.

Before you can use an SQL descriptor area, you need to allocate one:

```
EXEC SOL ALLOCATE DESCRIPTOR identifier END-EXEC.
```

The identifier serves as the "variable name" of the descriptor area. When you don't need the descriptor anymore, you should deallocate it:

```
EXEC SQL DEALLOCATE DESCRIPTOR identifier END-EXEC.
```

To use a descriptor area, specify it as the storage target in an INTO clause, instead of listing host variables:

```
EXEC SQL FETCH NEXT FROM mycursor INTO SQL DESCRIPTOR mydesc END-EXEC.
```

If the result set is empty, the Descriptor Area will still contain the metadata from the query, i.e. the field names.

For not yet executed prepared queries, the DESCRIBE statement can be used to get the metadata of the result set:

```
EXEC SQL BEGIN DECLARE SECTION END-EXEC.

01 SQL-STMT PIC X(30) VARYING.

EXEC SQL END DECLARE SECTION END-EXEC.

MOVE "SELECT * FROM table1" TO ARR OF SQL-STMT.

COMPUTE LEN OF SQL-STMT = FUNCTION STORED-CHAR-LENGTH ( ARR OF SQL-STMT ) .

EXEC SQL PREPARE STMT1 FROM :SQL-STMT END-EXEC.

EXEC SQL DESCRIBE STMT1 INTO SQL DESCRIPTOR MYDESC END-EXEC.
```

Before PostgreSQL 9.0, the SQL keyword was optional, so using DESCRIPTOR and SQL DESCRIPTOR produced named SQL Descriptor Areas. Now it is mandatory, omitting the SQL keyword is regarded as the syntax that produces SQLDA Descriptor Areas. However, ecobpg does not support SQLDA and it causes an error.

In DESCRIBE and FETCH statements, the INTO and USING keywords can be used to similarly: they produce the result set and the metadata in a Descriptor Area.

Now how do you get the data out of the descriptor area? You can think of the descriptor area as a group item with named fields. To retrieve the value of a field from the header and store it into a host variable, use the following command:

```
EXEC SQL GET DESCRIPTOR name :hostvar = field END-EXEC.
```

Currently, there is only one header field defined: COUNT, which tells how many item descriptor areas exist (that is, how many columns are contained in the result). The host variable needs to be of an integer type as PIC S9(9) COMP-5. To get a field from the item descriptor area, use the following command:

```
EXEC SQL GET DESCRIPTOR name VALUE num :hostvar = field END-EXEC.
```

num can be a host variable containing an integer as PIC S9(9) COMP-5.

hostvar must be PIC S9(9) COMP-5 if type of the field is integer. Possible fields are:

CARDINALITY (integer)

number of rows in the result set

DATA

actual data item (therefore, the data type of this field depends on the query)

```
DATETIME_INTERVAL_CODE (integer)
```

When TYPE is 9, DATETIME_INTERVAL_CODE will have a value of 1 for DATE, 2 for TIME, 3 for TIMESTAMP, 4 for TIME WITH TIME ZONE, or 5 for TIMESTAMP WITH TIME ZONE.

DATETIME_INTERVAL_PRECISION (integer)

not implemented

INDICATOR (integer)

the indicator (indicating a null value or a value truncation)

KEY_MEMBER (integer)

not implemented

```
LENGTH (integer)
```

length of the datum in characters

NAME (string)

name of the column

NULLABLE (integer)

not implemented

OCTET_LENGTH (integer)

length of the character representation of the datum in bytes

PRECISION (integer)

precision (for type numeric)

RETURNED_LENGTH (integer)

length of the datum in characters

RETURNED_OCTET_LENGTH (integer)

length of the character representation of the datum in bytes

SCALE (integer)

scale (for type numeric)

TYPE (integer)

numeric code of the data type of the column

In EXECUTE, DECLARE and OPEN statements, the effect of the INTO and USING keywords are different. A Descriptor Area can also be manually built to provide the input parameters for a query or a cursor and USING SQL DESCRIPTOR name is the way to pass the input parameters into a parametrized query. The statement to build a named SQL Descriptor Area is below:

```
EXEC SQL SET DESCRIPTOR name VALUE num field = :hostvar END-EXEC.
```

PostgreSQL supports retrieving more than one record in one FETCH statement and storing the data in host variables in this case assumes that the variable is an array. E.g.:

```
EXEC SQL BEGIN DECLARE SECTION END-EXEC.

01 GROUP-ITEM.

05 IDNUM PIC S9(9) COMP OCCURS 5.

EXEC SQL END DECLARE SECTION END-EXEC.

EXEC SQL FETCH 5 FROM mycursor INTO SQL DESCRIPTOR mydesc END-EXEC.

EXEC SQL GET DESCRIPTOR mydesc VALUE 1 :IDNUM = DATA END-EXEC.
```

D.7 Error Handling

This section describes how you can handle exceptional conditions and warnings in an embedded SQL program. There are two nonexclusive facilities for this.

- Callbacks can be configured to handle warning and error conditions using the WHENEVER command.
- Detailed information about the error or warning can be obtained from the sqlca variable.

D.7.1 Setting Callbacks

One simple method to catch errors and warnings is to set a specific action to be executed whenever a particular condition occurs. In general:

```
EXEC SQL WHENEVER condition action END-EXEC.
```

condition can be one of the following:

SQLERROR

The specified action is called whenever an error occurs during the execution of an SQL statement.

SQLWARNING

The specified action is called whenever a warning occurs during the execution of an SQL statement.

NOT FOUND

The specified action is called whenever an SQL statement retrieves or affects zero rows. (This condition is not an error, but you might be interested in handling it specially.)

action can be one of the following:

CONTINUE

This effectively means that the condition is ignored. This is the default.

GOTO label

GO TO label

Jump to the specified label (using a COBOL goto statement).

SQLPRINT

Print a message to standard error. This is useful for simple programs or during prototyping. The details of the message cannot be configured.

STOP

Call STOP, which will terminate the program.

CALL name usingargs

DO name usingargs

Call the specified functions with the following characters including arguments. Thus, syntaxes (including compiler depending) are able to be placed as well as the arguments. Though, there are some limitation as following:

- You can't use RETURNING, ON EXCEPTION or OVER FLOW clauses.
- In the called subprogram, you must specify CONTINUE for every action with WHENEVER statement.

The SQL standard only provides for the actions CONTINUE and GOTO (and GO TO).

Here is an example that you might want to use in a simple program. It prints a simple message when a warning occurs and aborts the program when an error happens:

```
EXEC SQL WHENEVER SQLWARNING SQLPRINT END-EXEC.
EXEC SQL WHENEVER SQLERROR STOP END-EXEC.
```

The statement EXEC SQL WHENEVER is a directive of the SQL preprocessor, not a COBOL statement. The error or warning actions that it sets apply to all embedded SQL statements that appear below the point where the handler is set, unless a different action was set for the same condition between the first EXEC SQL WHENEVER and the SQL statement causing the condition, regardless of the flow of control in the COBOL program. So neither of the two following COBOL program excerpts will have the desired effect:

```
*

* WRONG

*

...

IF VERBOSE = 1 THEN

EXEC SQL WHENEVER SQLWARNING SQLPRINT END-EXEC

END-IF.
```

```
EXEC SQL SELECT ... END-EXEC.
...

* WRONG

* ...
CALL SET-ERROR-HANDLER.

* (and execute "EXEC SQL WHENEVER SQLERROR STOP" in SET-ERROR-HANDLER)
...
EXEC SQL SELECT ... END-EXEC.
...
```

D.7.2 sqlca

For more powerful error handling, the embedded SQL interface provides a global variable with the name sqlca (SQL communication area) that has the following group item:

```
01 sqlca_t.
    10 sqlcaid PIC X(8).
    10 sqlabc PIC S9(9) COMP-5.
    10 sqlcode PIC S9(9) COMP-5.
    10 sqlerrm.
        20 sqlerrml PIC S9(9) COMP-5.
        20 sqlerrmc PIC X(150).
    10 sqlerrp PIC X(8).
    10 sqlerrd PIC S9(9) COMP-5 OCCURS 6.
    10 sqlwarn PIC X(8).
    10 sqlstate PIC X(5).
```

sqlca covers both warnings and errors. If multiple warnings or errors occur during the execution of a statement, then sqlca will only contain information about the last one.

If no error occurred in the last SQL statement, SQLCODE will be 0 and SQLSTATE will be "00000". If a warning or error occurred, then SQLCODE will be negative and SQLSTATE will be different from "00000". A positive SQLCODE indicates a harmless condition, such as that the last query returned zero rows. SQLCODE and SQLSTATE are two different error code schemes; details appear below.

If the last SQL statement was successful, then SQLERRD(2) contains the OID of the processed row, if applicable, and SQLERRD(3) contains the number of processed or returned rows, if applicable to the command.

In case of an error or warning, SQLERRMC will contain a string that describes the error. The field SQLERRML contains the length of the error message that is stored in SQLERRMC (the result of FUNCTION STORED-CHAR-LENGTH. Note that some messages are too long to fit in the fixed-size sqlerrmc array; they will be truncated.

In case of a warning, the 3rd character of SQLWARN is set to W. (In all other cases, it is set to something different from W.) If the 2nd character of SQLWARN is set to W, then a value was truncated when it was stored in a host variable. The 1st character of SQLWARN is set to W if any of the other elements are set to indicate a warning.

The fields sqlcaid, sqlcabc, sqlerrp, and the remaining elements of sqlerrd and sqlwarn currently contain no useful information.

The structure sqlca is not defined in the SQL standard, but is implemented in several other SQL database systems. The definitions are similar at the core, but if you want to write portable applications, then you should investigate the different implementations carefully.

Here is one example that combines the use of WHENEVER and sqlca, printing out the contents of sqlca when an error occurs. This is perhaps useful for debugging or prototyping applications, before installing a more "user-friendly" error handler.

```
EXEC SQL WHENEVER SQLERROR GOTO PRINT_SQLCA END-EXEC.

PRINT_SQLCA.

DISPLAY "==== sqlca ====".

DISPLAY "SQLCODE: " SQLCODE.
```

```
DISPLAY "SQLERRML: " SQLERRML.
   DISPLAY "SQLERRMC: " SQLERRMC.
   DISPLAY "SQLERRD: " SQLERRD(1) " " SQLERRD(2) " " SQLERRD(3)" " SQLERRD(4) " "
SQLERRD(5) " " SQLERRD(6).
   DISPLAY "SQLSTATE: " SQLSTATE.
   DISPLAY "=======".
```

The result could look as follows (here an error due to a misspelled table name):

```
==== sglca ====
salcode: -000000400
SOLERRML: +000000064
SQLERRMC: relation "pg_databasep" does not exist (10292) on line 93
sqlstate: 42P01
_____
```

D.7.3 SQLSTATE vs. SQLCODE

The fields SQLSTATE and SQLCODE are two different schemes that provide error codes. Both are derived from the SQL standard, but SQLCODE has been marked deprecated in the SQL-92 edition of the standard and has been dropped in later editions. Therefore, new applications are strongly encouraged to use SQLSTATE.

SQLSTATE is a five-character array. The five characters contain digits or upper-case letters that represent codes of various error and warning conditions. SQLSTATE has a hierarchical scheme: the first two characters indicate the general class of the condition, the last three characters indicate a subclass of the general condition. A successful state is indicated by the code 00000. The SQLSTATE codes are for the most part defined in the SQL standard. The PostgreSQL server natively supports SQLSTATE error codes; therefore a high degree of consistency can be achieved by using this error code scheme throughout all applications.



Refer to "PostgreSQL Error Codes" in "Appendixes" in the PostgreSQL Documentation for further information.

SQLCODE, the deprecated error code scheme, is a simple integer. A value of 0 indicates success, a positive value indicates success with additional information, and a negative value indicates an error. The SQL standard only defines the positive value +100, which indicates that the last command returned or affected zero rows, and no specific negative values. Therefore, this scheme can only achieve poor portability and does not have a hierarchical code assignment. Historically, the embedded SQL processor for PostgreSQL has assigned some specific SQLCODE values for its use, which are listed below with their numeric value and their symbolic name. Remember that these are not portable to other SQL implementations. To simplify the porting of applications to the SQLSTATE scheme, the corresponding SQLSTATE is also listed. There is, however, no one-to-one or one-to-many mapping between the two schemes (indeed it is many-to-many), so you should consult the global SQLSTATE in each case.



Refer to "PostgreSQL Error Codes" in "Appendixes" in the PostgreSQL Documentation.

These are the assigned SQLCODE values:

Indicates no error. (SQLSTATE 00000)

100

This is a harmless condition indicating that the last command retrieved or processed zero rows, or that you are at the end of the cursor. (SQLSTATE 02000)

When processing a cursor in a loop, you could use this code as a way to detect when to abort the loop, like this:

```
PERFORM NO LIMIT

EXEC SQL FETCH ... END-EXEC

IF SQLCODE = 100 THEN

GO TO FETCH-END

END-IF

END-PERFORM.
```

But WHENEVER NOT FOUND GOTO ... effectively does this internally, so there is usually no advantage in writing this out explicitly.

-12

Indicates that your virtual memory is exhausted. The numeric value is defined as -ENOMEM. (SQLSTATE YE001)

-200

Indicates the preprocessor has generated something that the library does not know about. Perhaps you are running incompatible versions of the preprocessor and the library. (SQLSTATE YE002)

-201

This means that the command specified more host variables than the command expected. (SQLSTATE 07001 or 07002)

-202

This means that the command specified fewer host variables than the command expected. (SQLSTATE 07001 or 07002)

-203

This means a query has returned multiple rows but the statement was only prepared to store one result row (for example, because the specified variables are not arrays). (SQLSTATE 21000)

-204

The host variable is of type signed int and the datum in the database is of a different type and contains a value that cannot be interpreted as a signed int. The library uses strtol() for this conversion. (SQLSTATE 42804)

-205

The host variable is of type unsigned int and the datum in the database is of a different type and contains a value that cannot be interpreted as an unsigned int. The library uses strtoul() for this conversion. (SQLSTATE 42804)

-206

The host variable is of type float and the datum in the database is of another type and contains a value that cannot be interpreted as a float. The library uses strtod() for this conversion. (SQLSTATE 42804)

-207

The host variable is of type DECIMAL and the datum in the database is of another type and contains a value that cannot be interpreted as a DECIMAL or DISPLAY value. For the case of DISPLAY, this error happens if values in the database is too large for converting to DISPLAY value. (SQLSTATE 42804)

-208

The host variable is of type interval and the datum in the database is of another type and contains a value that cannot be interpreted as an interval value. (SQLSTATE 42804)

-209

The host variable is of type date and the datum in the database is of another type and contains a value that cannot be interpreted as a date value. (SQLSTATE 42804)

-210

The host variable is of type timestamp and the datum in the database is of another type and contains a value that cannot be interpreted as a timestamp value. (SQLSTATE 42804)

-211

This means the host variable is of type bool and the datum in the database is neither 't' nor 'f'. (SQLSTATE 42804)

-212

The statement sent to the PostgreSQL server was empty. (This cannot normally happen in an embedded SQL program, so it might point to an internal error.) (SQLSTATE YE002)

-213

A null value was returned and no null indicator variable was supplied. (SQLSTATE 22002)

-214

An ordinary variable was used in a place that requires an array. (SQLSTATE 42804)

-215

The database returned an ordinary variable in a place that requires array value. (SQLSTATE 42804)

-220

The program tried to access a connection that does not exist. (SQLSTATE 08003)

-221

The program tried to access a connection that does exist but is not open. (This is an internal error.) (SQLSTATE YE002)

-230

The statement you are trying to use has not been prepared. (SQLSTATE 26000)

-240

The descriptor specified was not found. The statement you are trying to use has not been prepared. (SQLSTATE 33000)

-241

The descriptor index specified was out of range. (SQLSTATE 07009)

-242

An invalid descriptor item was requested. (This is an internal error.) (SQLSTATE YE002)

-243

During the execution of a dynamic statement, the database returned a numeric value and the host variable was not numeric. (SQLSTATE 07006)

-244

During the execution of a dynamic statement, the database returned a non-numeric value and the host variable was numeric. (SQLSTATE 07006)

-400

 $Some\ error\ caused\ by\ the\ PostgreSQL\ server.\ The\ message\ contains\ the\ error\ message\ from\ the\ PostgreSQL\ server.$

-401

 $The\ PostgreSQL\ server\ signaled\ that\ we\ cannot\ start,\ commit,\ or\ rollback\ the\ transaction.\ (SQLSTATE\ 08007)$

-402

The connection attempt to the database did not succeed. (SQLSTATE 08001)

-403

Duplicate key error, violation of unique constraint. (SQLSTATE 23505)

-404

A result for the subquery is not single row. (SQLSTATE 21000)

-602

An invalid cursor name was specified. (SQLSTATE 34000)

-603

Transaction is in progress. (SQLSTATE 25001)

-604

There is no active (in-progress) transaction. (SQLSTATE 25P01)

-605

An existing cursor name was specified. (SQLSTATE 42P03)

D.8 Preprocessor Directives

Several preprocessor directives are available that modify how the ecobpg preprocessor parses and processes a file.

D.8.1 Including Files

To include an external file into your embedded SQL program, use:

```
EXEC SQL INCLUDE filename END-EXEC.

EXEC SQL INCLUDE filename
EXEC.

EXEC SQL INCLUDE "filename
END-EXEC.
```

The embedded SQL preprocessor will look for a file named filename.pco, preprocess it, and include it in the resulting COBOL output. Thus, embedded SQL statements in the included file are handled correctly.

By default, the ecobpg preprocessor will search a file at the current directory. This behavior can be changed by the ecobpg commandline option.

First, the preprocessor tries to locate a file by specified file name at the current directory. If it fails and the file name does not end with .pco, the preprocessor also tries to locate a file with the suffix at the same directory.

The difference between EXEC SQL INCLUDE and COPY statement is whether precompiler processes embedded SQLs in the file, or not. If the file contains embedded SQLs, use EXEC SQL INCLUDE.



The include file name is case-sensitive, even though the rest of the EXEC SQL INCLUDE command follows the normal SQL case-sensitivity rules.

D.8.2 The define and undef Directives

Similar to the directive #define that is known from C, embedded SQL has a similar concept:

```
EXEC SQL DEFINE name END-EXEC.
EXEC SQL DEFINE name value END-EXEC.
```

So you can define a name:

```
EXEC SQL DEFINE HAVE_FEATURE END-EXEC.
```

And you can also define constants:

```
EXEC SQL DEFINE MYNUMBER 12 END-EXEC.
EXEC SQL DEFINE MYSTRING 'abc' END-EXEC.
```

Use undef to remove a previous definition:

```
EXEC SQL UNDEF MYNUMBER END-EXEC.
```

Note that a constant in the SQL statement is only replaced by EXEC SQL DEFINE. The replacement may change the number of characters in a line, but ecobpg does not validate it after the replacement. Pay attention to the limitation of the number of characters in a line.

D.8.3 ifdef, ifndef, else, elif, and endif Directives

You can use the following directives to compile code sections conditionally:

```
EXEC SQL ifdef name END-EXEC.
```

Checks a name and processes subsequent lines if name has been created with EXEC SQL define name.

```
EXEC SQL ifndef name END-EXEC.
```

Checks a name and processes subsequent lines if name has not been created with EXEC SQL define name.

```
EXEC SQL else END-EXEC.
```

Starts processing an alternative section to a section introduced by either EXEC SQL ifdef name or EXEC SQL ifndef name.

```
EXEC SQL elif name END-EXEC.
```

Checks name and starts an alternative section if name has been created with EXEC SQL define name.

```
EXEC SQL endif END-EXEC.
```

Ends an alternative section.

Example:

```
EXEC SQL ifndef TZVAR END-EXEC.

EXEC SQL SET TIMEZONE TO 'GMT' END-EXEC.

EXEC SQL elif TZNAME END-EXEC.

EXEC SQL SET TIMEZONE TO TZNAME END-EXEC.

EXEC SQL else END-EXEC.

EXEC SQL SET TIMEZONE TO TZVAR END-EXEC.

EXEC SQL endif END-EXEC.
```

D.9 Processing Embedded SQL Programs

Now that you have an idea how to form embedded SQL COBOL programs, you probably want to know how to compile them. Before compiling you run the file through the embedded SQL COBOL preprocessor, which converts the SQL statements you used to special function calls. After compiling, you must link with a special library that contains the needed functions. These functions fetch information from the arguments, perform the SQL command using the libpq interface, and put the result in the arguments specified for output.

The preprocessor program is called ecobpg and is included in a normal PostgreSQL installation. Embedded SQL programs are typically named with an extension .pco. If you have a program file called prog1.pco, you can preprocess it by simply calling:

```
ecobpg progl.pco
```

This will create a file called prog1.cob. If your input files do not follow the suggested naming pattern, you can specify the output file explicitly using the -o option.

The preprocessed file can be compiled normally, following the usage of the compiler.

The generated COBOL source files include library files from the PostgreSQL installation, so if you installed PostgreSQL in a location that is not searched by default, you have to add an option such as -I/usr/local/pgsql/include to the compilation command line.

To link an embedded SQL program, you need to include the libecpg library.

Again, you might have to add an option for library search like -L/usr/local/pgsql/lib to that command line.

If you manage the build process of a larger project using make, it might be convenient to include the following implicit rule to your makefiles:

The complete syntax of the ecobpg command is detailed in "D.12.1 ecobpg".

Currently, ecobpg does not support multi threading.

D.10 Large Objects

Large objects are not supported by ECOBPG.

If you need to access large objects, use large objects interfaces of libpq instead.

D.11 Embedded SQL Commands

This section describes all SQL commands that are specific to embedded SQL.

Command	Description
ALLOCATE DESCRIPTOR	Allocate an SQL descriptor area
CONNECT	Establish a database connection
DEALLOCATE DESCRIPTOR	Deallocate an SQL descriptor area
DECLARE	Define a cursor
DESCRIBE	Obtain information about a prepared statement or result set
DISCONNECT	Terminate a database connection
EXECUTE IMMEDIATE	Dynamically prepare and execute a statement
GET DESCRIPTOR	Get information from an SQL descriptor area
OPEN	Open a dynamic cursor
PREPARE	Prepare a statement for execution
SET AUTOCOMMIT	Set the autocommit behavior of the current session
SET CONNECTION	Select a database connection
SET DESCRIPTOR	Set information in an SQL descriptor area
TYPE	Define a new data type
VAR	Define a variable
WHENEVER	Specify the action to be taken when an SQL statement causes a specific class condition to be raised



Refer to the SQL commands listed in "SQL Commands" under "Reference" in the PostgreSQL Documentation, which can also be used in embedded SQL, unless stated otherwise.

D.11.1 ALLOCATE DESCRIPTOR

Name

ALLOCATE DESCRIPTOR -- allocate an SQL descriptor area

Synopsis

ALLOCATE DESCRIPTOR name

Description

ALLOCATE DESCRIPTOR allocates a new named SQL descriptor area, which can be used to exchange data between the PostgreSQL server and the host program.

Descriptor areas should be freed after use using the DEALLOCATE DESCRIPTOR command.

Parameters

name

A name of SQL descriptor. This can be an SQL identifier or a host variable.

Examples

```
EXEC SQL ALLOCATE DESCRIPTOR mydesc END-EXEC.
```

Compatibility

ALLOCATE DESCRIPTOR is specified in the SQL standard.

See Also

DEALLOCATE DESCRIPTOR, GET DESCRIPTOR, SET DESCRIPTOR

D.11.2 CONNECT

Name

CONNECT -- establish a database connection

Synopsis

```
CONNECT TO connection_target [ AS connection_name ] [ USER connection_user_name ]

CONNECT TO DEFAULT

CONNECT connection_user_name

DATABASE connection_target
```

Description

The CONNECT command establishes a connection between the client and the PostgreSQL server.

Parameters

```
connection_target

connection_target specifies the target server of the connection on one of several forms.

[ database_name ] [ @host ] [ :port ]

Connect over TCP/IP

unix:postgresql://host [ :port ] / [ database_name ] [ ?connection_option ]

Connect over Unix-domain sockets

tcp:postgresql://host [ :port ] / [ database_name ] [ ?connection_option ]

Connect over TCP/IP
```

SQL string constant

containing a value in one of the above forms

host variable

host variable of fixed-length string (trailing spaces are ignored) containing a value in one of the above forms

connection name

An optional identifier for the connection, so that it can be referred to in other commands. This can be an SQL identifier or a host variable.

connection_user_name

The user name for the database connection.

This parameter can also specify user name and password, using one the forms user_name/password, user_name IDENTIFIED BY password, or user_name USING password.

User name and password can be SQL identifiers, string constants, or host variables (fixed-length string, trailing spaces are ignored).

DEFAULT

Use all default connection parameters, as defined by libpq.

Examples

Here a several variants for specifying connection parameters:

```
EXEC SQL CONNECT TO "connectdb" AS main END-EXEC.
EXEC SQL CONNECT TO "connectdb" AS second END-EXEC.
EXEC SQL CONNECT TO "unix:postgresql://localhost/connectdb" AS main USER connectuser END-
EXEC.
EXEC SQL CONNECT TO 'connectdb' AS main END-EXEC.
EXEC SQL CONNECT TO 'unix:postgresql://localhost/connectdb' AS main USER :user END-EXEC.
EXEC SQL CONNECT TO :dbn AS :idt END-EXEC.
EXEC SQL CONNECT TO :dbn USER connectuser USING :pw END-EXEC.
EXEC SQL CONNECT TO @localhost AS main USER connectdb END-EXEC.
EXEC SQL CONNECT TO REGRESSDB1 as main END-EXEC.
EXEC SOL CONNECT TO connectdb AS :idt END-EXEC.
EXEC SQL CONNECT TO connectdb AS main USER connectuser/connectdb END-EXEC.
EXEC SQL CONNECT TO connectdb AS main END-EXEC.
EXEC SQL CONNECT TO connectdb@localhost AS main END-EXEC.
EXEC SQL CONNECT TO tcp:postgresql://localhost/ USER connectdb END-EXEC.
EXEC SQL CONNECT TO tcp:postgresql://localhost/connectdb USER connectuser IDENTIFIED BY
connectpw END-EXEC.
EXEC SQL CONNECT TO tcp:postgresql://localhost:20/connectdb USER connectuser IDENTIFIED BY
connectpw END-EXEC.
EXEC SQL CONNECT TO unix:postgresql://localhost/ AS main USER connectdb END-EXEC.
EXEC SQL CONNECT TO unix:postgresgl://localhost/connectdb AS main USER connectuser END-
EXEC SQL CONNECT TO unix:postgresql://localhost/connectdb USER connectuser IDENTIFIED BY
"connectpw" END-EXEC.
EXEC SQL CONNECT TO unix:postgresql://localhost/connectdb USER connectuser USING
"connectow" END-EXEC.
EXEC SQL CONNECT TO unix:postgresql://localhost/connectdb?connect_timeout=14 USER
connectuser END-EXEC.
```

Here is an example program that illustrates the use of host variables to specify connection parameters:

```
EXEC SQL BEGIN DECLARE SECTION END-EXEC.

* database name

01 DBNAME PIC X(6).

* connection user name
```

```
01 USER PIC X(8).
   connection string
   01 CONNECTION PIC X(38).
   01 VER PIC X(256).
EXEC SQL END DECLARE SECTION END-EXEC.
   MOVE "testdb" TO DBNAME.
   MOVE "testuser" TO USER.
   MOVE "tcp:postgresql://localhost:5432/testdb" TO CONNECTION.
   EXEC SQL CONNECT TO :DBNAME USER :USER END-EXEC.
   EXEC SQL SELECT version() INTO :VER END-EXEC.
   EXEC SQL DISCONNECT END-EXEC.
   DISPLAY "version: " VER.
   EXEC SQL CONNECT TO :CONNECTION USER :USER END-EXEC.
   EXEC SQL SELECT version() INTO :VER END-EXEC.
   EXEC SQL DISCONNECT END-EXEC.
   DISPLAY "version: " VER.
```

Compatibility

CONNECT is specified in the SQL standard, but the format of the connection parameters is implementation-specific.

See Also

DISCONNECT, SET CONNECTION

D.11.3 DEALLOCATE DESCRIPTOR

Name

DEALLOCATE DESCRIPTOR -- deallocate an SQL descriptor area

Synopsis

DEALLOCATE DESCRIPTOR name

Description

DEALLOCATE DESCRIPTOR deallocates a named SQL descriptor area.

Parameters

name

The name of the descriptor which is going to be deallocated. This can be an SQL identifier or a host variable.

Examples

EXEC SQL DEALLOCATE DESCRIPTOR mydesc END-EXEC.

Compatibility

DEALLOCATE DESCRIPTOR is specified in the SQL standard.

See Also

ALLOCATE DESCRIPTOR, GET DESCRIPTOR, SET DESCRIPTOR

D.11.4 DECLARE

Name

DECLARE -- define a cursor

Synopsis

```
DECLARE cursor_name [ BINARY ] [ INSENSITIVE ] [ [ NO ] SCROLL ] CURSOR [ { WITH | WITHOUT } HOLD ] FOR prepared_name

DECLARE cursor_name [ BINARY ] [ INSENSITIVE ] [ [ NO ] SCROLL ] CURSOR [ { WITH | WITHOUT } HOLD ] FOR query
```

Description

DECLARE declares a cursor for iterating over the result set of a prepared statement. This command has slightly different semantics from the direct SQL command DECLARE: Whereas the latter executes a query and prepares the result set for retrieval, this embedded SQL command merely declares a name as a "loop variable" for iterating over the result set of a query; the actual execution happens when the cursor is opened with the OPEN command.

Parameters

cursor_name

A cursor name. This can be an SQL identifier or a host variable.

prepared_name

The name of a prepared query, either as an SQL identifier or a host variable.

query

A SELECT or VALUES command which will provide the rows to be returned by the cursor.

For the meaning of the cursor options, see DECLARE.



See

Refer to "SQL Commands" in "Reference" in the PostgreSQL Documentation for information on the SELECT, VALUES and DECLARE command.

Examples

Examples declaring a cursor for a query:

```
EXEC SQL DECLARE C CURSOR FOR SELECT * FROM My_Table END-EXEC.

EXEC SQL DECLARE C CURSOR FOR SELECT Item1 FROM T END-EXEC.

EXEC SQL DECLARE cur1 CURSOR FOR SELECT version() END-EXEC.
```

An example declaring a cursor for a prepared statement:

```
EXEC SQL PREPARE stmt1 AS SELECT version() END-EXEC.

EXEC SQL DECLARE cur1 CURSOR FOR stmt1 END-EXEC.
```

Compatibility

DECLARE is specified in the SQL standard.

See Also

OPEN, CLOSE, DECLARE



Refer to "SQL Commands" in "Reference" in the PostgreSQL Documentation for information on the CLOSE and DECLARE command.

D.11.5 DESCRIBE

Name

DESCRIBE -- obtain information about a prepared statement or result set

Synopsis

```
DESCRIBE [ OUTPUT ] prepared_name USING SQL DESCRIPTOR descriptor_name

DESCRIBE [ OUTPUT ] prepared_name INTO SQL DESCRIPTOR descriptor_name
```

Description

DESCRIBE retrieves metadata information about the result columns contained in a prepared statement, without actually fetching a row.

Parameters

prepared_name

The name of a prepared statement. This can be an SQL identifier or a host variable.

descriptor_name

A descriptor name. It can be an SQL identifier or a host variable.

Examples

```
EXEC SQL ALLOCATE DESCRIPTOR mydesc END-EXEC.

EXEC SQL PREPARE stmt1 FROM :sql_stmt END-EXEC.

EXEC SQL DESCRIBE stmt1 INTO SQL DESCRIPTOR mydesc END-EXEC.

EXEC SQL GET DESCRIPTOR mydesc VALUE 1 :charvar = NAME END-EXEC.

EXEC SQL DEALLOCATE DESCRIPTOR mydesc END-EXEC.
```

Compatibility

DESCRIBE is specified in the SQL standard.

See Also

ALLOCATE DESCRIPTOR, GET DESCRIPTOR

D.11.6 DISCONNECT

Name

DISCONNECT -- terminate a database connection

Synopsis

```
DISCONNECT connection_name
DISCONNECT [ CURRENT ]
DISCONNECT DEFAULT
```

Description

DISCONNECT closes a connection (or all connections) to the database.

Parameters

connection_name

A database connection name established by the CONNECT command.

CURRENT

Close the "current" connection, which is either the most recently opened connection, or the connection set by the SET CONNECTION command. This is also the default if no argument is given to the DISCONNECT command.

DEFAULT

Close the default connection.

ALL

Close all open connections.

Examples

```
EXEC SQL CONNECT TO testdb AS DEFAULT USER testuser END-EXEC.

EXEC SQL CONNECT TO testdb AS con1 USER testuser END-EXEC.

EXEC SQL CONNECT TO testdb AS con2 USER testuser END-EXEC.

EXEC SQL CONNECT TO testdb AS con3 USER testuser END-EXEC.

* close con3

EXEC SQL DISCONNECT CURRENT END-EXEC.

* close DEFAULT

EXEC SQL DISCONNECT DEFAULT END-EXEC.

* close con2 and con1

EXEC SQL DISCONNECT ALL END-EXEC.
```

Compatibility

DISCONNECT is specified in the SQL standard.

See Also

CONNECT, SET CONNECTION

D.11.7 EXECUTE IMMEDIATE

Name

EXECUTE IMMEDIATE -- dynamically prepare and execute a statement

Synopsis

EXECUTE IMMEDIATE string

Description

EXECUTE IMMEDIATE immediately prepares and executes a dynamically specified SQL statement, without retrieving result rows.

Parameters

string

A literal string or a host variable containing the SQL statement to be executed.

Examples

Here is an example that executes an INSERT statement using EXECUTE IMMEDIATE and a host variable named command:

```
MOVE "INSERT INTO test (name, amount, letter) VALUES ('db: ''rl''', 1, 'f')" TO ARR OF cmd. COMPUTE LEN OF cmd = FUNCTION STORED-CHAR-LENGTH(ARR OF cmd). EXEC SQL EXECUTE IMMEDIATE :cmd END-EXEC.
```

Compatibility

EXECUTE IMMEDIATE is specified in the SQL standard.

D.11.8 GET DESCRIPTOR

Name

GET DESCRIPTOR -- get information from an SQL descriptor area

Synopsis

```
GET DESCRIPTOR descriptor_name :hostvariable = descriptor_header_item [, ...]

GET DESCRIPTOR descriptor_name VALUE column_number :hostvariable = descriptor_item [, ...]
```

Description

GET DESCRIPTOR retrieves information about a query result set from an SQL descriptor area and stores it into host variables. A descriptor area is typically populated using FETCH or SELECT before using this command to transfer the information into host language variables.

This command has two forms: The first form retrieves descriptor "header" items, which apply to the result set in its entirety. One example is the row count. The second form, which requires the column number as additional parameter, retrieves information about a particular column. Examples are the column name and the actual column value.

Parameters

descriptor_name

A descriptor name.

descriptor_header_item

A token identifying which header information item to retrieve. Only COUNT, to get the number of columns in the result set, is currently supported.

column_number

The number of the column about which information is to be retrieved. The count starts at 1.

descriptor_item

A token identifying which item of information about a column to retrieve. See Section 33.7.1 for a list of supported items.

hostvariable

A host variable that will receive the data retrieved from the descriptor area.

Examples

An example to retrieve the number of columns in a result set:

```
EXEC SQL GET DESCRIPTOR d :d_count = COUNT END-EXEC.
```

An example to retrieve a data length in the first column:

```
EXEC SQL GET DESCRIPTOR d VALUE 1 :d_returned_octet_length = RETURNED_OCTET_LENGTH END-EXEC.
```

An example to retrieve the data body of the second column as a string:

```
EXEC SQL GET DESCRIPTOR d VALUE 2 :d_data = DATA END-EXEC.
```

Here is an example for a whole procedure of executing SELECT current_database(); and showing the number of columns, the column data length, and the column data:

```
EXEC SQL BEGIN DECLARE SECTION END-EXEC.
   01 D-COUNT PIC S9(9) COMP-5.
   01 D-DATA PIC X(1024).
   01 D-RETURNED-OCTET-LENGTH PIC S9(9) COMP.
EXEC SQL END DECLARE SECTION END-EXEC.
   EXEC SQL CONNECT TO testdb AS con1 USER testuser END-EXEC.
   EXEC SQL ALLOCATE DESCRIPTOR d END-EXEC.
   Declare, open a cursor, and assign a descriptor to the cursor
   EXEC SQL DECLARE cur CURSOR FOR SELECT current_database() END-EXEC.
   EXEC SQL OPEN cur END-EXEC.
   EXEC SQL FETCH NEXT FROM cur INTO SQL DESCRIPTOR d END-EXEC.
   Get a number of total columns
   EXEC SQL GET DESCRIPTOR d :D-COUNT = COUNT END-EXEC.
   DISPLAY "d_count
                                   = " D-COUNT.
  Get length of a returned column
   EXEC SQL GET DESCRIPTOR d VALUE 1 :D-RETURNED-OCTET-LENGTH = RETURNED_OCTET_LENGTH END-
EXEC.
   DISPLAY "d_returned_octet_length = " D-RETURNED-OCTET-LENGTH.
  Fetch the returned column as a string
   EXEC SQL GET DESCRIPTOR d VALUE 1 :D-DATA = DATA END-EXEC.
   DISPLAY "d_data
                                    = " D-DATA.
   Closing
   EXEC SQL CLOSE cur END-EXEC.
   EXEC SQL COMMIT END-EXEC.
   EXEC SQL DEALLOCATE DESCRIPTOR d END-EXEC.
   EXEC SQL DISCONNECT ALL END-EXEC.
```

When the example is executed, the result will look like this:

Compatibility

GET DESCRIPTOR is specified in the SQL standard.

See Also

ALLOCATE DESCRIPTOR, SET DESCRIPTOR

D.11.9 OPEN

Name

OPEN -- open a dynamic cursor

Synopsis

```
OPEN cursor_name

OPEN cursor_name USING value [, ... ]

OPEN cursor_name USING SQL DESCRIPTOR descriptor_name
```

Description

OPEN opens a cursor and optionally binds actual values to the placeholders in the cursor's declaration. The cursor must previously have been declared with the DECLARE command. The execution of OPEN causes the query to start executing on the server.

Parameters

cursor_name

The name of the cursor to be opened. This can be an SQL identifier or a host variable.

value

A value to be bound to a placeholder in the cursor. This can be an SQL constant, a host variable, or a host variable with indicator.

descriptor_name

The name of a descriptor containing values to be bound to the placeholders in the cursor. This can be an SQL identifier or a host variable.

Examples

```
EXEC SQL OPEN a END-EXEC.

EXEC SQL OPEN d USING 1, 'test' END-EXEC.

EXEC SQL OPEN c1 USING SQL DESCRIPTOR mydesc END-EXEC.

EXEC SQL OPEN :curname1 END-EXEC.
```

Compatibility

OPEN is specified in the SQL standard.

See Also

DECLARE, CLOSE



See

Refer to "SQL Commands" in "Reference" in the PostgreSQL Documentation for information on the CLOSE command.

D.11.10 PREPARE

Name

PREPARE -- prepare a statement for execution

Synopsis

```
PREPARE name FROM string
```

Description

PREPARE prepares a statement dynamically specified as a string for execution. This is different from the direct SQL statement PREPARE, which can also be used in embedded programs. The EXECUTE command is used to execute either kind of prepared statement.

Parameters

prepared_name

An identifier for the prepared query.

string

A literal string or a host variable containing a preparable statement, one of the SELECT, INSERT, UPDATE, or DELETE.

Examples

```
MOVE "SELECT * FROM test1 WHERE a = ? AND b = ?" TO ARR OF STMT.

COMPUTE LEN OF STMT = FUNCTION STORED-CHAR-LENGTH (ARR OF STMT).

EXEC SQL ALLOCATE DESCRIPTOR indesc END-EXEC.

EXEC SQL ALLOCATE DESCRIPTOR outdesc END-EXEC.

EXEC SQL PREPARE foo FROM :STMT END-EXEC.

EXEC SQL EXECUTE foo USING SQL DESCRIPTOR indesc INTO SQL DESCRIPTOR outdesc END-EXEC.
```

Compatibility

PREPARE is specified in the SQL standard.

See Also

EXECUTE



See

Refer to "SQL Commands" in "Reference" in the PostgreSQL Documentation for information on the EXECUTE command.

D.11.11 SET AUTOCOMMIT

Name

SET AUTOCOMMIT -- set the autocommit behavior of the current session

Synopsis

```
SET AUTOCOMMIT \{ = | TO \} \{ ON | OFF \}
```

Description

SET AUTOCOMMIT sets the autocommit behavior of the current database session. By default, embedded SQL programs are not in autocommit mode, so COMMIT needs to be issued explicitly when desired. This command can change the session to autocommit mode, where each individual statement is committed implicitly.

Compatibility

SET AUTOCOMMIT is an extension of PostgreSQL ECOBPG.

D.11.12 SET CONNECTION

Name

SET CONNECTION -- select a database connection

Synopsis

```
SET CONNECTION [ TO | = ] connection_name
```

Description

SET CONNECTION sets the "current" database connection, which is the one that all commands use unless overridden.

Parameters

connection_name

A database connection name established by the CONNECT command.

DEFAULT

Set the connection to the default connection.

Examples

```
EXEC SQL SET CONNECTION TO con2 END-EXEC.

EXEC SQL SET CONNECTION = con1 END-EXEC.
```

Compatibility

SET CONNECTION is specified in the SQL standard.

See Also

CONNECT, DISCONNECT

D.11.13 SET DESCRIPTOR

Name

SET DESCRIPTOR -- set information in an SQL descriptor area

Synopsis

```
SET DESCRIPTOR descriptor_name descriptor_header_item = value [, ...]

SET DESCRIPTOR descriptor_name VALUE number descriptor_item = value [, ...]
```

Description

SET DESCRIPTOR populates an SQL descriptor area with values. The descriptor area is then typically used to bind parameters in a prepared query execution.

This command has two forms: The first form applies to the descriptor "header", which is independent of a particular datum. The second form assigns values to particular datums, identified by number.

Parameters

descriptor_name

A descriptor name.

descriptor_header_item

A token identifying which header information item to set. Only COUNT, to set the number of descriptor items, is currently supported.

number

The number of the descriptor item to set. The count starts at 1.

descriptor_item

A token identifying which item of information to set in the descriptor. See Section 33.7.1 for a list of supported items.

value

A value to store into the descriptor item. This can be an SQL constant or a host variable.

Examples

```
EXEC SQL SET DESCRIPTOR indesc COUNT = 1 END-EXEC.

EXEC SQL SET DESCRIPTOR indesc VALUE 1 DATA = 2 END-EXEC.

EXEC SQL SET DESCRIPTOR indesc VALUE 1 DATA = :val1 END-EXEC.

EXEC SQL SET DESCRIPTOR indesc VALUE 2 DATA = 'some string', INDICATOR = :val1 END-EXEC.

EXEC SQL SET DESCRIPTOR indesc VALUE 2 INDICATOR = :val2 END-EXEC.
```

Compatibility

SET DESCRIPTOR is specified in the SQL standard.

See Also

ALLOCATE DESCRIPTOR, GET DESCRIPTOR

D.11.14 TYPE

Name

TYPE -- define a new data type

Synopsis

```
TYPE type_name IS ctype
```

Description

The TYPE command defines a new COBOL type. It is equivalent to putting a typedef into a declare section.

This command is only recognized when ecobpgpg is run with the -c option.

A level number of 01 is automatically added to type_name item. Thus, the level number must not to be specified externally. To define a group item, a level number needs to be specified to the each subordinate items.

For reasons of internal implementation, "TYPE" must be placed just after "EXEC SQL", without containing newline. For other place, you can use newline.

Parameters

type_name

The name for the new type. It must be a valid COBOL type name.

ctype

A COBOL type specification (including expression format specification).

Examples

```
EXEC SQL TYPE CUSTOMER IS

02 NAME PIC X(50) VARYING.

02 PHONE PIC S9(9) COMP. END-EXEC.

EXEC SQL TYPE CUST-IND IS

02 NAME_IND PIC S9(4) COMP.

02 PHONE_IND PIC S9(4) COMP. END-EXEC.

EXEC SQL TYPE INTARRAY IS

02 INT PIC S9(9) OCCURS 20. END-EXEC.

EXEC SQL TYPE STR IS PIC X(50) VARYING. END-EXEC.

EXEC SQL TYPE STRING IS PIC X(10). END-EXEC.
```

Here is an example program that uses EXEC SQL TYPE:

```
EXEC SQL TYPE TT IS
    02 V PIC X(256) VARYING.
    02 I PIC S9(9) COMP. END-EXEC.
EXEC SQL TYPE TT-IND IS
    02 V-IND PIC S9(4) COMP.
    02 I-IND PIC S9(4) COMP. END-EXEC.
EXEC SQL BEGIN DECLARE SECTION END-EXEC.
   01 T TYPE TT.
   01 T-IND TYPE TT-IND.
EXEC SQL END DECLARE SECTION END-EXEC.
   EXEC SQL CONNECT TO testdb AS con1 END-EXEC.
   EXEC SQL SELECT current_database(), 256 INTO :T :T-IND LIMIT 1 END-EXEC.
   DISPLAY "t.v = " ARR OF V OF T.
   DISPLAY "t.i = " I OF T.
   DISPLAY "t_ind.v_ind = " V-IND OF T-IND.
   DISPLAY "t_ind.i_ind = " I-IND OF T-IND.
   EXEC SQL DISCONNECT con1 END-EXEC.
```

Compatibility

The TYPE command is a PostgreSQL extension.

D.11.15 VAR

Name

VAR-define a variable

Synopsis

VAR varname IS ctype

Description

The VAR command defines a host variable. It is equivalent to an ordinary COBOL variable definition inside a declare section.

When translating, a level number 01 is added. Thus, the level number must not to be specified externally.

To define a group item, a level number needs to be specified to the each subordinate items.

For reasons of internal implementation, "VAR" must be placed just after "EXEC SQL", without containing newline. For other place, you can use newline.

Parameters

varname

A COBOL variable name.

ctype

A COBOL type specification.

Examples

```
EXEC SQL VAR VC IS PIC X(10) VARYING. END-EXEC.
EXEC SQL VAR BOOL-VAR IS BOOL. END-EXEC.
```

Compatibility

The VAR command is a PostgreSQL extension.

D.11.16 WHENEVER

Name

WHENEVER -- specify the action to be taken when an SQL statement causes a specific class condition to be raised

Synopsis

```
WHENEVER { NOT FOUND | SQLERROR | SQLWARNING } action
```

Description

Define a behavior which is called on the special cases (Rows not found, SQL warnings or errors) in the result of SQL execution.

Parameters

See Section "D.7.1 Setting Callbacks" or a description of the parameters.

Examples

```
EXEC SQL WHENEVER NOT FOUND CONTINUE END-EXEC.

EXEC SQL WHENEVER SQLWARNING SQLPRINT END-EXEC.

EXEC SQL WHENEVER SQLERROR sqlprint END-EXEC.

EXEC SQL WHENEVER SQLERROR Sqlprint END-EXEC.

EXEC SQL WHENEVER SQLERROR CALL "print2" END-EXEC.

EXEC SQL WHENEVER SQLERROR DO handle_error USING "select" END-EXEC.

EXEC SQL WHENEVER SQLERROR DO sqlnotice USING 0 1 END-EXEC.

EXEC SQL WHENEVER SQLERROR DO "sqlprint" END-EXEC.

EXEC SQL WHENEVER SQLERROR GOTO error_label END-EXEC.

EXEC SQL WHENEVER SQLERROR STOP END-EXEC.
```

A typical application is the use of WHENEVER NOT FOUND GOTO to handle looping through result sets:

```
EXEC SQL CONNECT TO testdb AS con1 END-EXEC.

EXEC SQL ALLOCATE DESCRIPTOR d END-EXEC.

EXEC SQL DECLARE cur CURSOR FOR SELECT current_database(), 'hoge', 256 END-EXEC.

EXEC SQL OPEN cur END-EXEC.

* when end of result set reached, break out of while loop

EXEC SQL WHENEVER NOT FOUND GOTO NOTFOUND END-EXEC.

PERFORM NO LIMIT

EXEC SQL FETCH NEXT FROM cur INTO SQL DESCRIPTOR d END-EXEC

...

END-PERFORM.

NOTFOUND.

EXEC SQL CLOSE cur END-EXEC.

EXEC SQL COMMIT END-EXEC.

EXEC SQL DEALLOCATE DESCRIPTOR d END-EXEC.

EXEC SQL DEALLOCATE DESCRIPTOR d END-EXEC.

EXEC SQL DISCONNECT ALL END-EXEC.
```

Compatibility

WHENEVER is specified in the SQL standard, but most of the actions are PostgreSQL extensions.

D.12 PostgreSQL Client Applications

This part contains reference information for PostgreSQL client applications and utilities. Not all of these commands are of general utility; some might require special privileges. The common feature of these applications is that they can be run on any host, independent of where the database server resides.

When specified on the command line, user and database names have their case preserved—the presence of spaces or special characters might require quoting. Table names and other identifiers do not have their case preserved, except where documented, and might require quoting.

D.12.1 ecobpg

Name

```
ecobpg -- embedded SQL COBOL preprocessor
```

Synopsis

ecobpg [option...] file...

Description

ecobpg is the embedded SQL preprocessor for COBOL programs. It converts COBOL programs with embedded SQL statements to normal COBOL code by replacing the SQL invocations with special function calls. The output files can then be processed with any COBOL compiler tool chain.

ecobpg will convert each input file given on the command line to the corresponding COBOL output file. Input files preferably have the extension .pco, in which case the extension will be replaced by .cob to determine the output file name. If the extension of the input file is not .pco, then the output file name is computed by appending .cob to the full file name. The output file name can also be overridden using the -o option.

Options

ecobpg accepts the following command-line arguments:

-c

Automatically generate certain COBOL code from SQL code. Currently, this works for EXEC SQL TYPE.

-I directory

Specify an additional include path, used to find files included via EXEC SQL INCLUDE. Defaults are: (current directory), /usr/local/include, the PostgreSQL include directory which is defined at compile time (default: /usr/local/pgsql/include), and /usr/include, in that order.

-o filename

Specifies that ecobpg should write all its output to the given filename.

-f format

Specifies the COBOL code notation. For "format", specify either of the following. If omitted, "fixed" is used.

fixed

Specifies fixed format notation. Up to 72 columns can be specified for area B. Characters in column 73 and beyond are deleted in the precompiled source.

variable

Specifies variable format notation. Up to 251 columns can be specified for area B. Characters in column 252 and beyond are deleted in the precompiled source.

-r option

Selects run-time behavior. Option can be one of the following:

prepare

Prepare all statements before using them. Libecpg will keep a cache of prepared statements and reuse a statement if it gets executed again. If the cache runs full, libecpg will free the least used statement.

questionmarks

Allow question mark as placeholder for compatibility reasons. This used to be the default long ago.

-t

Turn on autocommit of transactions. In this mode, each SQL command is automatically committed unless it is inside an explicit transaction block. In the default mode, commands are committed only when EXEC SQL COMMIT is issued.

--varchar-with-named-member

When converting VARCHAR host variable, adding name of the variable to members as prefix. Instead of LEN and ARR, (varname)-ARR and (varname)-LEN will be used.

--bytea-with-named-member

When converting bytea host variable, adding name of the variable to members as prefix. Instead of LEN and ARR, (varname)-ARR and (varname)-LEN will be used.

-E encode

Specify the COBOL source encoding: "UTF8", "SJIS", or "EUC_JP".

If this option is omitted, the encoding is processed based on the locale.

-V

Print additional information including the version and the "include" path.

--version

Print the ecobpg version and exit.

-?

--help

Show help about ecobpg command line arguments, and exit.

Notes

When compiling the preprocessed COBOL code files, the compiler needs to be able to find the library files in the PostgreSQL include directory.

Programs using COBOL code with embedded SQL have to be linked against the libecpg library, for example using the linker options.

The value of either of these directories that is appropriate for the installation can be found out using pg_config.



Refer to "pg_config" in "Reference" in the PostgreSQL Documentation.

Examples

If you have an embedded SQL COBOL source file named prog1.pco, you can create an executable program using the following command:

ecobpg progl.pco

Appendix E Quantitative Limits

This appendix lists the quantitative limits of Fujitsu Enterprise Postgres.

Table E.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)
Table space name	Up to 63 bytes (*1) (*2)
Cursor name	Up to 63 bytes (*1) (*2)
Function name	Up to 63 bytes (*1) (*2)
Aggregate function name	Up to 63 bytes (*1) (*2)
Trigger name	Up to 63 bytes (*1) (*2)
Constraint name	Up to 63 bytes (*1) (*2)
Conversion name	Up to 63 bytes (*1) (*2)
Role name	Up to 63 bytes (*1) (*2)
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)

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

Table E.2 Database object

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

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

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

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

Table E.5 Data types and attributes that can be hand 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	2 bytes	From -32,768 to 32,767
	expression	4 bytes	From -2,147,483,648 to 2,147,483,647
	8 bytes	From -9,223,372,036,854,775,808 to 9,223,372,036,854,775,807	
	Internal decimal expression		Up to 13,1072 digits before the decimal point, and up to 16,383 digits after the decimal point
	Floating point expression	4 bytes	From -3.4E+38 to -7.1E-46, 0, or from 7.1E-46 to 3.4E+38
	8 bytes	From -1.7E+308 to -2.5E-324, 0, or from 2.5E-324 to 1.7E+308	
bytea		Up to one gigabyte minus 53 bytes	
Large object		Up to two gigabytes	

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

Table E.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

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

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

Table E.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

Item	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.

Table E.8 Data sizes

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.

^{*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.

Appendix F Reference

F.1 JDBC Driver

Refer to the Java API Reference for information on PostgreSQL JDBC driver.

F.2 ODBC Driver

F.2.1 List of Supported APIs

The following table shows the support status of APIs:

Function name	Support status
SQLAllocConnect	Y
SQLAllocEnv	Y
SQLAllocHandle	Y
SQLAllocStmt	Y
SQLBindCol	Y
SQLBindParameter	Y
SQLBindParam	Y
SQLBrowseConnect	Y
SQLBulkOperations	Y
SQLCancel	Y
SQLCancelHandle	N
SQLCloseCursor	Y
SQLColAttribute	Y
SQLColAttributeW	Y
SQLColAttributes	Y
SQLColAttributesW	Y
SQLColumnPrivileges	Y
SQLColumnPrivilegesW	Y
SQLColumns	Y
SQLColumnsW	Y
SQLCompleteAsync	N
SQLConnect	Y
SQLConnectW	Y
SQLCopyDesc	Y
SQLDataSources	Y
SQLDataSourcesW	Y
SQLDescribeCol	Y

Function name	Support status
SQLDescribeColW	Y
SQLDescribeParam	Y
SQLDisconnect	Y
SQLDriverConnect	Y
SQLDriverConnectW	Y
SQLDrivers	Y
SQLEndTran	Y
SQLError	Y
SQLErrorW	Y
SQLExecDirect	Y
SQLExecDirectW	Y
SQLExecute	Y
SQLExtendedFetch	Y
SQLFetch	Y
SQLFetchScroll	Y
SQLForeignKeys	Y
SQLForeignKeysW	Y
SQLFreeConnect	Y
SQLFreeEnv	Y
SQLFreeHandle	Y
SQLFreeStmt	Y
SQLGetConnectAttr	Y
SQLGetConnectAttrW	Y
SQLGetConnectOption	Y
SQLGetConnectOptionW	Y
SQLGetCursorName	Y
SQLGetCursorNameW	Y
SQLGetData	Y
SQLGetDescField	Y
SQLGetDescFieldW	Y
SQLGetDescRec	Y
SQLGetDescRecW	Y
SQLGetDiagField	Y
SQLGetDiagFieldW	Y
SQLGetDiagRec	Y
SQLGetDiagRecW	Y
SQLGetEnvAttr	Y
SQLGetFunctions	Y
SQLGetInfo	Y

Function name	Support status
SQLGetInfoW	Y
SQLGetStmtAttr	Y
SQLGetStmtAttrW	Y
SQLGetStmtOption	Y
SQLGetTypeInfo	Y
SQLGetTypeInfoW	Y
SQLMoreResults	Y
SQLNativeSql	Y
SQLNativeSqlW	Y
SQLNumParams	Y
SQLNumResultCols	Y
SQLParamData	Y
SQLParamOptions	Y
SQLPrepare	Y
SQLPrepareW	Y
SQLPrimaryKeys	Y
SQLPrimaryKeysW	Y
SQLProcedureColumns	Y
SQLProcedureColumnsW	Y
SQLProcedures	Y
SQLProceduresW	Y
SQLPutData	Y
SQLRowCount	Y
SQLSetConnectAttr	Y
SQLSetConnectAttrW	Y
SQLSetConnectOption	Y
SQLSetConnectOptionW	Y
SQLSetCursorName	Y
SQLSetCursorNameW	Y
SQLSetDescField	Y
SQLSetDescRec	Y
SQLSetEnvAttr	Y
SQLSetParam	Y
SQLSetPos	Y
SQLSetScrollOptions	N
SQLSetStmtAttr	Y
SQLSetStmtAttrW	Y
SQLSetStmtOption	Y
SQLSpecialColumns	Y

Function name	Support status
SQLSpecialColumnsW	Y
SQLStatistics	Y
SQLStatisticsW	Y
SQLTablePrivileges	Y
SQLTablePrivilegesW	Y
SQLTables	Y
SQLTablesW	Y
SQLTransact	Y

Y: Supported

N: Not supported

F.3 .NET Data Provider

There are the following ways to develop applications using Fujitsu Npgsql .NET Data Provider:

- Use the Fujitsu Npgsql.NET Data Provider API (classes and methods) directly.

Fujitsu Npgsql .NET Data Provider is created based on the open source software Npgsql. Refer to the "Npgsql - .Net Data Provider for Postgresql" for information on the APIs:



See

Refer to the Installation and Setup Guide for Client for the version of Npgsql that Fujitsu Npgsql .NET Data Provider is based on.

- Use the API of the .NET System.Data.Common namespace

It is possible to create applications that do not rely on a provider when you use the System.Data.Common namespace. Refer to the "Writing Provider-Independent Code in ADO.NET" in MSDN Library for information.



See

Refer to the Npgsql API Reference for information on Npgsql API classes and methods.

F.4 C Library (libpq)



See

Refer to "libpq - C Library" in "Client Interfaces" in the PostgreSQL Documentation.

F.5 Embedded SQL in C



See

 $Refer \ to \ "ECPG - Embedded \ SQL \ in \ C" \ in \ "Client \ Interfaces" \ in \ the \ PostgreSQL \ Documentation.$

Index

[A] Additional Notes on each Type Plugin43
[B] BIND_VARIABLE89
[C]
CLOSE_CURSOR89
Code examples for applications56,66
COLUMN_VALUE
[D] DBMS_SQL87
DECODE
DEFINE COLUMN
DUAL Table82
[E]
Encoding System Settings24,49
Entry information of subprogram69
Example of specifying the hint clause57,68
EXECUTE90
[F]
FETCH_ROWS91
[L]
Language settings
Libraries to use
[N]
Notes on automatically generating update-type SQL statements
Notes on metadata
Notes on the Query Builder
1172
[O]
OPEN_CURSOR 91 Outer Join Operator (+) 79
Outer Join Operator (1)
[P]
PARSE
Path of library 69 Path of the library file 69
Pattern matching
Precompiling example
[S]
Scan Using a Vertical Clustered Index (VCI)
Settings for encrypting communication data for connection to the
server
String functions and operators
SUBSTR84
[T]
Type Plugins. 42

[U]	
Use the API of the .NET System.Data.Common namespace	e. <mark>20</mark> 3
Use the Fujitsu Npgsql.NET Data Provider API (classes ar	ıd
methods) directly	203
[W]	
When setting from outside with environment variables	50
When specifying in the connection URI	50
When using Add-OdbcDsn	20
When using ODBCConf.exe	18