

The Fujitsu Supported PostgreSQL (FSP) Database Solution

*Fujitsu's Specialist Services and Support for the Open Source PostgreSQL
Database*

Solution Overview Whitepaper

Version 1.02

www.fastware.com/postgresql.html

Fujitsu Australia Software Technology Pty Ltd



THE POSSIBILITIES ARE INFINITE

This document contains technology relating to strategic products controlled by export control laws of the producing and/ or exporting countries. This document or a portion thereof should not be exported (or re-exported) without authorization from the appropriate government authorities in accordance with such laws.

No part of this publication may be reproduced or transmitted in any form or by any mean, electronic or mechanical, including photocopying, recording, or storing in a retrieval system, or translated into any language in any form without the express written permission of Fujitsu Australia Software Technology Pty Ltd.

Fujitsu Australia Software Technology Pty Ltd endeavours to ensure that the information in this document is correct and fairly stated, but does not accept liability for any error or omission.

The development of Fujitsu Supported PostgreSQL is continuous and published information may not be up to date. It is important to check the current position with Fujitsu Australia Software Technology Pty Ltd. This document is not part of a contract or license insofar as may be expressly agreed.

Fujitsu Australia Software Technology Pty Ltd

Fujitsu Supported PostgreSQL

Solution Overview Whitepaper Version 1.01 (January 2005)

The contents of this whitepaper may be revised without prior notice.

All Rights Reserved

Copyright © Fujitsu Australia Software Technology PTY LTD 2005

Table of Contents

1	ABSTRACT	5
2	INTRODUCTION TO THE FUJITSU SUPPORTED POSTGRESQL (FSP) DATABASE SOLUTION	5
2.1	The Open Source PostgreSQL Database in Brief	6
2.2	Fujitsu’s PostgreSQL Value Add – Support, Training, Consultancy, System Integration, Migration, and Hosting	6
2.3	Fujitsu’s Extended Storage Management (ESM) Plug-in Module for PostgreSQL	7
3	THE FSP DATABASE SOLUTION WITHIN THE ENTERPRISE	8
3.1	Support	8
3.2	Training	8
3.3	Hosting Services	8
3.4	Consulting, Migration Services and Systems Integration	8
3.5	Ongoing PostgreSQL Development	8
4	THE FSP DATABASE SOLUTION WITHIN FUJITSU’S INFORMATION ARCHITECTURE	9
4.1	Reasons to Choose the FSP Database Solution	9
4.2	Enterprise Ready	10
4.3	Independent Software Vendor (ISV) Support	10
5	THE ESM PLUG-IN MODULE FOR POSTGRESQL – EXTENDED RELIABILITY, PERFORMANCE AND FUNCTIONALITY	11
5.1	Efficient Storage Equals Better Performance	11
5.1.1	Scalable Solution	11
5.1.2	Superior Reliability - “Stealth Sync” Enables Stable Response Times	12
5.1.3	Pipeline Control of ESM Processes	13
5.1.4	Multi Version Concurrency Control (MVCC)	13
5.2	Ease of Operation	14
5.2.1	All-in-one Backup – No Vacuuming Required	14
5.2.2	Media Recovery	14
5.3	Other Features of the ESM Plug-in Module	15
5.3.1	Cross Platform Deployment	15
5.3.2	Application Server Support	15
5.3.3	Ultra-fast Initial Database Loading	15
5.3.4	Easy Upgrade from PostgreSQL	15
5.3.5	Tools	15
6	THE OPEN SOURCE POSTGRESQL ADVANTAGE	17
6.1	Reliability	18
6.1.1	ACID Compliance	18
6.1.2	Multi-version Concurrency Control	18
6.1.3	Write Ahead Logging (WAL)	18
6.1.4	Point in Time Recovery (PITR)	18

6.1.5	Replication	18
6.1.6	Data Integrity	19
6.1.7	The Open Source PostgreSQL Development Model	19
6.1.8	Open Source	19
6.2	Performance	19
6.2.1	Indexing Support	19
6.2.2	The PostgreSQL Query Planner	20
6.2.3	Granular Locking	20
6.2.4	Buffer Management and Caching	20
6.2.5	Scalability	20
6.3	Extensibility	20
6.4	SQL Support	21
6.5	Rich Data Type System	21
6.5.1	Character Types	21
6.5.2	Numeric Types	21
6.5.3	Arrays	21
6.5.4	Large Objects	21
6.5.5	Geometric Types	22
6.5.6	Geographic Information System (GIS) types	22
6.5.7	Network Types	22
6.5.8	Composite Types	22
6.5.9	High Precision Dates and Time	22
6.6	Extensive Character Encoding Support	22
6.7	Native Language Support	22
6.8	Interfaces	22
6.9	Procedural Languages	22
6.10	Ease of Use	23
6.11	Data Security	23
7	FUJITSU'S INTERSTAGE AND FSP	24
7.1	The Interstage Suite	24
7.2	Interstage Application Server Foundation	24
7.3	Interstage Integration Products	24
8	FUJITSU AND OPEN SOURCE SOFTWARE	25
8.1	Our Commitment	25
8.2	Working with Partners	25
8.3	Support of FSP	25
9	ABOUT FUJITSU AUSTRALIA SOFTWARE TECHNOLOGY (FAST)	26
9.1	Contact Us	26
9.2	On the Web	26
10	APPENDIX I – GLOSSARY OF TERMS AND ABBREVIATIONS	27

1 Abstract

The Fujitsu Supported PostgreSQL (FSP) Database Solution featuring the Extended Storage Management (ESM) module is the latest database initiative from the Fujitsu Software Group. As part of FSP, Fujitsu supports the Open Source community's PostgreSQL database, and by enhancing the storage mechanisms of PostgreSQL using specialist technology from Fujitsu, FSP with ESM offers the best of Open Source combined with the best of Fujitsu. This powerful combination is backed by Fujitsu's world-class support services, and delivers to customers a first-rate, dependable database solution, at a very affordable price.

2 Introduction to the Fujitsu Supported PostgreSQL (FSP) Database Solution

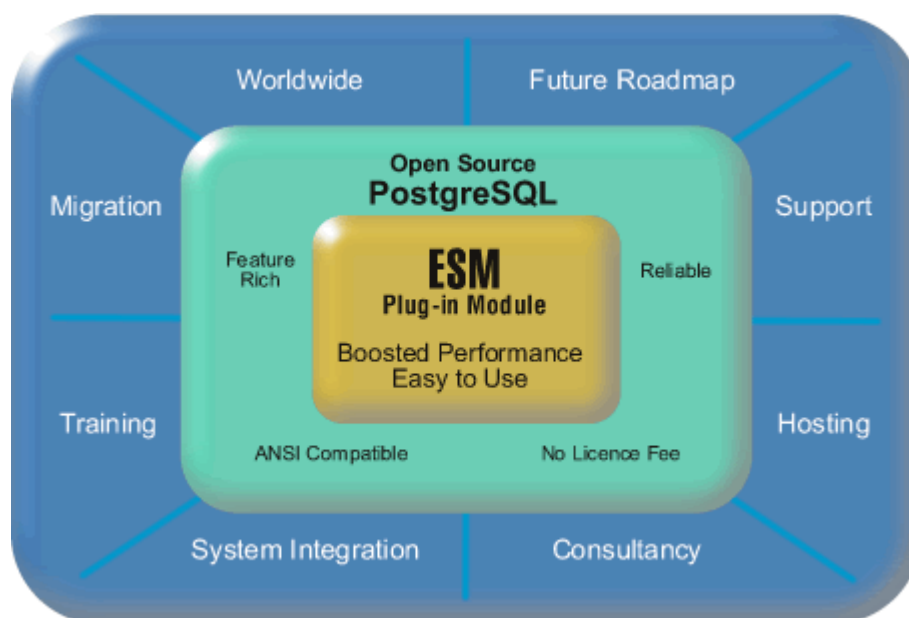


Figure 1 – FSP Conceptual View

Fujitsu, as part of its commitment to Open Source Software initiatives, is now offering comprehensive support arrangements to organisations wishing to capitalise on the established power and reliability of the Open Source PostgreSQL database. In addition to a wide range of specialist services, Fujitsu Supported PostgreSQL also features the Extended Storage Management (ESM) plug-in module, placing PostgreSQL firmly on the map as the ideal low-cost database solution for businesses of all sizes.

2.1 The Open Source PostgreSQL Database in Brief

PostgreSQL is the world's most advanced open source database. It is developed and distributed by a voluntary team of active participants in the Open Source PostgreSQL community. It is widely used at all levels of business and is renowned for its reliability, stability, high performance, ease of use, rich feature set and SQL compliance. This alone makes it a serious player in the database market, but combined with its transparent open source development process, PostgreSQL ensures that users can control the future of their database and avoid lock-in to expensive middleware, tools and upgrade paths.

In addition to its well-proven, wide-ranging capabilities, Open Source PostgreSQL is governed by the BSD (Berkeley Software Design) licence which delivers the following key benefits:

- There are no licence fees for the software.
- No fees accrue for on-sold products which incorporate an instance of PostgreSQL. This commercially favourable licensing model may not apply to other open source database products.

A more detailed discussion of PostgreSQL follows later in this document in the section entitled 'The Open Source PostgreSQL Advantage'.

Further information on Open Source PostgreSQL is available from the community website:

<http://www.postgresql.org/>

2.2 Fujitsu's PostgreSQL Value Add – Support, Training, Consultancy, System Integration, Migration, and Hosting

Fujitsu Supported PostgreSQL (FSP) complements the Open Source PostgreSQL database system by providing a comprehensive range of products and services which address the complete needs of businesses deploying PostgreSQL.

New and established users of PostgreSQL can decisively benefit from FSP packages in the following ways:

- Reliance on Fujitsu's established support reputation
- Assurance of dealing with well-trained and experienced support staff
- Fast turnaround times on support calls
- Rapid escalation of high-impact issues
- Additional training, consultancy, migration and hosting services as required

Fujitsu's strength and experience as one of the world's leading IT vendors will allow organisations to focus their expertise on core business, comfortable in the knowledge that their PostgreSQL database has all the support it needs.

2.3 Fujitsu's Extended Storage Management (ESM) Plug-in Module for PostgreSQL




Figure 2 – The ESM Plug-in Module for Open Source PostgreSQL

Note: The developer-level package name for the ESM module is PostgreSQL Plus.

The ESM module is built upon the Open Source community's PostgreSQL database. It is enhanced with the storage mechanisms of Fujitsu's own database technologies. In combining this specialised storage capability with the features and functionality of PostgreSQL, users who require high-end database availability, enterprise performance and scalability on anything from two to thirty two CPUs, now have an alternative that can bring substantial cost benefits to their operations.

A more detailed description of the ESM module is provided the section entitled 'The ESM Plug-in Module for PostgreSQL – Extended Reliability, Performance and Functionality'.

3 The FSP Database Solution within the Enterprise

As part of the FSP offering, Fujitsu can provide technical support, training, migration services, hosting services, systems integration and consulting services to businesses deploying PostgreSQL. In this way, FSP customers benefit not only from the reliability and transparency of open source, but also the confidence, reputation, and business and industry know-how that comes with an experienced vendor such as Fujitsu.

3.1 Support

Businesses which adopt open source software solutions can only do so if there are reliable vendor support options available for these solutions. Fujitsu Supported PostgreSQL offers a suite of support options which meets the needs of businesses wishing to use PostgreSQL in any environment.

3.2 Training

FSP also makes available training services to new and experienced PostgreSQL users alike. Fujitsu's extensive experience in software development and migration allows us to tailor our training to your requirements and answer your questions based on experience.

3.3 Hosting Services

In conjunction with Fujitsu's existing data centre and data management facilities, an FSP package can provide PostgreSQL users with the option of having their database managed in a world-class Internet Data Centre with 24 hour monitoring, redundancy at all levels, and a variety of network access points.

3.4 Consulting, Migration Services and Systems Integration

In order to bring the benefits of PostgreSQL to your existing applications a number of tasks may need to be completed. Fujitsu has experience and industry knowledge to move applications from other databases in the market to PostgreSQL.

Fujitsu also has extensive knowledge and access to world-class integration tools, including its own Interstage Application Server (discussed later in this document).

3.5 Ongoing PostgreSQL Development

Fujitsu is also committed to making PostgreSQL the most advanced database system it can be. During the 8.0 development cycle, it sponsored key developers to see the introduction of tablespaces and savepoints -- functionality which is commonly required by customers.

To continue this commitment, Fujitsu now employs a number of PostgreSQL development experts and consultants from the Open Source Community. As well as assisting clients, these developers will be contributing to the future of PostgreSQL by identifying features needed by users and implementing them.

4 The FSP Database Solution within Fujitsu's Information Architecture

When deployed as part of a complete solution, FSP and ESM can be used with other Fujitsu components such as the Interstage suite of products, or integrated with software from other vendors, or other open-source systems.

More information on Fujitsu's Interstage Suite can be found later in this document.

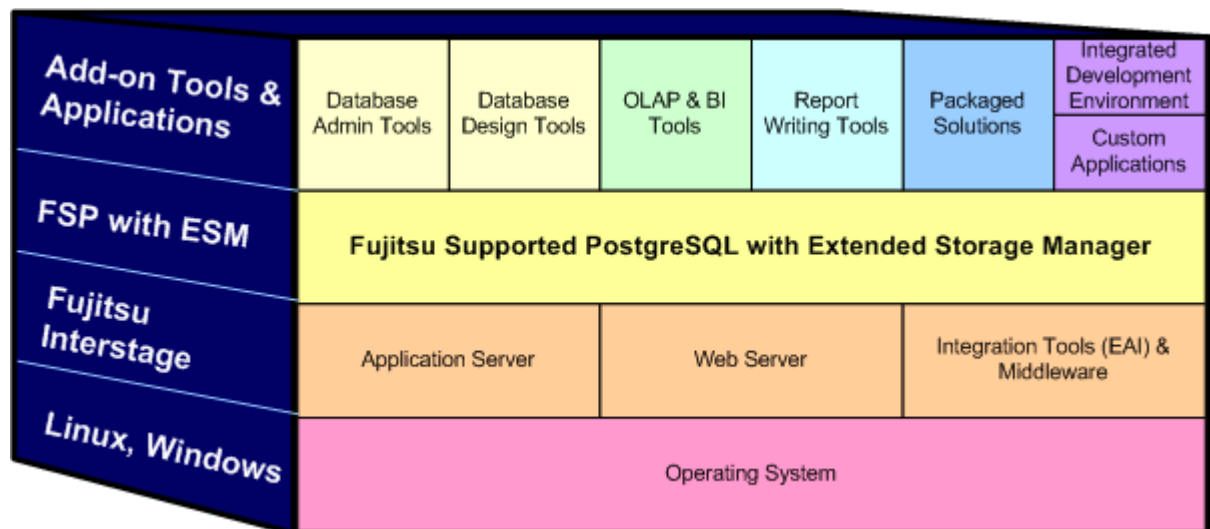


Figure 3 – Fujitsu Information Architecture Featuring FSP and ESM Technology

4.1 Reasons to Choose the FSP Database Solution

When you choose to migrate from an expensive commercial database to a free or low-cost database system, there are many factors that need to be considered:

- The FSP database solution is highly suitable for migrating from expensive proprietary databases because it supports a rich SQL language that is compatible with the ANSI SQL standards.
- The FSP database solution provides high-end Database scalability to users. Upgrading from the original PostgreSQL to incorporate the ESM plug-in module is a trouble-free process. PostgreSQL with ESM supports all the original PostgreSQL SQL and API commands.
- The FSP database solution offers a mix of features comparable to those of high-cost commercial databases. This includes triggers, functions, transaction processing, and different transaction isolation levels.
- Lower cost: Switching to the FSP Database solution will add up to a significant saving in licensing costs, over commercial products. In addition, PostgreSQL has been designed and created to have much lower maintenance and tuning requirements than more expensive proprietary databases.
- The FSP database solution will benefit from the future direction and growth of the Open Source PostgreSQL database. Additionally, Fujitsu will implement highly-specialised enhancements to the ESM module ensuring that FSP will always feature the latest innovations in Database Technology giving users the best possible future upgrade options.

4.2 Enterprise Ready

PostgreSQL is widely considered to be the most advanced of the open source databases. It is scalable, robust and secure, and its first-rate technical pedigree has led to its adoption as the database of choice for thousands of small, medium and large organisations worldwide.

PostgreSQL is easy to operationally administer and, importantly for development teams, is readily usable with widely available software engineering toolsets. This further benefits many organisations which make the switch to PostgreSQL as part of a strategy to reduce development costs.

Fujitsu's intentions are to make the FSP database solution into a product that scales to the levels required of very large enterprise systems. Some of the enhancements incorporated into the ESM plug-in module for PostgreSQL stem from Fujitsu's own Symfoware database product, which was first released in 1995.

4.3 Independent Software Vendor (ISV) Support

Fujitsu has a strategy of working with ISVs to extend the applicability of the FSP database solution. Rich ISV support is envisaged in order to integrate the FSP solution with the widest possible range of business application systems. Migration from databases by other vendors will be made easier with a set of migration and conversion tools.

5 The ESM Plug-in Module for PostgreSQL – Extended Reliability, Performance and Functionality

Seamless upgradeability to an enterprise-ready database makes the FSP database solution cost and feature competitive with expensive proprietary databases and the natural choice for PostgreSQL users who need high-end scalability.

5.1 Efficient Storage Equals Better Performance

The introduction of Fujitsu's complementary ESM (Extended Storage Manager) plug-in module into Open Source PostgreSQL gives the FSP solution added capabilities to meet the database needs of organisations requiring specialised high-end Enterprise features.

The ESM module enhances the original PostgreSQL storage mechanism component and extends it to include a sophisticated parallel processing sub-system.

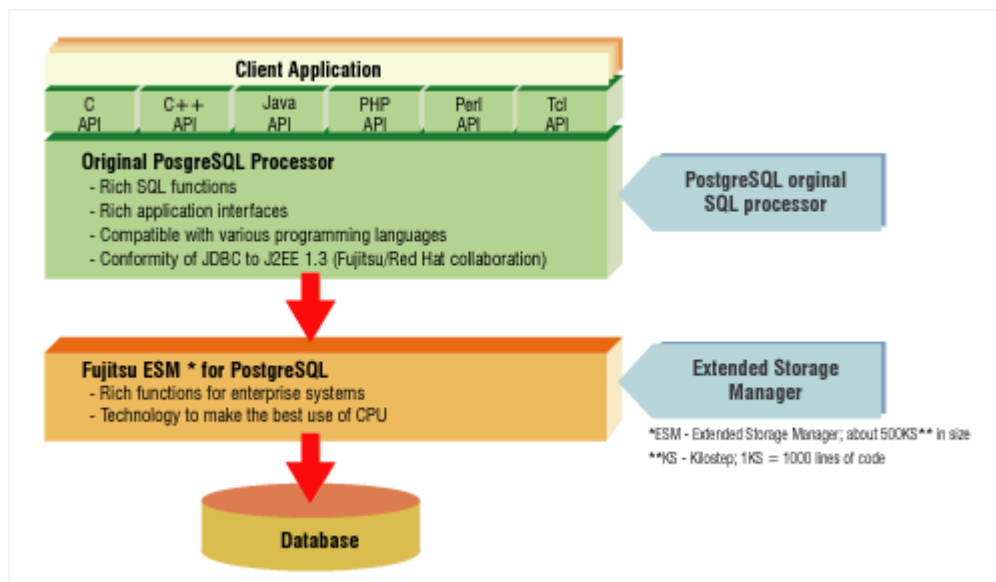


Figure 4 – Fujitsu's ESM technology enhances PostgreSQL

5.1.1 Scalable Solution

The Fujitsu storage mechanism transforms ESM-enhanced systems into highly scalable database solutions.

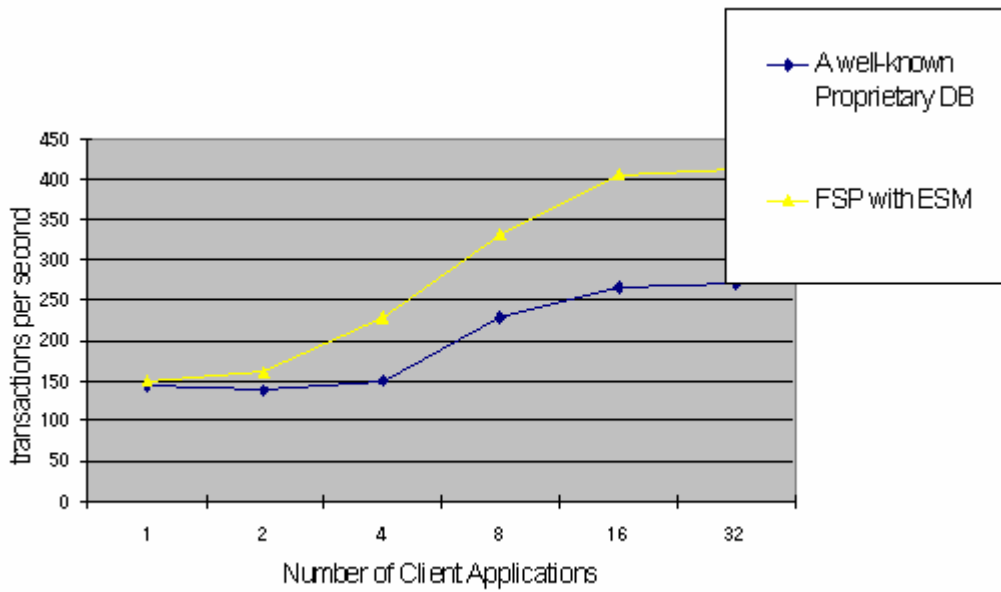


Figure 5 – Benchmark tests have shown that PostgreSQL enhanced with ESM can have greater scalability than other well-known proprietary databases.

Note: Tests were undertaken using the pgbench utility provided in the PostgreSQL source distribution. This test is an implementation, by the PostgreSQL community, of TPC-B[1].

5.1.2 Superior Reliability - "Stealth Sync" Enables Stable Response Times

ESM delivers stable transaction response times, even when your database load increases significantly.

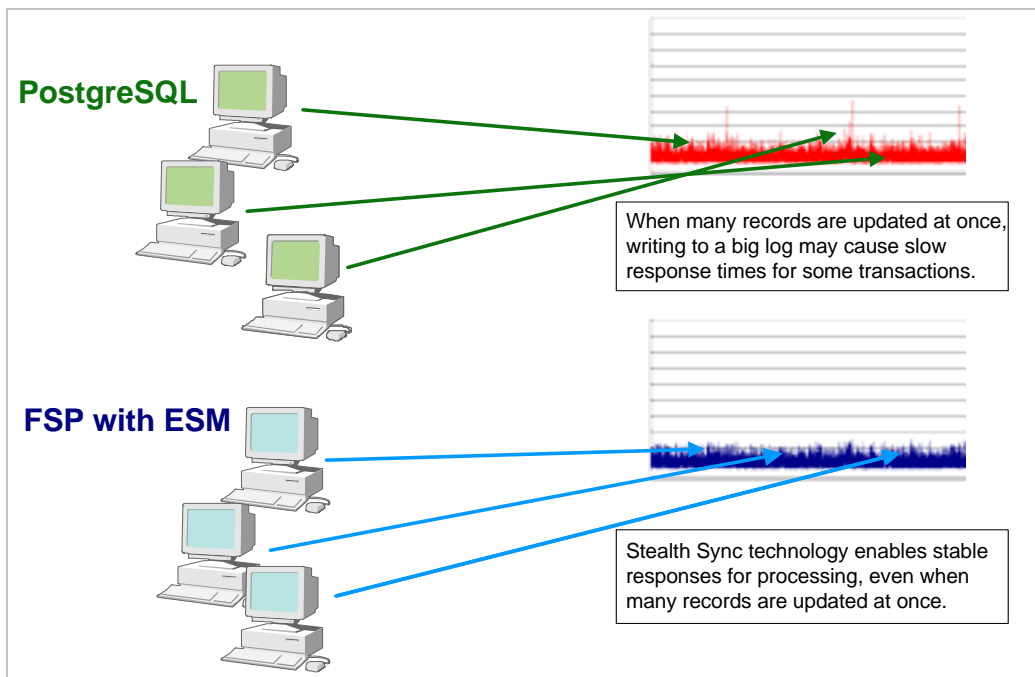


Figure 6 – Fujitsu’s Stealth Sync technology enables stable responses

5.1.3 Pipeline Control of ESM Processes

ESM delivers enhanced performance by parallel execution of several database processes.

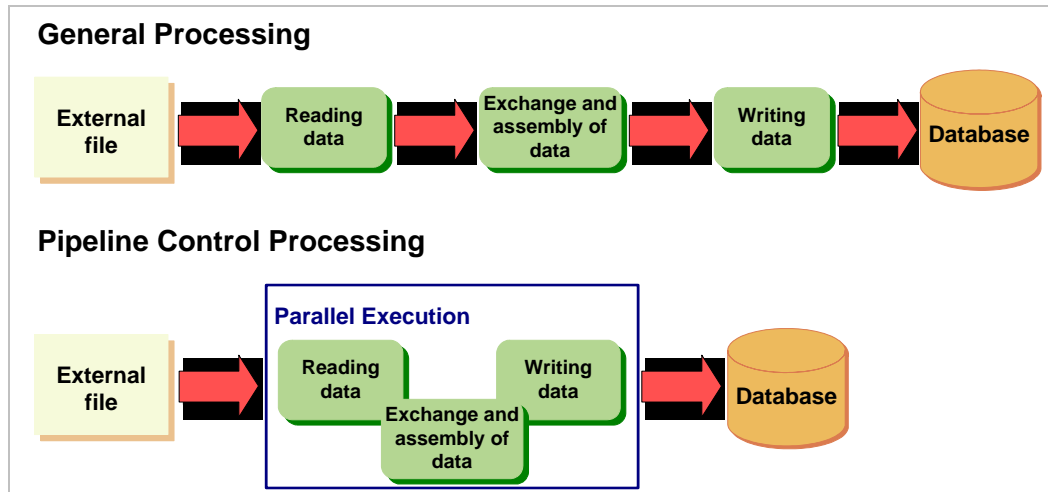


Figure 7 – Fujitsu’s Pipeline Technology Enables Parallel Processing

5.1.4 Multi Version Concurrency Control (MVCC)

The ESM module takes advantage of PostgreSQL's MVCC functionality to guarantee the consistency of reads (database accesses) and improve the efficiency of executing multiple concurrent transactions.

In MVCC, there is no lock on data that is being read, or on information that is not committed. (Instead, the data version itself is managed.) This means that there is no need for a lock or for blocking of any operation when information is accessed for read. All business processing runs simultaneously; the information seen by database users is always consistent, and the processing time of database operations is greatly reduced.

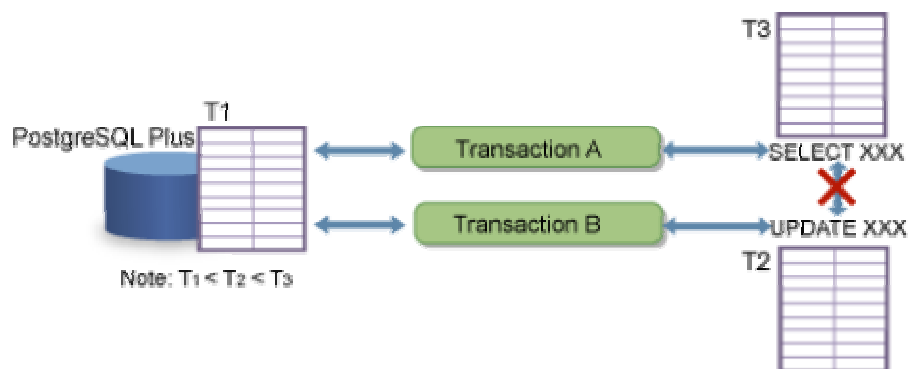


Figure 8 – Guarantee of Read Consistency

However, one scenario tackled successfully by ESM is the case when the number of transactions in MVCC reaches a high level with the potential to adversely affect performance. ESM solves this using the previously mentioned Stealth Sync control capability to write data to disk automatically and with optimum timing efficiency. In this way, ESM ensures a consistent response time even if the number of transactions increases. Superior scalability is realized for large numbers of transactions; the processing time is greatly reduced. MVCC can be used with a wide range of high-throughput applications.

5.2 Ease of Operation

5.2.1 All-in-one Backup – No Vacuuming Required

The ESM module takes care of managing a large-scale database without the need for periodic manually-activated Vacuuming operations. Automatic re-use of deleted spaces provides stable database performance without administration hassles.

5.2.2 Media Recovery

ESM offers ultra-fast online backup, with no interruptions to your business applications while you backup your valuable data. This is achieved by doing backups at file level, rather than record level.

In case of a disk crash, all committed transactions can be recovered from the backups and the transaction log.

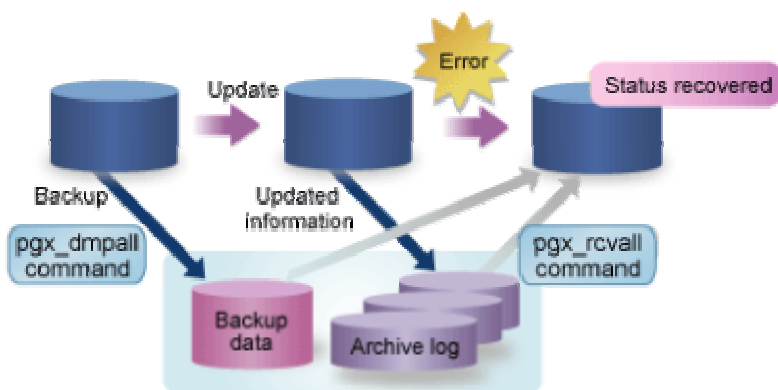


Figure 9 – Media Recovery

More Performance Improvements due in Future Versions

Fujitsu's is not stopping at the enhancements mentioned above. More are planned for future releases.

- Hint statements
- Statistical tools
- Stored procedures ('in', 'out', 'inout' parameters)
- Java stored procedures
- Updatable views
- Updatable cursors
- JDBC driver enhancements

5.3 Other Features of the ESM Plug-in Module

5.3.1 Cross Platform Deployment

ESM is available for RedHat™ Linux, Turbo Linux™, and Microsoft Windows™.

5.3.2 Application Server Support

ESM can be integrated with Fujitsu's leading edge, high-performance Interstage Application Server. At the JDBC level, ESM is compatible with any Java based Application Server.

5.3.3 Ultra-fast Initial Database Loading

ESM provides initial loading performance improvements of over 10 times faster than the original PostgreSQL data load

5.3.4 Easy Upgrade from PostgreSQL

It takes around 40 minutes to upgrade from an original PostgreSQL database of 5 million records to ESM (database size of 1.3GB)

5.3.5 Tools

ESM comes with a set-up and management GUI to make installation easy. In addition, several open-source and third party, high quality GUI tools exist to both administer the database and carry out database design.

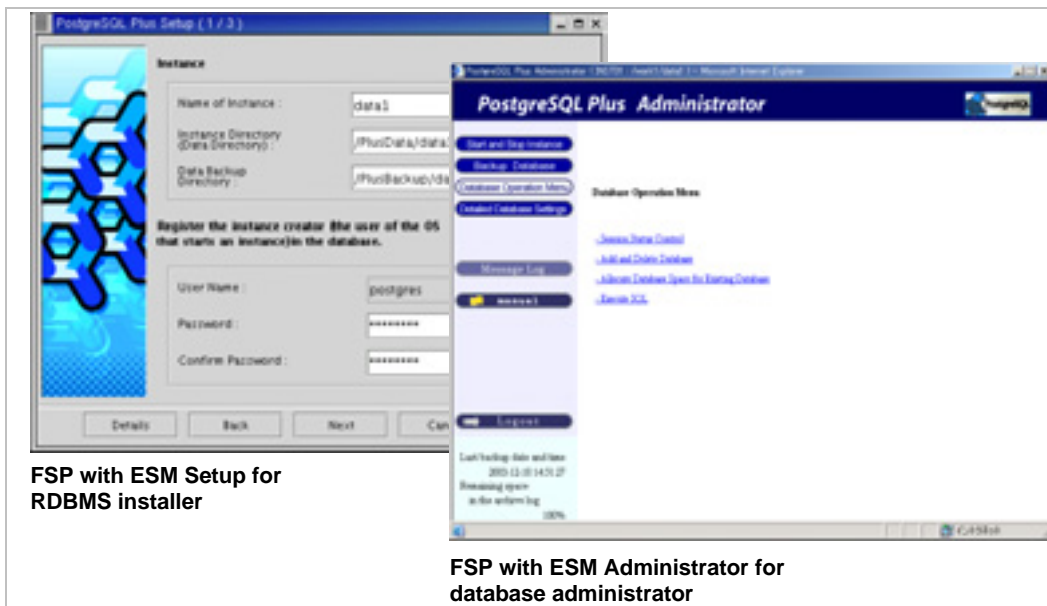


Figure 10 - Fujitsu Provides Enhanced GUI Set-up and Management Tools

Several performance-analysing features, including performance-related system views, are built in to ESM which help you get the most out of PostgreSQL.

6 The Open Source PostgreSQL Advantage

The ESM module integrates seamlessly with PostgreSQL and all the features of original PostgreSQL that have made it so successful remain available.

PostgreSQL is renowned for its reliability, performance, extensibility, SQL support, usability from a variety of popular languages and platforms and its support for a large number of procedural languages.

A Point List of Some Technical Features that PostgreSQL Offers	
Standard SQL Language Features	<ul style="list-style-type: none"> • ANSI SQL compliance • Views • Sequences • Functional and Partial indexes • Sub-selects • Support for UNION, UNION ALL and EXCEPT queries
Extensibility	<ul style="list-style-type: none"> • Outer Joins • Inheritance • Native interfaces for ODBC, JDBC, C, C++, PHP, Perl, TCL, ECPG, Python, and Ruby • Procedural languages • An open, documented backend customisation API • Stored Procedures • User defined types, operators, functions, indexes and aggregates.
Transaction Processing and Concurrency	<ul style="list-style-type: none"> • Fully ACID compliant • Better than row-level locking • Hot stand-by • Security and Data Integrity • Native SSL support • Native Kerberos authentication • Rules • Triggers • Referential Integrity
Integration and Connectivity	<ul style="list-style-type: none"> • Replication (non-commercial and commercial solutions) allowing the duplication of the master database to multiple slave machines • Tools for generating portable SQL to share with other SQL-compliant systems • Cross-database compatibility functions for easing the transition from other, less SQL-compliant RDBMS • Loadable extensions offering SHA1, MD5, XML, and other functionality

6.1 Reliability

Many users of PostgreSQL report having never experienced crashes or availability problems with PostgreSQL, even after many years in production. There are eight key reasons why PostgreSQL is so reliable.

6.1.1 ACID Compliance

PostgreSQL is an ACID compliant database. This means that:

- A – Database modifications are ATOMIC (An atomic operation is one which cannot ever be in an incomplete or inconsistent state. That is, the operation either succeeds or fails. This is important when working with databases since without atomic data modification; the potential for data integrity/corruption issues is present)
- C – Modifications to data are ensured to be CONSISTENT
- I – Data modifications are ISOLATED from one another using a versioning system
- D – PostgreSQL takes care to ensure that all modifications are guaranteed to be written to hard disk, to ensure that they are DURABLE

These four factors combine to ensure that, in a concurrent environment, data is not lost or corrupted.

6.1.2 Multi-version Concurrency Control

Multi-version concurrency control (MVCC) is a system by which the visibility of data modifications to isolated concurrent transactions is achieved by storing multiple versions of a row, along with data describing which transactions can see that row. This is a highly reliable system which scales extremely well.

6.1.3 Write Ahead Logging (WAL)

Write Ahead Logging (WAL) is a mechanism of efficiently journaling all data modifications so that in the case of a crash, all recent data modifications can be restored.

6.1.4 Point in Time Recovery (PITR)

From version 8.0, PostgreSQL supports Point in Time Recovery (PITR). Working in conjunction with WAL, this technology allows database administrators to restore the database to any point in the past. Better yet, no additional performance overhead is incurred by using PITR.

6.1.5 Replication

The replication capability of PostgreSQL databases further bolsters its reliability. The open and extensible architecture of the system has led to the design of several replication systems, both open and proprietary. The Slony PostgreSQL replication solution is the most well-known. Further information on this and other PostgreSQL replication products is available from the PostgreSQL Open Source Community web site (<http://www.postgresql.org/>).

6.1.6 Data Integrity

Data Integrity is at the heart of PostgreSQL. On top of MVCC, PostgreSQL also supports schema level mechanisms, including foreign keys and constraints, to ensure that data remains in a consistent state.

PostgreSQL also uses an internal dependency tracking model to ensure that any object which relies on the existence of another object cannot be deleted or removed without the dependency itself being removed.

6.1.7 The Open Source PostgreSQL Development Model

PostgreSQL has a rigorous, transparent development model. Public discussion between users and developers is held whenever significant feature additions are made. Details are hammered out and the design of the feature is examined closely.

The developer or developers then work, often closely with other developers and users, often submitting evaluation 'patches' so that other users can gauge their progress and assist with implementation details.

Once a modification or update is ready for inclusion in the software, it is released for review. This allows other developers to scrutinise the modification or update line by line and gauge the quality of the code and impact it will have on the rest of the source code tree. As necessary, patches are continuously reworked until the quality matches the rigorous and highly-exacting standards which PostgreSQL developers and users demand.

The patch must also pass a comprehensive series of tests, stored in the software source code tree, which any user has the ability to run in order to verify the durability of the patch.

PostgreSQL also has an extensive community of beta testers who further test these feature additions by running them in a development environment, with real world data and applications. This ensures that, upon release, PostgreSQL has extensive real world testing on top of renowned code level scrutiny and peer review.

6.1.8 Open Source

As an open source project, all source code for PostgreSQL is available to the public. This means that PostgreSQL can be obtained at no cost, with no on going license fees.

Moreover, PostgreSQL is licensed under a business-friendly Berkeley-style license. This means it can be used, modified or compiled and redistributed in any form – including binary format – without the need to provide source code to users.

6.2 Performance

PostgreSQL is a database system designed with performance in mind. The basis of this is its support for indexing and index partitioning, a sophisticated query planner with performance debugging capabilities, granular locking, next generation buffer management and caching, and excellent scalability under concurrent workloads.

6.2.1 Indexing Support

Indexing ensures that as the data in your database grows the cost of searching it remains near constant. PostgreSQL supports the industry standard B-Tree algorithm for indexing, as well as the hash, L-Tree and R-Tree algorithms. It also supports GiST, which allows users to write their own indexing systems (See 'Extensibility' below).

PostgreSQL also supports functional indexes. Such indexes allow you to create an index on the output of a function called on a column or columns of a table, greatly improving performance for many real world applications.

On top of this, PostgreSQL supports partitioning of data through partial indexes. A partial index is one which covers only a subset of the data in a table and as such is smaller than an index over the whole table. This means that applications only interested in part of a table can access data much faster.

6.2.2 The PostgreSQL Query Planner

PostgreSQL supports a cost based query planning system. This planner determines the cost of the different methods of accessing data from tables, the cost of joining them using a variety of algorithms, and the cost of sorting them. It then chooses the least expensive plan and uses it to retrieve the data.

PostgreSQL also provides users an interface to this process, so that they can easily debug the performance of their queries. This debugging information allows the users to determine where indexes should be created, whether memory dedicated to caching and sorting needs to be increased and whether the user can use more appropriate data types or SQL in their query. Such control over performance is vital to those evaluating PostgreSQL for enterprise use.

6.2.3 Granular Locking

PostgreSQL supports low level locking which ensures a high level of concurrent data access while protecting data integrity. Locking is supported at the table and row level. Lower level locking, for shared resources, is optimised for each software and hardware platform PostgreSQL supports.

6.2.4 Buffer Management and Caching

From 8.0, PostgreSQL supports a buffer management and caching system which employs a sophisticated algorithm to keep the database caching populated with the most in demand data. Since caching can greatly increase performance, this feature is fundamental to the success of applications in a performance critical environment.

6.2.5 Scalability

The above-mentioned features combine to create a highly scalable, highly available database system. A small memory footprint and granular locking mean that PostgreSQL can easily scale to tens of thousands of concurrent transactions on commodity hardware. Index and buffer management mean that the system remains responsive, even at high load levels.

6.3 Extensibility

Extensibility in PostgreSQL refers to the ability of the user or DBA to modify and customise many aspects of the system.

Users can define their own functions, aggregates, types, indexes and operators. This object orientated approach to data allows users to move their application logic into PostgreSQL itself and present complex functionality to all applications with access to the system. This also greatly simplifies client applications whilst ensuring that the business logic underlying them is tied to the transactional semantics and proven stability of PostgreSQL.

6.4 SQL Support

PostgreSQL has world class SQL support. On top of the features expected from any SQL database, PostgreSQL supports:

- A very high level of compliance with ANSI SQL 92, ANSI SQL 99 and ANSI SQL 2003 with documented extensions and enhancements
- Schemas, which provide an SQL-level namespacing mechanism
- Sub queries
- Outer joins
- Rules
- Views
- Cursors
- Prepared statements
- Stored procedures
- Savepoints/nested transactions
- Permissions based privilege system for each object in the system
- An interprocess messaging system
- Triggers

6.5 Rich Data Type System

PostgreSQL has one of the most feature rich type systems. Some highlights are:

6.5.1 Character Types

PostgreSQL supports SQL standard character types as well as text-type which can support strings of effectively unlimited length.

6.5.2 Numeric Types

PostgreSQL supports an arbitrary precision numeric type suitable for anything from basic accounting systems to the most demanding scientific and financial market applications.

6.5.3 Arrays

PostgreSQL has a first-rate SQL2003 array implementation. This makes dealing with complex application data easy and is critical in data warehousing contexts.

6.5.4 Large Objects

PostgreSQL supports large objects which allow you to store up to two gigabyte (2 GB) of binary data as a single retrieval object in the system.

6.5.5 Geometric Types

PostgreSQL is able to store two-dimensional spatial data. This includes points in space, lines (both finite and infinite), circles, polygons and boxes.

6.5.6 Geographic Information System (GIS) types

Proof of PostgreSQL's extensibility is a first class GIS implementation. This extension allows PostgreSQL to be used with industry standard client and middle ware applications as well as allow users to store their own three-dimensional data.

6.5.7 Network Types

PostgreSQL also supports types for Internet Protocol (IP) types – both version 4 and version 6. It also supports CIDR and mac address types.

6.5.8 Composite Types

A composite type is analogous to an object. It is composed of one or more elementary types and allows users to store complex data in a discrete format.

6.5.9 High Precision Dates and Time

PostgreSQL supports high precision date, time and interval types. This allows users to store dates and times from 4173 BC to 587489 AD with resolution to 1 microsecond.

6.6 Extensive Character Encoding Support

PostgreSQL supports over 25 different character encoding formats including ASCII, LATIN encoding types and UNICODE.

6.7 Native Language Support

PostgreSQL has documentation and error messages in a number of languages, including Japanese, German, Italian, French, Mandarin, Cantonese, Russian, Spanish, Portuguese, Slovenian and Slovakian.

6.8 Interfaces

PostgreSQL can be accessed from a variety of languages and data access mechanisms, such as JDBC and ODBC. Interfaces are available for at least the following languages: C, C++, Java, PHP, Perl, Python, Ruby, TCL, C# and other .NET languages, Lisp, as well as a variety of others. Many more, including COBOL, ColdFusion and Visual Basic can access PostgreSQL from ODBC or JDBC interfaces.

6.9 Procedural Languages

A procedural language is one which can be used by PostgreSQL users to develop functions. The database itself executes the function in the language it was written. These languages are tightly coupled with PostgreSQL meaning that you will have close contact to data and metadata.

Currently, there is support for Perl, Python, PHP, Java, Ruby, the R statistical language, the Unix shell, TCL and PostgreSQL's own PL/PgSQL language. Functions can also be written in C and C++ and run as binaries.

The real benefit of using a procedural language is that you can encode the logic of your application into the database and allow all parts of your application to make use of it. Better yet, you can expose that logic to other applications for data mining and integration purposes. This can dramatically decrease development time with little or no training, since developers can use the same language as that of the application they are writing or another language they are familiar with.

6.10 Ease of Use

PostgreSQL has a strong focus on ease of use. Its SQL console provides simplified mechanisms to display object information, a SQL quick reference, tab-completion, query history and buffering, and the ability to work in an interactive or streaming mode.

PostgreSQL is easy to tune and configure. PostgreSQL's setup tool determines the cache size suited to the system it is being installed upon and provides a single configuration file from which resource usage can be configured.

A significant benefit of this is that developers and administrators can spend less time on configuration issues.

6.11 Data Security

PostgreSQL has a four-tier data security system.

The first level of security is that PostgreSQL cannot be run as a privileged user. This guards against potential system-takeover problems.

The second level of security is SSL encryption of client/server connections. This guards against sniffing of data in a network context.

Third is PostgreSQL's sophisticated authentication system. Users can be authenticated by host or IP address/network. PostgreSQL can force password authentication, encrypted password authentication, Kerberos authentication, IDENT based authentication and a range of other mechanisms using Pluggable Authentication Modules.

Finally, PostgreSQL supports granular permissions based access to all database objects based on user and access method. Combined with schemas, which provides namespace isolation for each user, PostgreSQL has a highly sophisticated security framework.

7 Fujitsu's Interstage and FSP

FSP and ESM can be used stand-alone, or they can form part of a complete Enterprise-level IT solution.

7.1 The Interstage Suite

Fujitsu's Interstage suite is a broad family of modular, flexible real-time software products that allow enterprises and their partners to design, develop, deploy, and manage scalable, customised mission-critical applications securely and reliably. It is a perfect complement to the FSP database solution. Fujitsu Interstage is an integrated solution that allows individual components to be deployed on an "as-needed" basis.

7.2 Interstage Application Server Foundation

Interstage provides a rich set of web-based middleware application server solution products with JDBC level integration. The Interstage Application Server is the foundation of secure, reliable and effective Internet and Intranet applications. It's accompanying components and development tools and other components make it a complete solution for enterprises of any size. It includes all the standardised technologies needed for web services today, including ebXML, RosettaNet, UDDI, SOAP, WSDL and J2EE.

7.3 Interstage Integration Products

Interstage also provides a rich set of Integration Products. This includes a full-featured, comprehensive Business Process Management (BPM) solution designed to improve productivity of businesses by maximising the value of IT infrastructure. It also includes:

- Interstage Portalworks (a front-end integration product that organises and integrates content, applications, people and processes),
- Interstage Content Integrator (a scalable application platform for enterprise content integration), and
- Interstage CollaborationRing (a conductor for business-critical systems, orchestrating between partners, applications and external services).

8 Fujitsu and Open Source Software

8.1 Our Commitment

Fujitsu has in recent times, initiated a number of successful collaborations with the open source community. This mutually beneficial relationship adds materially to community resources, assuring the long-term future, and continuing high quality of products released under the Open Source licensing model. The proven reliability and applicability of products developed in collaboration with Fujitsu, ensures that wherever possible, new opportunities to extend and grow this relationship, will be actively pursued.

As an active member of the Open Source Development Lab (ODSL) Fujitsu is also strongly behind the push to increase adoption of Open Source technology across the widest possible range of industry and government sectors.

8.2 Working with Partners

As an intrinsic part of its long-term strategic commitment to Open Source Software, Fujitsu has partnered with a number of global organisations such as RedHat, SuSE Linux and SRA, to add value and extend products developed by the Open Source community.

This commitment will be the hallmark of partnerships and alliances that FAST will build in Australia in the near future as it works to establish PostgreSQL as the automatic choice of organisation's seeking a low-cost Enterprise-ready Database solution.

8.3 Support of FSP

In addition to the direct support options offered by Fujitsu, there is a vibrant community of PostgreSQL professionals and enthusiasts that your company can call upon as needed.

Fujitsu Australia Software Technology (FAST) will operate support of PostgreSQL in Australia.

9 About Fujitsu Australia Software Technology (FAST)

Fujitsu Australia Software Technology (FAST) was established in April 1989 by Fujitsu Limited, a world leader in Information Technology and was the first Fujitsu software development company to be formed outside Japan.

Our mission is to build solutions using advanced or emerging software technologies to achieve leadership in information technology for which FAST is renowned today. Our core technical competencies include:

- Enterprise computing
- Information Management
- Digital Content Scheduling and delivery Systems

Products developed by FAST include Interstage, a full Enterprise-grade Application server with end-to-end Business Process Integration capabilities, and TELentice, the world-leading Digital Media Networking solution.

9.1 Contact Us

To find out about how FSP can deliver real business benefits to your organisation, please call:

Liam O'Duibhir, FSP Product Manager, on +61-2-9452-9068

Email: postgresql@fast.fujitsu.com.au

Fujitsu Australia Software Technology Pty Ltd

14 Rodborough Rd, Frenchs Forest, NSW, Australia 2086

9.2 On the Web

Additional information on FSP can be viewed at the website:

www.fastware.com/postgresql.html

10 Appendix I – Glossary of Terms and Abbreviations

FSP	Fujitsu Supported PostgreSQL
DB	Database
ESM	Extended Storage Management
Open Source PostgreSQL Community	The participants (principally software developers) in the Open Source initiative who develop, popularise and promote the PostgreSQL database. There are hundreds of individuals involved in this initiative from almost every country within the world.
FAST	Fujitsu Australia Software Technology
PostgreSQL Plus	The developer-level package name for the ESM module is PostgreSQL Plus.
TCO	Total Cost of Ownership – effectively the ‘true cost’ of a system or technology when all costs, direct and indirect, are taken into account.
Source Code	All software applications are written using specialised languages which define the instructions to be followed by a computer and interactions with users. Popular source code computer languages include ‘C’ and Java.
CIDR	Classless Inter-Domain Routing - a new addressing scheme for the Internet which allows for more efficient allocation of IP addresses than the old Class A, B, and C address scheme