FUJITSU

Fujitsu Enterprise Postgres
Open source value, enterprise quality

Strong growth in Database Management Systems (DBMSs) is expected to continue, making DBMSs the largest single cost factor in enterprise software

Fujitsu Enterprise Postgres is the latest DBMS offering from Fujitsu, built using Postgres technology. Fujitsu Enterprise Postgres provides enterprise features, quality, performance and reliability, while significantly reducing infrastructure software costs.

Don't sacrifice service quality to cut costs

Most industries now operate in highly competitive markets; organizations must offer customers value for money to stay in business. To remain competitive, it is common practice for organizations to implement measures to improve their efficiency and reduce costs. Infrastructure is one of the first areas where a Chief Information Officer (CIO) looks to make cuts. While reducing costs is a good way to help improve an organization's ability to be competitive, there are perceived risks involved.

Fujitsu Enterprise Postgres minimizes this risk by combining Fujitsu's 40 years of experience in developing enterprise databases with the strength and reliability of PostgreSQL, providing a secure, reliable, feature-rich enterprise-ready database. PostgreSQL has millions of users globally with its legendary reliability; and now an enterprise version packaged with a "no lock-in" commercial level support arrangement can provide CIOs with a very low risk option to make significant savings and improve operational efficiency.

Another reason that organizations shy away from changing expensive proprietary DBMS vendors to other options is the perception that the change process would be costly and difficult. Fujitsu Enterprise Postgres is made to mitigate this risk. Not only are the operational costs of Fujitsu Enterprise Postgres significantly lower than other proprietary databases, but the initial cost of migration is minimal. Full compatibility with open source PostgreSQL databases is a major factor in ensuring smooth and low cost migration.

The migration process:

- Analyze existing databases and produce reports identifying changes required for successful migration
- Estimate workload and time frames required for migration
- Identify syntax changes required to Data Definition Language and stored procedures for successful migration
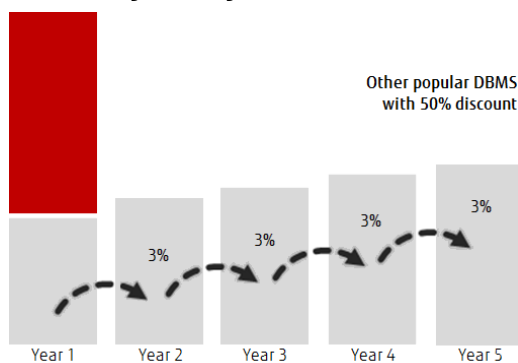- Perform and manage the migration

Figure 1: High costs due to yearly increase in maintenance costs
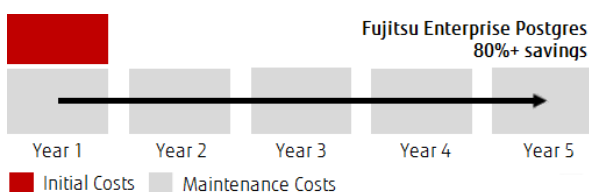

Figure 2: Fixed costs with Fujitsu Enterprise Postgres

Features

Fujitsu Enterprise Postgres is strongly ANSI SQL:2011 compliant and shares all the features of open source PostgreSQL, including full ACID compliance, inheritance, database server clustering, Unicode, multi-version concurrency control (MVCC), point-in-time recovery, asynchronous replication, locale awareness, nested transactions, WAL fault tolerance, and works on all major operating systems[†].

Fujitsu Enterprise Postgres extends those features with the following:

- Transparent Data Encryption >
- Data Masking >
- Dedicated audit log >
- Mirroring Controller >
- Connection Manager >
- In-Memory Columnar Index[*] >
- Global Meta Cache >
- High-Speed Data Load >
- High-Speed Backup/Recovery >
- One-click backup and recovery
- WAL duplication
- Enhanced GUI management tool
- System usage statistics

High performance is synonymous with enterprise solutions; Fujitsu Enterprise Postgres leverages all aspects that it can to deliver enterprise-quality performance.

[†]: Fujitsu Enterprise Postgres supports Windows and Linux operating systems. For a list of supported environments, refer to the Fujitsu Enterprise Postgres Datasheet⧉.

[*]: Implemented via Vertical Clustered Index (VCI) by Fujitsu Laboratories Limited. This feature is available only in the Advanced Edition of Fujitsu Enterprise Postgres.

Improved ease of use

One of Fujitsu's overarching software design goals is to build intuitive products. Larger and more complex systems are intrinsically susceptible to failures in reliability and performance. By building smart-safety layers into Fujitsu Enterprise Postgres to configure important features of the system and monitor various layers of the system for potential failures, the overall reliability and performance of the system is improved.
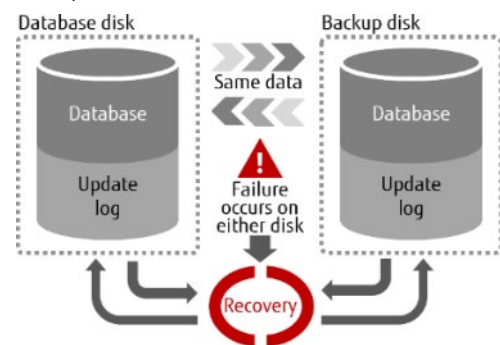

Figure 3: Recovery in Fujitsu Enterprise Postgres

An example is the default database cluster configuration performed by the WebAdmin tool. At a minimum, an installation of a Database Management System should be capable of performing backup and recovery operations. For this reason, Fujitsu Enterprise Postgres has this built into its database cluster creation utilities. When creating a database cluster with the WebAdmin tool that comes with Fujitsu Enterprise Postgres, you specify not only the location of where the database cluster will be stored on disk, but also where the backup files and copies of the Write Ahead Log (WAL) are to be stored.

Backups of the database cluster can be executed with ease through the WebAdmin user interface, and WALs are automatically streamed by the system. Fujitsu Enterprise Postgres performs smart monitoring of the disks, network and server processes, looking for potential issues. When an issue has been identified by the system, a point-in-time system recovery can be initiated by the click of a button.

The database server is automatically configured using optimal settings for the machine on which it is being installed. Database cluster creation also examines the host machine specifications such as processor, memory, and network connections to determine the optimal settings. Of course, these can be overridden by an experienced administrator, who will need to account for factors undeterminable from environmental specifications, such as how the database will be used. For example, will the database be executing low-frequency long-running BI-style queries or high-frequency short-running OLTP-style queries?

Enhanced security

Transparent Data Encryption

Fujitsu Enterprise Postgres comes with Transparent Data Encryption out of the box; unlike some other proprietary DBMS products, you are not required to purchase additional products to use this functionality.

Transparent Data Encryption provides encryption at the file level, essentially solving the issue of protecting data at rest. Its main purpose is to fulfil requirements for the Payment Card Industry Data Security Standard (PCI DSS) and allows confidential information such as credit card numbers to be made unrecognizable on disk.
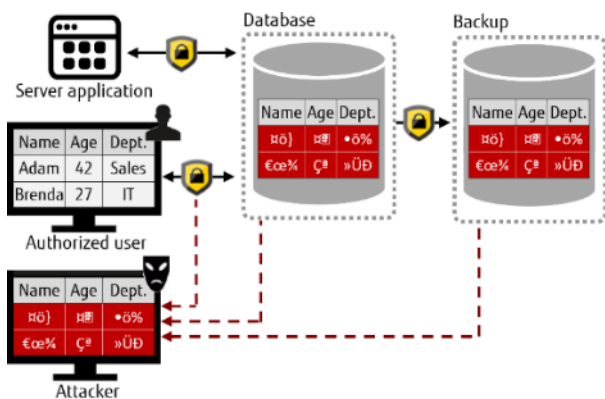

Figure 4: Production and backup data are protected even when under attack

Data is automatically encrypted and decrypted when it is written and read; manual key management is not required. Even if an attacker gets through all the access controls and connects to the server, they will not be able to access the data as the OS file is encrypted.

Backups of encrypted data stay encrypted, WAL data is encrypted, streaming replication data is encrypted and even temporary files generated through sort operations are encrypted.

Read more White paper: Transparent Data Encryption ☑

Data Masking

Data masking is the obscuring of data in a database. Fujitsu Enterprise Postgres uses masking policies to obscure data returned from queries, making it available for reference without exposing the actual data.

The data masking feature can be implemented in a variety of methods, including: character shuffling, nulling or deletion, encryption, masking, and word substitution.

Data masking makes it safe to use production data in a test or development environment.
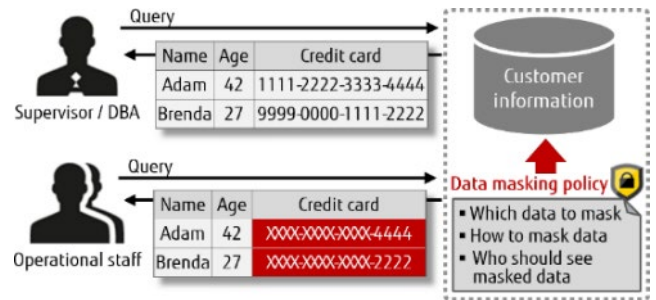

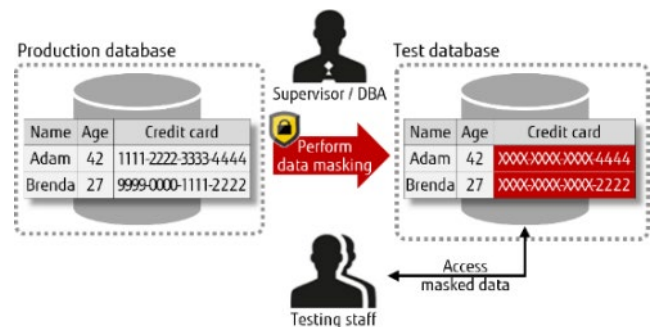Figure 5: Data masking in a production environment

Figure 6: Data masking in a test environment

Read more White paper: Data Masking ☑

Dedicated Audit Log

Fujitsu Enterprise Postgres provides a dedicated audit log that is compliant with all Payment Card Industry Data Security Standard (PCI DSS) requirements. This audit log can be used to counter security threats such as spoofing, unauthorized access to the database, application manipulation, and misuse of privileges.
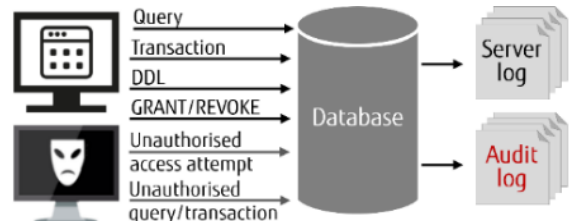

Figure 7: Dedicated audit log

PostgreSQL outputs audit log records to the server log only, which makes operation and analysis complex (since some logs, such as SQL runtime logs, do not output the schema name) and time consuming (because output criteria cannot be specified, and since log volumes can be large, performance will deteriorate). On the other hand, the Fujitsu Enterprise Postgres audit log allows a flexible and efficient way to accurately monitor logs.

Read more White paper: Dedicated Audit Log ☑

High performance options

In-Memory Columnar Index

The VCI engine, which is integrated with Fujitsu Enterprise Postgres, provides significantly faster analytical query processing by storing a columnar representation of row-oriented data in memory. Our tests show that for a 280 GB dataset on a 56-core Linux node, this results in almost 5 times the throughput of analytical queries while maintaining equivalent transaction volumes.

VCI reflects updated row-oriented data to its columnar equivalent, stores it without being dependent on

Figure 8: In-Memory Columnar Index

memory capacity, and quickly conducts analysis of that data. Massive volumes of column-oriented data can be stored by taking advantage of this new technique for managing data.
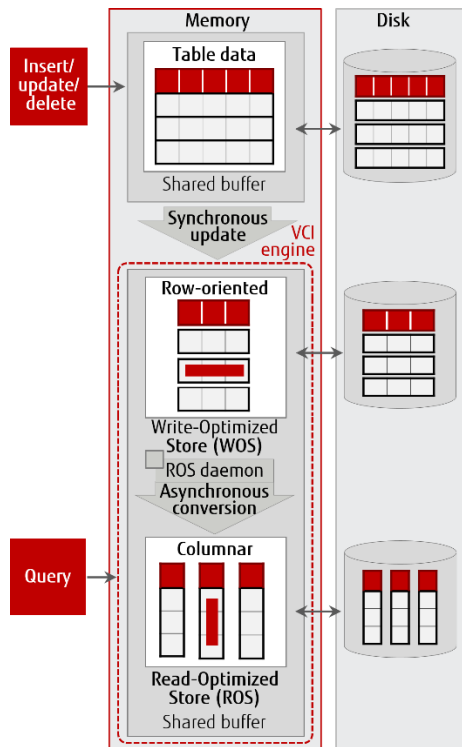
> **Read more** White paper: Improve data analysis performance without impacting business transaction processing ⬈

High-Speed Data Load

Fujitsu Enterprise Postgres offers the ability to load large volumes of data using as many parallel processes as possible, given the number of cores and CPU availability.

Data load features of other database systems can only use a fixed number of parallel process, regardless of CPU availability. In an era of mission-critical systems that are able to take advantage of multi-core technology, this is not the optimal solution. Fujitsu Enterprise Postgres high-speed data load sends data from the input file to several parallel workers, each of which will simultaneously perform data conversion, table creation, and index creation.
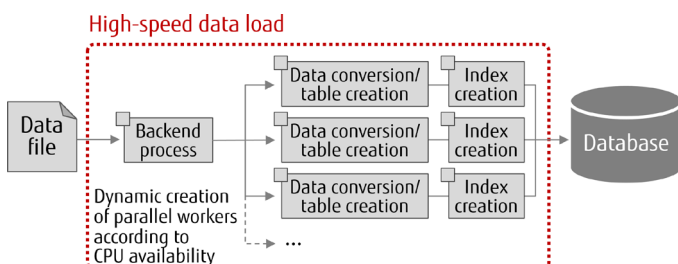

Figure 9: High-speed data load and dynamic creation of parallel workers

> **Read more** White paper: High-Speed Data Load ⬈

High-Speed Backup/Recovery

The backup/recovery utility now enables backup/recovery of clusters and selected tablespaces using any copy method specified by the user, to replace the conventional, file-level copy process and allow TB-scale operations. The pgx_dmpall command and the pgx_rcvall command can now perform backup/recovery using user exits, which can redirect the process to use high-speed copy technology.
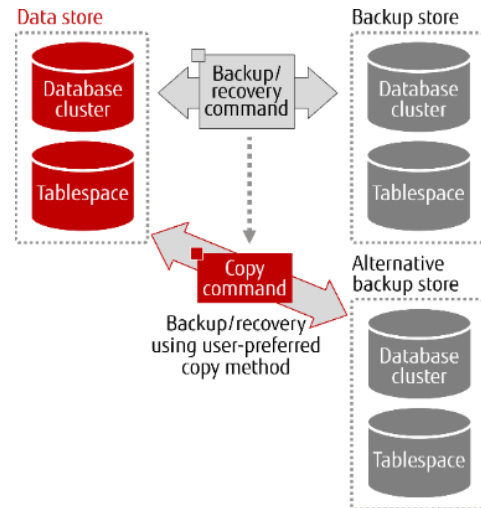

Figure 10: Backup/recovery using high-speed copy technology of choice

> **Read more** Brochure: High-Speed Backup/Recovery ⬈

Global Meta Cache

Before Fujitsu Enterprise Postgres 11, information about system catalogs and table definitions was cached in per-process memory. As for version 12, you can cache some of this information in shared memory by using the Global Meta Cache feature. This reduces overall system memory usage without causing loss of data consistency.
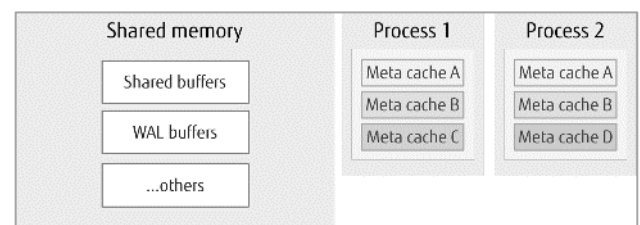

Figure 11: Cache memory allocation without using Global Meta Cache
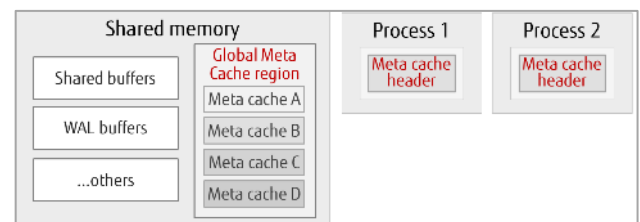

Figure 12: Cache memory allocation using Global Meta Cache

High Availability options

High availability (HA) is not just about keeping a server up and running; high availability requirements can be different from organization to organization or even system to system. Some organizations put a high emphasis on data durability, some on system reliability, and others on performance stability.

Fujitsu Enterprise Postgres has the flexibility to be configured to meet all the different HA requirements that an organization might have. An organization can utilize the many open source and third-party products available for open source PostgreSQL with Fujitsu Enterprise Postgres, maintaining that flexibility to modify storage architecture without being "locked in".

Fujitsu Enterprise Postgres also provides some additional failover options using proven proprietary Enterprise clustering products including:

- PRIMECLUSTER
- Microsoft Cluster Server

> **Read more** White paper: Database High Availability - Solutions and Implementations ⬀

Mirroring Controller

Mirroring Controller performs two main functions:

- The Mirroring Controller agents constantly monitor your operating system, server, disk, network, and database, and notify you if something is amiss.
- If an abnormality is detected, Mirroring Controller performs automatic switch/disconnect — if the primary server fails, processing is switched to the standby server to ensure operational continuity, and the primary server is disconnected.

You have the option of switching from the primary server to the standby server manually, or setting the Mirroring Controller to do so automatically, ensuring there are no interruptions to your database operations.

Mirroring Controller is installed on each database server. The Mirroring Controller agents communicate with each other and continually monitor server processes and operating system, network, and disk state — they even monitor the actual Mirroring Controller processes.
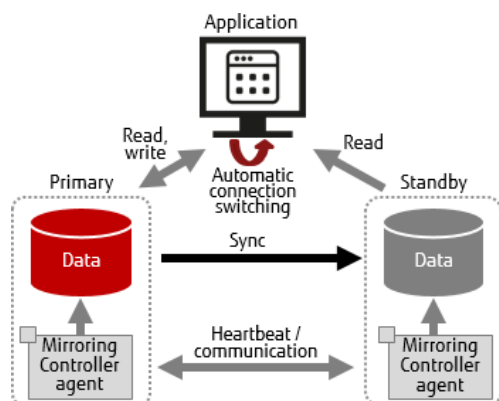

Figure 13: Mirroring Controller agents performing heartbeat monitoring

> **Read more** White paper: Mirroring Controller ⬀

Connection Manager

Fujitsu Enterprise Postgres also provides Connection Manager to enhance system availability. It implements heartbeat monitoring and transparent connection features so that applications can connect to the appropriate database server without being aware of the server state. It further improves business continuity by allowing faster database server switchover.
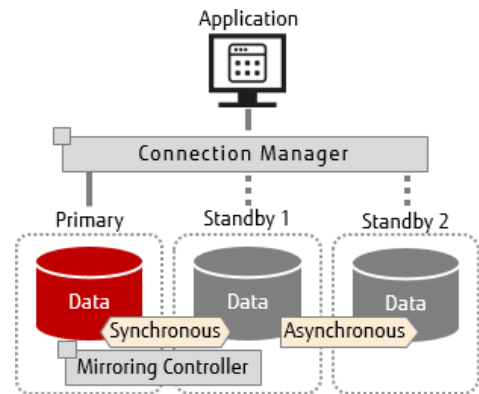

Figure 14: Application transparently connecting to database server

Application development

In addition to ODBC, Fujitsu Enterprise Postgres supports the same rich choice of client APIs available to PostgreSQL. There are multiple APIs available for the more popular development languages, including, but not limited to the following:

- C/C++
- .NET languages (such as C#, VB.NET)
- Perl
- Python
- Java
- PHP
- TCL
- LISP
- Ruby on Rails

Most of the mainstream IDEs have support for the Fujitsu Enterprise Postgres interface, allowing important features such as browsing databases and implementations of Object Relational Mapping. For example, all three models (code first, model first and database first) are supported under Microsoft's Visual Studio.

Entity Framework is now developed by an open source community separate from, but in consultation with, Microsoft. This enables support for Fujitsu Enterprise Postgres to be provided in a timely and accurate manner.

The community-supported framework provides a Data Designer Extensibility (DDEX) package allowing the developer to use the existing Visual Studio designer components such as the model browser and database explorer.

The diagram below illustrates how the low-level framework reduces the effort required to change .NET applications to use Fujitsu Enterprise Postgres. With the abstraction provided by Language Integrated Query, the Entity Client Data Provider manages translation to Fujitsu Enterprise Postgres.
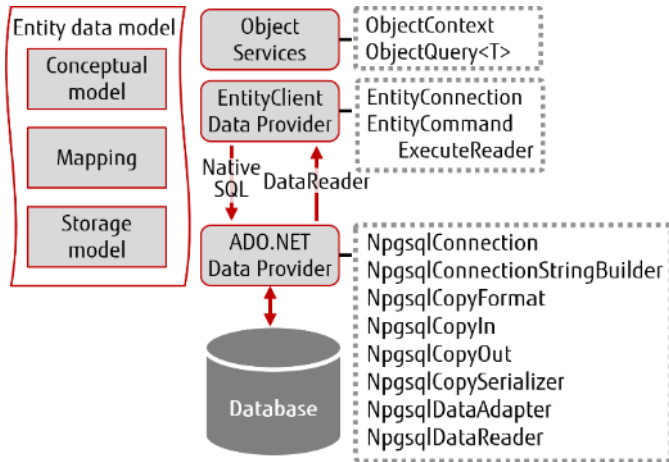

Figure 15: DDEX components

The Eclipse IDE also provides support for Fujitsu Enterprise Postgres through the Data Tools Platform Extender (DTP) plugin and JDBC. Database objects can be browsed through the Data Source Explorer where data can be viewed and edited. Queries can be executed via the SQL Results window and Query plans can be examined via the Execution Plan window.

Java ORM frameworks such as Hibernate work with Fujitsu Enterprise Postgres without the need to install additional libraries or drivers beyond the community-provided JDBC driver.

Integrating openly with widely adopted technologies such as Visual Studio, Eclipse, Entity Framework and Hibernate, allows organizations to utilize more up-to-date versions of technology by not having to wait for releases of proprietary wrappers and interfaces to catch up.

Commercial support offering

Fujitsu's commitment to its pledge "Shaping tomorrow with you", which reflects the desire for customer success, is evident by its extension of the commercial support and training arrangements beyond Fujitsu Enterprise Postgres. Fujitsu also provides flexible support arrangements to its customers for the open source PostgreSQL product.

Fujitsu is part of the PostgreSQL open source community and contributes not only to PostgreSQL but also to many open source projects. Fujitsu is an active sponsor of key PostgreSQL events such as PGCon, PgConf.ASIA and pgDay Asia. At these events, Fujitsu readily provides feedback in pursuit of better solutions to customer issues.

Timely and cost-effective migration

Changing vendors is never easy, especially with something as foundational as a DBMS. The decision to change vendors can become quite complex when looking beyond the basic ongoing operational return on investment (ROI). Things can go wrong with migrations, workload can be greatly underestimated, and technical problems seem to appear even though they should not. Company growth in the future may require different features; or scalability of products can change previously predicted licensing costs. These risks make the decision to change vendors very muddied.

Fujitsu Enterprise Postgres has been designed to mitigate some of the complexity, risk and cost associated with a migration, restoring some clarity in the decision of whether to change vendors:
- No "vendor lock-in"
- Full compatibility with open source PostgreSQL databases
- Enhanced compatibility with major proprietary database systems
- Library of tools that simplify the analysis, workload planning, execution and reporting of migrations

Summary

In today's increasingly competitive marketplace, organizations need to be more agile and responsive than ever before. A rich feature set like that available in Fujitsu Enterprise Postgres covers all your enterprise database requirements from extendable database typing to the strongest 256-bit transparent data encryption in the marketplace to protect and utilize your data. Delivery of a full feature set of enterprise quality has not compromised ease of use; DBAs and developers can be productive in a very short time frame thanks to compliance with industry standards and established open source interfaces. Like any other part of an organization, information technology infrastructure should be continuously evaluated for efficiency, and part of that can mean reducing cost. Not only has Fujitsu Enterprise Postgres been proven to have much lower operating costs than other proprietary DBMS vendors, but it reduces the risk and cost associated with changing vendors without "vendor lock-in"; organizations can now have the flexibility and freedom that is required in today's business world.

Contact us

If you have any questions about the Transparent Data Encryption and other enterprise security features of Fujitsu Enterprise Postgres, please contact us at enterprisepostgresql@fujitsu.com ✉.

About Fujitsu

Fujitsu is the 5th largest IT service provider in the world, offering a full range of technology products, solutions, and services. Over 126,000 Fujitsu employees support customers in over 100 countries.