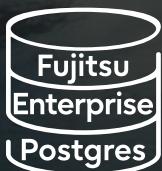


Fujitsu Enterprise Postgres 18

Enterprise AI
without
compromise





AI-Driven. Multi-Master Strong. Enterprise-Ready.



With the rapid evolution of AI technology, adoption of AI in the enterprise applications is steadily growing. The databases that were designed to store and process vast amount of data efficiently, are evolving to meet the new demands of AI based applications. Leveraging their strength in high availability, reliability and security, databases are increasingly adding new features to ensure a robust data infrastructure for AI applications and collaborations.

Fujitsu Enterprise Postgres is based on open-source PostgreSQL and includes additional, Fujitsu's developed, proprietary features to enhance security and reliability to serve enterprise applications. The enhanced functional capabilities aim to robustly respond to the market demands and support corporate data utilization at elevated level. Relational databases are no longer mere data repositories. These are increasingly taking role of robust enterprise data infrastructure for AI applications and collaborations and creating business value.

Fujitsu Enterprise Postgres, PostgreSQL reliability with Fujitsu innovation

Fujitsu Enterprise Postgres elevates PostgreSQL by adding powerful AI, stronger security, improved operability, and robust performance and DevOps features

| Fujitsu Enterprise Postgres unique feature | | PostgreSQL extension |
|---|--|---|
|  AI | <ul style="list-style-type: none">• In-database inference ►• Generative AI for enterprise ►• RAG application development ►• Hybrid search ►• Knowledge Data Management ► |  <ul style="list-style-type: none">• pgvector• pgvectorschale• pgai |
|  Reliability | <ul style="list-style-type: none">• Multi-master replication ►• Connection Manager ►• Mirroring Controller ►• WAL duplication ►• High-Speed Backup/Recovery ► | <ul style="list-style-type: none">• Pgpool-II• pgactive |
|  Security | <ul style="list-style-type: none">• Transparent Data Encryption ►• Data Masking ►• Dedicated Audit Log ►• Confidentiality management ►• Policy-based login security ►• Cloud-based key management ► | <ul style="list-style-type: none">• ldap2pg |
|  Operability | <ul style="list-style-type: none">• Global Meta Cache ►• High-Speed Backup/Recovery ►• Operator• WebAdmin | <ul style="list-style-type: none">• pgvector• PostGIS• pgBadger• pgAdmin4• pgBackRest• pg_repack |
|  Performance | <ul style="list-style-type: none">• Vertical Clustered Index ►• High-speed Data load ►• Parallel scan ►• Scale out | <ul style="list-style-type: none">• pg_dbms_stats• pg_hint_plan• pg_bigm |
|  DevOps | <ul style="list-style-type: none">• 3rd-party database conversion syntax• NCHAR• ECOBPG | <ul style="list-style-type: none">• orafce |

Ready to take your
business to the next level?

View the main features of Fujitsu Enterprise Postgres
and how they take PostgreSQL to the next level

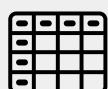
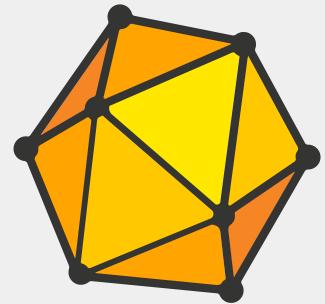
Learn more >



In-database Inference

Boosting the Knowledge Data Management function

Retrieval-Augmented Generation (RAG) is widely used to improve the accuracy of responses for LLM based applications. **Fujitsu Enterprise Postgres** provides the built-in capability for supporting RAG applications that query text data stored in its tables. It provides set of functions and procedures, categorized as Knowledge Data Management. These can be used to setup, administer and monitor execution of embedding models for vectorization of text data stored in the postgres table and query them for semantic searches from psql sessions.



Relation data

Accessed via
SELECT queries



Text data

Accessed via
full-text searches



Vector data

Semantic text and
vector similarity searches



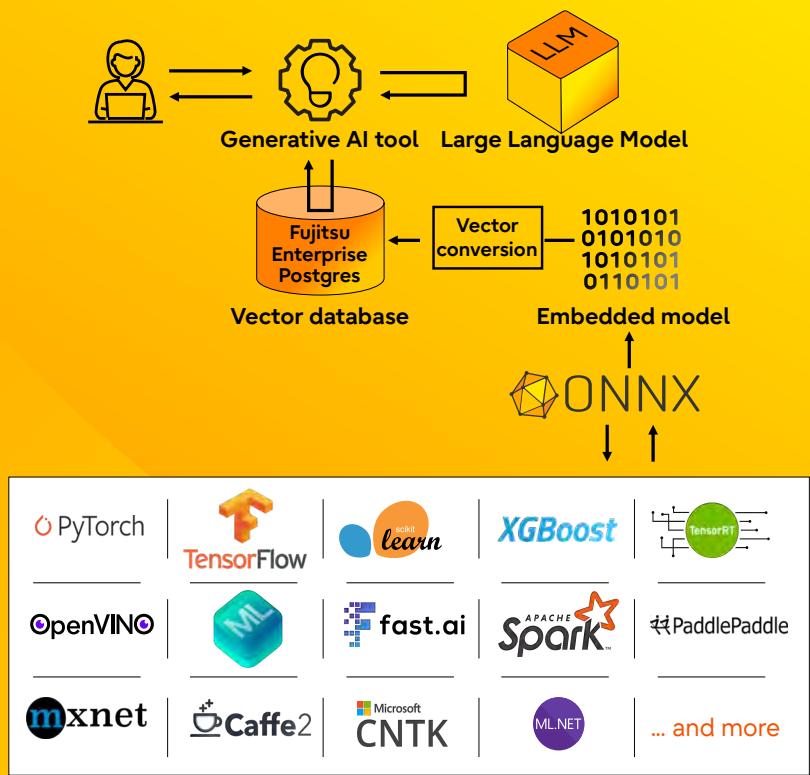
Graph data

Accessed via
graph searches

Fujitsu Enterprise Postgres 18 has enhanced its Knowledge Data Management capabilities to strongly support AI application development. In addition to supporting execution of embedding models from Ollama and OpenAI (capability available since version 17 SP1) for vectorization of text data, it now supports execution of ONNX format embedding models through Triton inference server.

Built-In ONNX powers advanced vectorization and RAG development

It is now possible to import ONNX (Open Neural Network eXchange) format embedding models directly into the database to perform vectorization of text data stored in tables. It provides comprehensive set of functions and procedures for administering and monitoring vectorization process from inside the database through psql commands. These built-in functions provide enhanced capability for RAG application development and



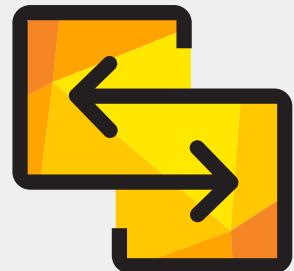
Supercharge your AI with fast, trusted, enterprise-grade data

View all Fujitsu Enterprise Postgres capabilities and how they create a rock-solid foundation for AI workloads

Learn more >

Multi-Master Replication

Enhancing support for mission-critical applications



In today's world, the advance of business globalization and digital transformation causes system downtime to directly translate to a loss of business opportunities.

No matter how good a service is, it's meaningless if data is unavailable. The importance of keeping data constantly available without system interruption is increasing. We address this requirement with **Fujitsu Enterprise Postgres** Multi-Master Replication, dramatically improving availability.

Data synchronization via bidirectional replication

Data is constantly synchronized between regions through bidirectional replication, ensuring every node holds the latest data. No complex setup or special tools are required

High Availability through database duplication

Within each region, high availability is ensured by database duplication (primary/standby configuration), so your systems can respond to a wide range of issues, from widespread disasters to localized failures

Business continuity even during disasters

By having each region mutually possess data, business can continue by switching to an active region even when a disaster occurs. This minimizes overall system downtime.

Rapid recovery with delta reflection only

During recovery, only delta changes need to be reflected, eliminating the need for long waits for resynchronization (full copy). This enables rapid recovery.

Replication settings

Our database allows you to specify which tables to include and exclude from replication. Since each region contains the latest data, business can continue by switching to an active region even when a disaster occurs.

Effective resource utilization

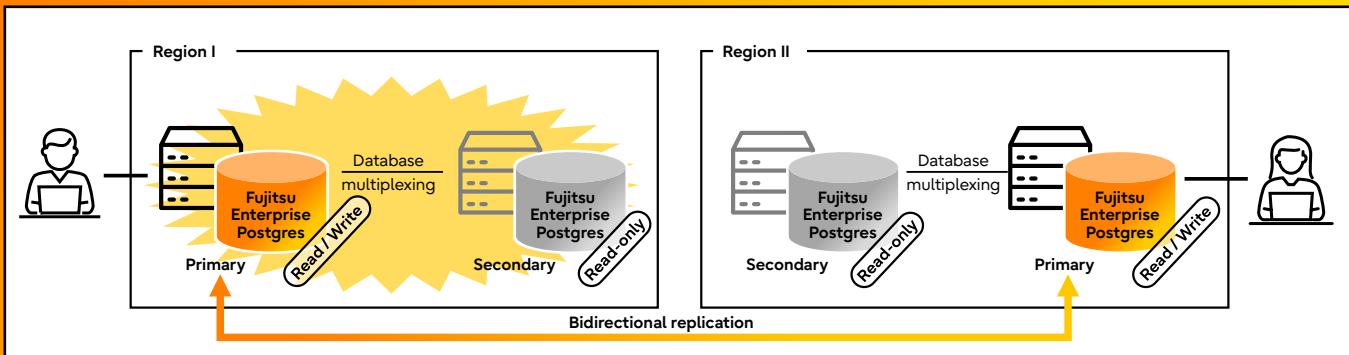
Since all nodes are operational, there's no need for a dedicated standby region for disaster recovery, maximizing resource utilization.



What is Multi-Master Replication?

Multiple database servers (nodes) act as masters and synchronize data in real-time, allowing data to be read from/written to any node, so that each node always holds the latest data.

In a typical single-master configuration, the entire system may stop if the master goes down. But with a multi-master configuration, even if one master node is affected by a disaster, business can continue with other active master nodes, minimizing business interruption.



Fujitsu Enterprise Postgres 18 - Enterprise AI without compromise

Fujitsu Enterprise Postgres 18 provides highly secure, reliable and available enterprise data infrastructure with enhanced AI capabilities.

Enterprises looking to accelerate their Digital Transformation (DX) utilizing AI in their applications for delivering critical enterprise functions with security, resilience, and flexibility can look forward to using Fujitsu Enterprise Postgres 18.



Discover more

Visit our website to learn more about **Fujitsu Enterprise Postgres** or to contact us and discuss your database needs.

Website: fast.fujitsu.com 

Contact: enterprisepostgresql@fujitsu.com 