


Fujitsu Enterprise Postgres for Kubernetes

Datasheet



Kubernetes is gaining attention as a virtualization technology that makes it easier to deploy and scale applications for shorter development cycles and more efficient operations. Fujitsu extends its own technologies to embrace containers and deliver a database system that empowers customer journeys.

Fujitsu's contribution to PostgreSQL

Fujitsu has been a keen player in open source development since open source database systems, Fujitsu is proud of its commitment to the promotion of PostgreSQL as a world-class enterprise database.

Fujitsu has been supporting PostgreSQL in various ways since 2004. One of the earliest contributions of Fujitsu to PostgreSQL was features in version 8.0. Since then, Fujitsu has been a proactive participant in the PostgreSQL community, providing Platinum sponsorship of key PostgreSQL events and supporting various PostgreSQL user groups.

The PostgreSQL-based database system is the outcome of Fujitsu's experience of developing enterprise databases for many years and contributions to the PostgreSQL community for the world's most advanced open source database system. The strength of PostgreSQL with Fujitsu's added enhanced features make Fujitsu Enterprise Postgres.


Fujitsu Enterprise Postgres

Fujitsu Enterprise Postgres is a PostgreSQL-based relational database management system for enterprises that offers open source value, enterprise quality, and mission-critical use.

Fujitsu Enterprise Postgres leverages PostgreSQL, a feature-rich open source database management system, used by millions of users globally, enabling integration with a wide range of software, information utilization systems, development tools, and application runtime environments.

Full compatibility with open source PostgreSQL allows users to enjoy the benefits of open source in enterprise quality. Database systems remain free from vendor lock-in, while achieving cost-effectiveness and high reliability. Fujitsu's strong track record in mission-critical enterprise systems supports this technology.


The improved high reliability has created substantial benefits to further compliment intelligent business data systems for enterprises.



Fujitsu Enterprise Postgres for Kubernetes

Fujitsu Enterprise Postgres for Kubernetes provides a framework for installing and managing your Fujitsu Enterprise Postgres on Red Hat OpenShift Container Platform. The Operator module manages the service on behalf of the Database Administrator.

Operators can deploy and manage Fujitsu Enterprise Postgres clusters, Pgpool-II clusters, and other components in an OCP environment.



The Operator is deployed as a standalone module, along with a highly available Fujitsu Enterprise Postgres cluster with pre-defined configuration to get started with a small workload. Users can adjust the configuration parameters at the time of deployment and during their use to optimize the instance according to workload.

The database is managed consistently according to predefined *custom resources*. Custom resources are YAML files that define the configuration of the system and are passed to the Operator.



As the name implies, the server container is intended to incorporate the Fujitsu Enterprise Postgres server component. In principle, a running server container is considered as an equivalent to a Fujitsu Enterprise Postgres Server instance.

Features and benefits

Feature	Details
Deployment	
• Data center	• Within one data center
• Model	• One master, two replicas / One master
• Replication type	• Asynchronous / Synchronous / Logical Replication
• Scaling	• CPU / Memory
Load balancing	
• Using Pgpool-II	
High availability	
• Failover type	• Automatic (read more below)
• Switchover type	• Manual (read more below)
• Auto-recovery	• Automatic (read more below)
Backup and Restore	
• Frequency	• Configurable
• Generation	• Configurable
• Backup type	• Full / Incremental
• Restore type	• Latest / PITR
• Restore to	• New cluster / Existing cluster
Upgrade	
• Supported	• Minor version
• Upgrade type	• Rolling update
Configuration change	
• Dynamic Configuration Change	
Monitoring & Alert	
• Monitoring	• Operator metrics
	• Operands(<i>i.e.</i> , cluster) metrics
• Alert	• Alert by metrics information
• Event Notification	• Custom resource creation events
Scalability	
• Scale-out read replica	• Manual / Automatic
• Scale-in read replica	• Manual

Fujitsu Enterprise Postgres features supported

- Transparent Data Encryption
- Data Masking
- Dedicated Audit Log
- Vertical Clustered Index
- Global Meta Cache
- High-speed Data Load
- Java, ODBC, and .NET Framework integration
- Embedded SQL integration (C, COBOL)



Fujitsu Enterprise Postgres
for Kubernetes

Containerize • Modernize

Topics


Start small and grow large

Fujitsu Enterprise Postgres for Kubernetes starts quickly with minimal configuration and easily scales as your business grows. It also scales automatically as transactions grow.

Cluster deployment

Users can instruct the operator to build clusters that include containers and volumes with Fujitsu Enterprise Postgres installed, alongside the required network resources. Clusters can be created with one master server or with multiple replicas (one master and two replicas). Synchronous or asynchronous replication can be selected for the replica servers.

Users can deploy Pgpool-II for load balancing and connection pooling with a Pgpool-II container. Multiple Pgpool-II pods may be deployed in a single deployment to increase availability.



Highly available configuration deployed by the Operator


Configurable volumes per cluster

To improve performance, users may want to separate the volumes storing database files and WAL files. Similarly, one may want to use a dedicated volume for a new tablespace. The Operator gives the end user the flexibility to create a cluster with multiple persistent volumes and select a suitable storage class for them. For example, the user can create a cluster with a data volume, a WAL volume on a storage class backed up on SSD, and a log volume on a storage class backed up on HDD.

Deploying Pgpool-II and connecting to cluster from Operator

Users can deploy the Pgpool-II container and access the database via Pgpool-II to use load balancing and connection pooling features.

Multiple Pgpool-II containers can be deployed for load sharing and high availability. Users can request a Kubernetes service to distribute their work across multiple Pgpool-II containers.



High availability features

Automatic failover

If an error is detected in the container or pod of the master server, the cluster will perform an automatic failover by promoting one of the replicas to master, and then switching the database connection destination. The database connection will be broken, but you can reconnect by re-establishing a connection from the application.


Automatic recovery

If an error occurs on the master server and an automatic failover occurs, the pod or container of the failed old master server is automatically restarted and reincorporated into the cluster as a replica server. If a replica server fails, it automatically restarts and rejoins the cluster as a replica server.

Manual switchover

You can manually switch any replica server to master server. In this case, the original master server becomes a replica server.

The high availability and failover management of Fujitsu Enterprise Postgres is provided by Patroni. Both Patroni and Fujitsu Enterprise Postgres will be installed on the same container image. Patroni will initialize and start a database instance, and then acquire a lock on a shared resource - in our case, it is a [Kubernetes configmap](#). Whichever pod acquires the lock will become the master. When a subsequent server container starts, Patroni will initialize that pod as a replica with streaming replication.



Kubernetes configmap for election of master

Scheduling backup with the Operator

Fujitsu Enterprise Postgres for Kubernetes supports full backup and incremental backup. Up to 5 backup schedules can be set up, with their retention period also being configurable.

A backup container is deployed as a complement to each server pod. Backup is performed at the scheduled time set by the user. The backup container checks if the server in the pod is a master or replica, and performs backup only if it is the master. The volume storing backup and archived WAL files must be on a shared storage such as NFS, AWS S3, or Red Hat OpenShift Container Storage. Backup and WAL archiving is performed by [pgBackRest](#).

Point-in-time Recovery and backup restore from Operator


Two types of restore are supported: restore backup data to an existing cluster, and create a new cluster and restore backup data. The former retains the attributes of the cluster, such as IP address and name, while the latter is created from scratch.

The restore process deploys a restore container. The restore container performs the pgBackRest restore operation from the backup data to be restored to the master server of the cluster. After the data is restored to the master server, the cluster is created by synchronizing the data to the replica servers.

Disaster recovery

By using OSS (pgBackRest) functionality to store backup data in object storage, data can be migrated to a database cluster in a different OCP environment.

Even if it is difficult to operate in an OCP environment with a database cluster due to a disaster, it is possible to continue operating in a different OCP environment.



Configuration change

Parameter change

The user can change the Fujitsu Enterprise Postgres parameters.

- For parameters that take effect immediately:
The Operator will apply the changes to all database pods and reload the Fujitsu Enterprise Postgres server process automatically. The cluster will not suffer outage.
- For parameters that require restart of the Fujitsu Enterprise Postgres server to take effect:
The Operator will update the configuration files on all database pods. However, users must perform a manual restart of the database process on all the database pods using the *FEPAction* custom resource.

Resource change

The user can change the amount of CPU and memory resources allocated to server containers, backup containers, and Pgpool-II containers by changing the *FEPCluster* custom resource. The Operator will apply the changes to the Statefulset.

Changing resource allocation will not take effect immediately - the user must restart all pods for the new resource allocation to take effect.

Version upgrades

Minor version upgrade

New and patched Fujitsu Enterprise Postgres releases are made available as new container images. When the latest container image is provided, the user can perform a minor version upgrade by changing the *FEPCluster* custom resource. The operator will perform a rolling update to enable the minor version upgrade with minimal system disruption.

A minor upgrade is done by updating the custom resource with a new Fujitsu Enterprise Postgres image name. The pod will be redeployed with the new image in a controlled manner.

- First, the replica servers is upgraded and restarted, and the system waits until they are ready, one server at a time.
- Then, a controlled switchover is performed to pick a new master.
- Finally, the old master will be upgraded.

Minor version upgrades will take effect immediately.

Major version upgrade

Upgrading the major version of the Operator and container builds a new major version of Fujitsu Enterprise Postgres in the same namespace as the old major version.

- Preparation** - Stop running applications and allow the upgrade execution container to connect.
- Operator upgrade** - Uninstall the old Operator and install the new version.
- Major version upgrade** - Create a new version of *FEPCluster* and migrate the data to it.
- Custom resource update** - Update custom resources as needed
- Resume applications** - Resume the stopped applications.

Monitoring & Alerts

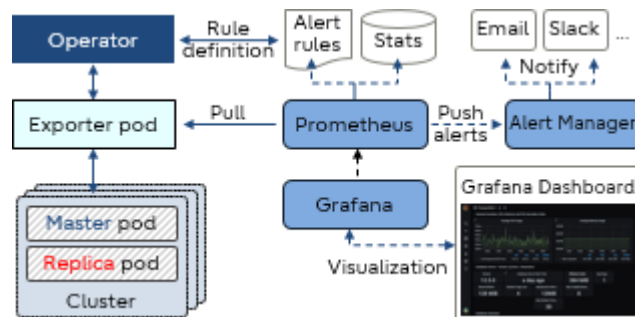
Monitoring

Monitoring is available from [Prometheus](#) and [Grafana](#), the de facto standard monitoring tools in Kubernetes.

The monitoring software Prometheus captures information about the cluster pod CPU, memory, and disk usage and health status. DBAs and Infrastructure Administrators can view monitoring data captured by Grafana features in an advanced graphical display.

Key monitoring items include:

- Database health
- OS performance information
- Disk usage
- Backup status
- Client connection information




Deep insights

Early detection of signs of failure is important. Prometheus sends alerts via the Alert Manager to DBAs and Infrastructure Administrators via email or Slack if the collected data (metrics) is unhealthy. Users can identify problems early and prevent problems before they occur.

Use case

- A Fujitsu Enterprise Postgres server container/pod CPU usage is exceeding 80% of the resource limits.
- A Fujitsu Enterprise Postgres Persistent Volume Claim has less than 10% disk available.
- A Fujitsu Enterprise Postgres server apparently went down or is not accessible.



Scaling replicas


Automatic scale out

Auto scaling enables the user to dynamically extend a database cluster according to a pre-defined scaling policy.

The user can pre-configure server load thresholds, so that the number of cluster replica pods is automatically increased when the threshold is reached. This allows the system to operate stably even when the system experience heavy access.

Use case

- Add one pod if CPU usage is more than 80%.
- Reduce the number of pods during off-season with low data volume.



Manual scale in/out

Users can scale out or scale in read replicas at any time. This can be done by applying changes to the `FEPCluster` custom resource.

DBAs and Infrastructure Administrators can manually increase or decrease the number of cluster pods to keep the system running at the minimum required number.

Supported Fujitsu Enterprise Postgres features

Category	Feature
Operation	Global Meta Cache 🔗
Security	Transparent Data Encryption 🔗
	Data Masking 🔗
	Dedicated Audit Log 🔗
High performance	Vertical Clustered Index 🔗
	High-speed Data Load 🔗
Application interface	Java integration 🔗
	ODBC integration 🔗
	.NET Framework integration 🔗
	Embedded SQL integration (C language) 🔗
	Embedded SQL integration (COBOL) 🔗

Supported OSS

OSS name	Version	Description
orafce	3.17.0	Oracle-compatible SQL features
patroni	2.1.2	Postgres cluster management
pg_bigm	1.2	Full-text search (multibyte)
pg_dbms_stats	1.5.0	Tuning (statistics management, query tuning)
pg_hint_plan	13.1.3.7	
pg_repack	1.4.7	Table reorganization
pg_rman	1.3.13	Backup and restore management
pg_statsinfo	13.0	Collection and accumulation of statistics
pgBackRest	2.36	Backup and restore management
pgBadger	11.6	Log analysis
Pgpool-II	4.2.8 (x86, s390x) 4.2.6 (ppc64le)	Failover, connection pooling, load balancing, etc.
postgres-exporter	0.9.0	PostgreSQL metrics monitoring capabilities for Prometheus with Fujitsu updated queries
PostgreSQL	14.9 (x86, s390x) 14.7 (ppc64le)	Database management system
PostgreSQL JDBC driver	42.2.23	JDBC driver
psqlODBC	13.02.0000	ODBC driver

Fujitsu Enterprise Postgres features not supported

Since Fujitsu Enterprise Postgres server containers are based on other components (like UBI and Patroni), there are certain limitations that do not allow it to be 100% functionally compatible with VM-based Fujitsu Enterprise Postgres server instances.

Limitation	Reason for limitation	Description
No support for JIT.	UBI8 does not provide required LLVM libraries.	It is not possible to enable JIT in postgresql.conf.
Impact on Fujitsu Enterprise Postgres parallelism improvement.	UBI8 does not host dstat binaries.	Fujitsu Enterprise Postgres improves parallelism by restricting the number of parallel workers when the CPU load is high. However, this is unlikely to have too much impact on a Fujitsu Enterprise Postgres container, since it runs in only one process.
CryptoExpress cards are not supported.	IBM LinuxOne does not currently support CryptoExpress cards in the OpenShift Container Platform.	The Fujitsu Enterprise Postgres TDEz extension cannot be used on LinuxOne OpenShift environment. However, the user can still use Transparent Data Encryption on both LinuxOne OpenShift as well as Azure (x86) OpenShift environments.
No support for Oracle foreign data wrapper.	The Oracle foreign data wrapper has a dependency on Instant Client package, which is not available.	The Oracle Instant Client package is not redistributed by Fujitsu Enterprise Postgres due to this limitation.

Fact sheet

Product name

Fujitsu Enterprise Postgres 14 for Kubernetes

Product information

Supported platforms

FEP Operator is tested on the following platforms.

Service	Platform
Self-Managed Kubernetes Service	<ul style="list-style-type: none"> Red Hat OpenShift Container Platform 4.12 - 4.14 (x86, s390x) , 4.11 - 4.13 (ppc64le) Rancher Kubernetes Engine (on Linux hosts) * SUSE Rancher 2.6
Full Managed Kubernetes Service	<ul style="list-style-type: none"> Red Hat OpenShift Service on AWS (ROSA) Red Hat OpenShift on IBM Cloud Azure Red Hat OpenShift (ARO) Azure Kubernetes Service (AKS) * Amazon Elastic Kubernetes Service (EKS) * IBM Cloud Kubernetes Service * Alibaba Cloud Container Service for Kubernetes (ACK) * Google Kubernetes Engine (GKE) *

* Kubernetes 1.25 - 1.27 (x86, s390x) , 1.24 - 1.26 (ppc64le)

Supported storage

Category	Storage
Type/interface	<ul style="list-style-type: none"> Container Storage Interface NFS Red Hat OpenShift Container Storage
Cloud Service	<ul style="list-style-type: none"> Azure Blob Storage Amazon S3

CPU

x86, s390x, ppc64le

Components embedded

Component	Version
Red Hat UBI minimal	8
Fujitsu Enterprise Postgres Server	14.10 (x86, s390x) 14.7 (ppc64le)

Collaboration tools

Integration with the monitoring and alerting tools below is supported:

Tool	Version	How to obtain
Alert Manager	Version installed with OpenShift	Pre-installed with OpenShift
Grafana	3.10.3 and later	Provided by OperatorHub
Prometheus	Version installed with OpenShift	Pre-installed with OpenShift

Further information

Information	Description	URL
Red Hat Ecosystem Catalog	Container catalog from Red Hat. Red Hat OpenShift Operator certified software is registered.	https://catalog.redhat.com/partners/Fujitsu
Red Hat Marketplace	Open cloud marketplace that facilitates the discovery and access of certified software for container-based environments.	https://marketplace.redhat.com/en-us/products/fujitsu-enterprise-postgres
Technical presentation	Technical details and capabilities of Fujitsu Enterprise Postgres for Kubernetes.	https://fast.fujitsu.com/FujitsuEnterprisePostgresForKubernetes-TechnicalPresentation
Webpage	Learn more about Fujitsu Enterprise Postgres for Kubernetes at our website.	https://fast.fujitsu.com/fujitsu-enterprise-postgres-for-kubernetes

fast.fujitsu.com/

Email: enterprisepostgresql@fujitsu.com

2024-03-28 WW EN



© Fujitsu Limited 2023. Fujitsu, the Fujitsu logo, and Fujitsu brand names are trademarks or registered trademarks of Fujitsu Limited in Japan and other countries. Other company, product and service names may be trademarks or registered trademarks of their respective owners. All rights reserved. No part of this document may be reproduced, stored, or transmitted in any form without prior written permission of Fujitsu Limited. Fujitsu Limited endeavors to ensure the information in this document is correct and fairly stated but does not accept liability for any errors or omissions.