


FUJITSU Enterprise Postgres 12 for Kubernetes



Overview

Preface

Purpose of this document

This document explains the FUJITSU Enterprise Postgres for Kubernetes concepts to those who are to operate databases using it.

This document explains the features of FUJITSU Enterprise Postgres for Kubernetes.

Intended readers

This document is intended for people who are:

- Considering installing FUJITSU Enterprise Postgres for Kubernetes
- Using FUJITSU Enterprise Postgres for Kubernetes for the first time
- Wanting to learn about the concept of FUJITSU Enterprise Postgres for Kubernetes
- Wanting to see a functional overview of FUJITSU Enterprise Postgres for Kubernetes

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

- Linux
- Kubernetes
- Containers
- Operators

Structure of this document

This document is structured as follows:

[Chapter 1 Know about the Product](#)

Explains the features of FUJITSU Enterprise Postgres for Kubernetes.

[Chapter 2 Know What it does](#)

Explains what you need to do.

[Appendix A OSS Supported by FUJITSU Enterprise Postgres for Kubernetes](#)

Explains the OSS supported by FUJITSU Enterprise Postgres for Kubernetes.

Abbreviations

The following abbreviations are used in this manual:

Full Name	Abbreviations
FUJITSU Software Enterprise Postgres for Kubernetes FUJITSU Software Enterprise Postgres	FEP or FUJITSU Enterprise Postgres
Custom Resource	CR
Custom Resource Definition	CRD
Persistent Volume	PV

Abbreviations of manual titles

The following abbreviations are used in this manual as manual titles:

Full Manual Title	Abbreviations
FUJITSU Software Enterprise Postgres for Kubernetes User's Guide	User's Guide

Trademarks

- Linux is a registered trademark or trademark of Mr. Linus Torvalds in the U.S. and other countries.
- Red Hat and all Red Hat-based trademarks and logos are trademarks or registered trademarks of Red Hat, Inc. in the United States and other countries.
- S/390 is a registered trademark of International Business Machines Corporation ("IBM") in the U.S. and other countries.

Other product and company names mentioned in this manual are the trademarks or registered trademarks of their respective owners.

Export restrictions

If this document is to be exported or provided overseas, confirm legal requirements for the Foreign Exchange and Foreign Trade Act as well as other laws and regulations, including U.S. Export Administration Regulations, and follow the required procedures.

Issue date and version

Edition 6.0: October 2022
Edition 5.0: June 2022
Edition 4.0: March 2022
Edition 3.0: December 2021
Edition 2.0: April 2021
Edition 1.0: March 2021

Copyright

Copyright 2021-2022 FUJITSU LIMITED

Contents

Chapter 1 Know about the Product.....	1
1.1 What is FUJITSU Software Enterprise Postgres for Kubernetes?.....	1
1.2 Operator Features.....	1
1.2.1 Cluster Deployment.....	2
1.2.1.1 Creating a FEPCluster.....	2
1.2.1.2 Creating a FEP Pgpool2 Container.....	3
1.2.2 Highly Available Feature.....	3
1.2.2.1 Automatic Failover.....	3
1.2.2.2 Automatic Recovery.....	3
1.2.2.3 Manual Switchover.....	3
1.2.3 Backup Recovery.....	4
1.2.3.1 Automatic Backup.....	4
1.2.3.2 Point-in-time Recovery.....	4
1.2.4 Configuration Change.....	4
1.2.4.1 Parameter Change.....	4
1.2.4.2 Resource Change.....	4
1.2.5 Minor Version Upgrade.....	4
1.2.5.1 Minor Version Upgrade.....	4
1.2.6 FEP Features.....	5
1.2.6.1 Scope of FEP Feature Support.....	5
1.3 Operator System Configuration.....	5
Chapter 2 Know What it does.....	7
2.1 Deployment.....	7
2.2 High Availability (Automatic failover and recovery).....	7
2.3 Configuration Change.....	7
2.4 Minor Version Upgrade.....	8
2.5 Configurable Volume Per Cluster.....	8
2.6 Deploying Pgpool-II and Connect to FEPCluster from Operator.....	8
2.7 Scheduling Backup from Operator.....	8
2.8 Perform PITR and Latest Backup Restore from Operator.....	9
Appendix A OSS Supported by FUJITSU Enterprise Postgres for Kubernetes.....	11

Chapter 1 Know about the Product

This chapter explains the features of FUJITSU Enterprise Postgres for Kubernetes.

1.1 What is FUJITSU Software Enterprise Postgres for Kubernetes?

FUJITSU Software Enterprise Postgres for Kubernetes provides an enterprise platform for installing and managing your FUJITSU Enterprise Postgres 12 on OpenShift Container Platform.

There are multiple components in the solution.

FEP operator: Manages the lifecycle of FEP server container, including deployment, configuration update, backup and recovery of FEP database.

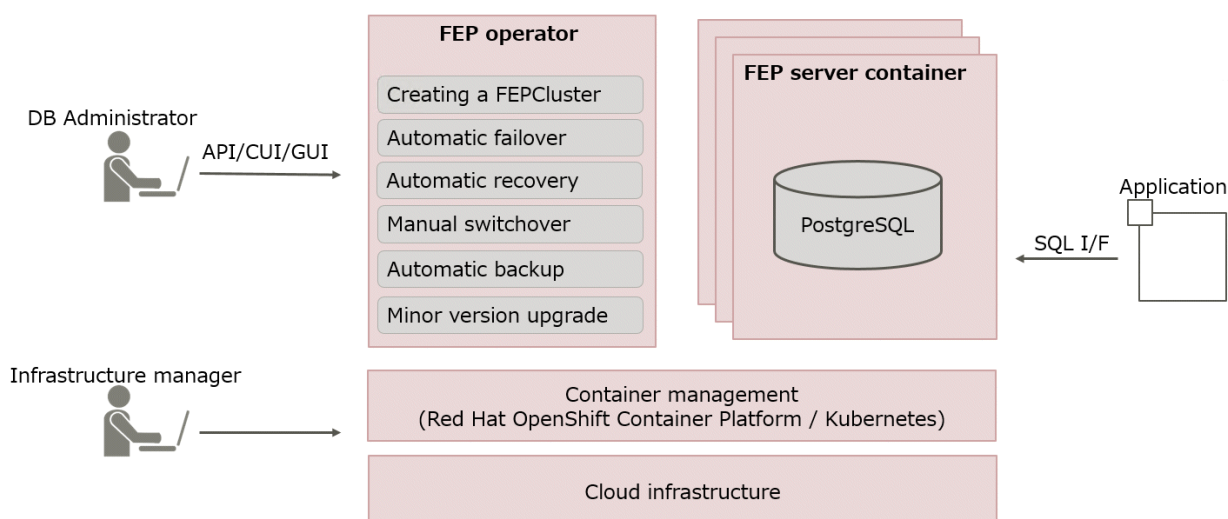
FEP server container: Contains the FEP server software to run the Postgres engine.

FEP backup container: Contains the FEP server software to perform scheduled backup operations.

FEP restore container: Contains the FEP server software to perform the restore operation.

FEP pgpool2 container: Contains the FEP server software to use Pgpool-II to provide load balancing and connection pooling.

Up and running in minutes, the operator provides the features required to maximise the benefits of this enterprise PostgreSQL solution.



This operator will deploy a standalone as well as highly available FUJITSU Enterprise Postgres cluster with pre-defined configuration to get started with small workload. User can adjust the configuration parameters at the time of deployment and after to make the instance suitable for the workload.

As the name implies, the FEP server container is intended to incorporate the FUJITSU Enterprise Postgres server component.

In principle, a running FEP server container is considered as equivalent to a FUJITSU Enterprise Postgres Server instance.

1.2 Operator Features

This product provides operator services to automate the construction and operation of databases on the customer's container management infrastructure. The features of the operator are as follows:

- Cluster Deployment
 - [Creating a FEPCluster](#)
 - [Creating a FEP Pgpool2 Container](#)

- Highly Available Feature
 - [Automatic Failover](#)
 - [Automatic Recovery](#)
 - [Manual Switchover](#)
- Backup Recovery
 - [Automatic Backup](#)
 - [Point-in-time Recovery](#)
- Configuration Change
 - [Parameter Change](#)
 - [Resource Change](#)
- [Minor Version Upgrade](#)
- [FEP Features](#)

1.2.1 Cluster Deployment

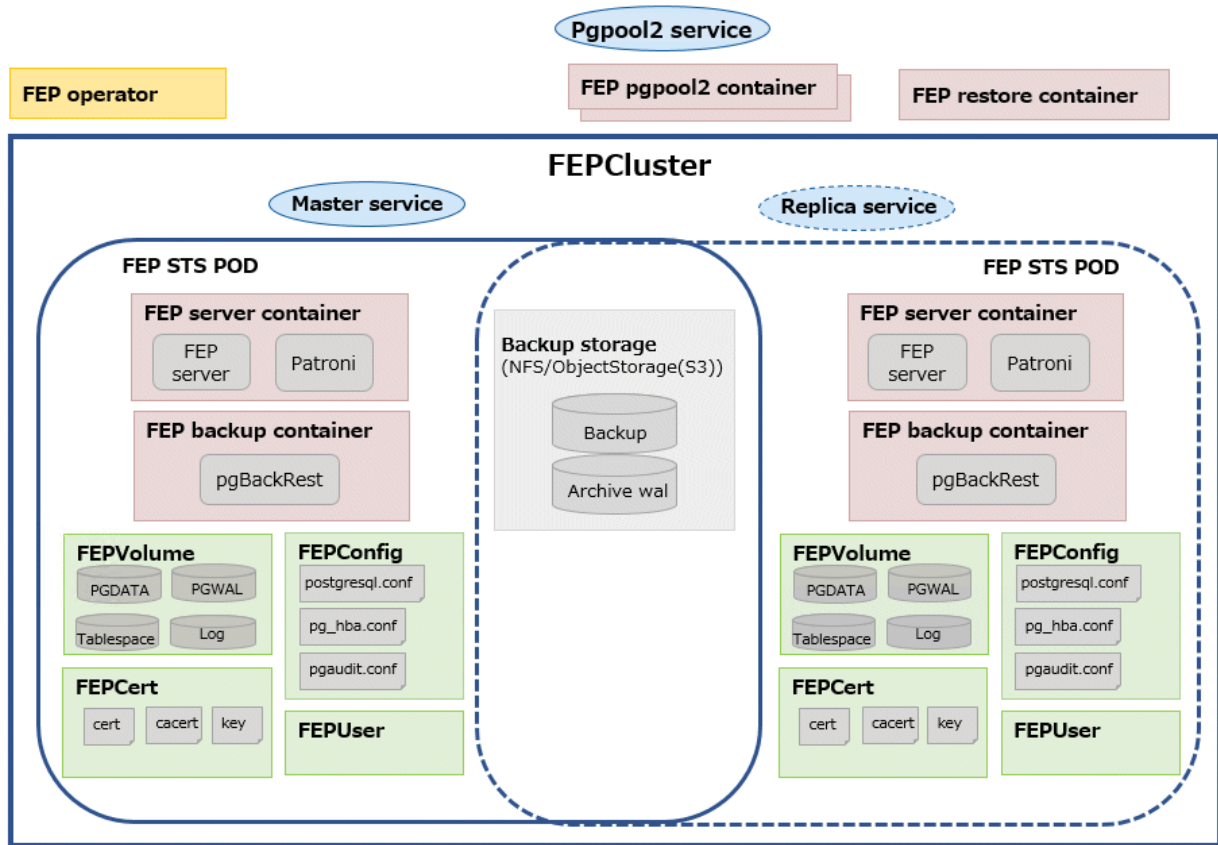
1.2.1.1 Creating a FEPCluster

Users can instruct the operator to build a system that includes the provisioning of containers and volumes with FEP installed, and network resources. The resulting system is called a FEPCluster. The FEPCluster can be created a single master server or multi-servers with one master and two replicas. You can choose between synchronous and asynchronous replication replica servers. The default is synchronous replication.

FEPCluster is composed of the following components:

- FEP server container
 - FEP server
 - Patroni
- FEP backup container
- CR FEPCVolume for volumes
- CR FEPCUser for database users
- CR FEPCConfig for Postgres configuration
- CR FEPCert for secrets such as TLS certificate, keystore passphrase

The Below diagram depicts a FEPCluster with one Master and one Replica POD.



1.2.1.2 Creating a FEP Pgpool2 Container

Users can deploy Pgpool-II for load balancing and connection pooling with FEP pgpool2 container.

Users can deploy multiple FEP pgpool2 Pods in a single deployment to increase availability.

1.2.2 Highly Available Feature

1.2.2.1 Automatic Failover

When 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 become the new master, and the connection destination of the database is switched. The database connection is broken, but you can reconnect by establishing a connection from the application again.

1.2.2.2 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.

1.2.2.3 Manual Switchover

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

1.2.3 Backup Recovery

1.2.3.1 Automatic Backup

Users can set an arbitrary schedule for automatic backup. The backup type can be a full backup or an incremental backup. You can back up the database to shared storage such as NFS persistent volume or AWS S3 compatible storage. Backups can be automatically deleted by setting a retention period of your choice.

1.2.3.2 Point-in-time Recovery

Allows point-in-time recovery from automated backup data to restore the cluster. You can choose between restoring data to an existing cluster and a new cluster. You can also choose to restore to the most recent data or to any time you specify.

1.2.4 Configuration Change

1.2.4.1 Parameter Change

You can change the parameters that make up the FEP. PostgreSQL provides two types of parameters: those that take effect immediately, and those that take effect after restarting FEP server process.

- postgresql.conf
- pg_hba.conf
- pgaudit.conf



For parameters that take effect immediately, operator will apply the change to all FEP Pods and reload the FEP server process automatically. There is no outage on the cluster.

For parameters that take effect after restarting FEP server process, operator will update the configuration files on all FEP Pods. However, users have to initiate a manual restart of FEP process on all the FEP Pods using the FEPAction CR. There is a momentary outage on the cluster and the users should perform this action at a time that has least disruption to the service.

1.2.4.2 Resource Change

You can change the amount of CPU and memory resources allocated to FEP server containers, FEP backup containers, or FEP pgpool2 containers by changing the FEPCluster CR. The operator will apply the change to the Statefulset. However, the users have to perform a manual recreation of statefulsets for the new resource allocation to take effect.



Changing resource allocation will not take effect immediately. The users have to recreate statefulsets for new resource allocation to take effect. There is a momentary outage on the cluster and the users should perform this action at a time that has least disruption to the service.

1.2.5 Minor Version Upgrade

1.2.5.1 Minor Version Upgrade

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

Note

The minor version upgrade will take effect immediately. There is a momentary outage on the cluster and the users should perform this action at a time that has least disruption to the service.

1.2.6 FEP Features

1.2.6.1 Scope of FEP Feature Support

Note

These features also require the FEP Client ISO.

The FEPCluster that is created supports the following features in addition to the PostgreSQL features of OSS. Details of each feature can be found in the FEP documentation.

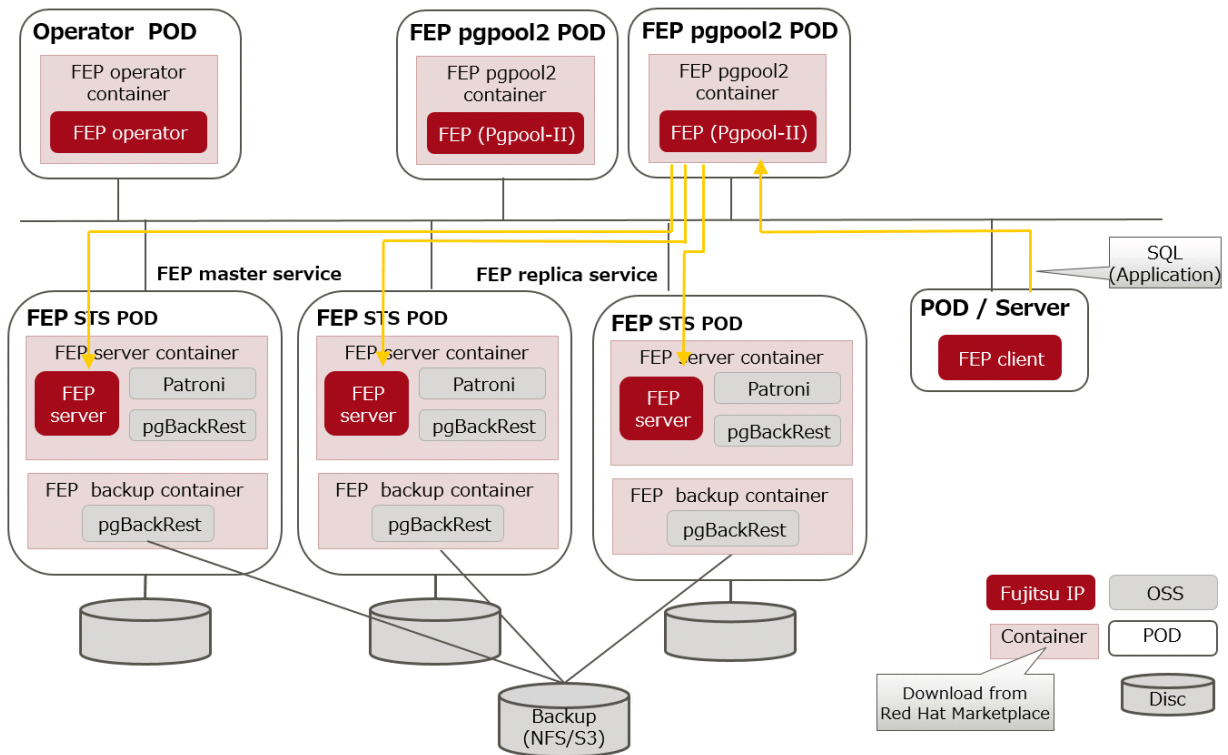
Category	Feature
Operation	pgAdmin
	Global Meta Cache
Data masking	Transparent Data Encryption
	Audit Log
	Secrecy
High Performance	In-memory feature
	High-speed data load
Application Interface	Java Integration
	ODBC Integration
	.NET Framework Integration
	Embedded SQL Integration (C language)
	Embedded SQL Integration (COBOL)

1.3 Operator System Configuration

The basic relationships among POD, containers and services are as follows.

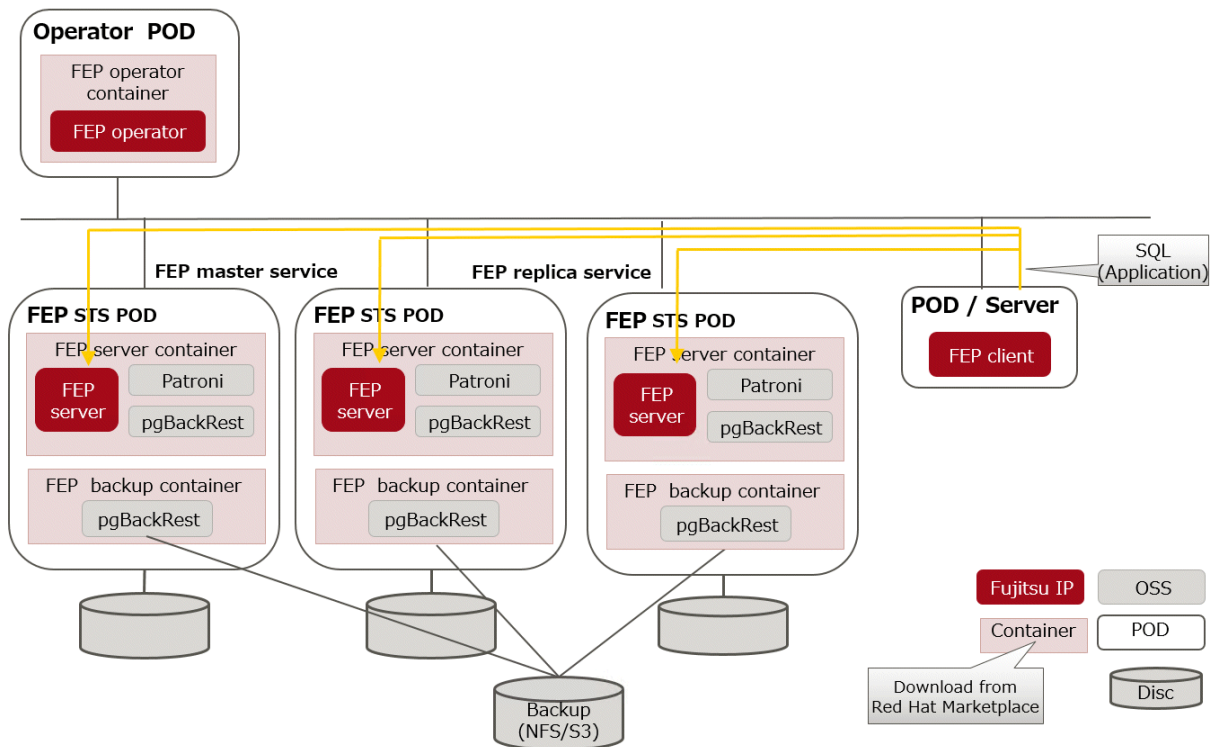
Example) Deployment with Pgpool-II

In this deployment scenario, Pgpool-II is used to provide connection pooling and load balancing. End user application will point its connection to Pgpool service. Depending on the transaction type, Pgpool will forward the connection to either the Master Pod or the Replica Pod. If a failover/switchover occurs, the FEP pgpool2 will direct traffic to the new FEP master Pod. This is transparent to the end user application.



Example) Deployment without Pgpool-II

Users can also run applications such as SQL directly against the FEPCluster without configuring Pgpool-II. In this deployment scenario, end user application will point its connection to the FEP master service. If a failover/switchover occurs, the FEP master service will point to the new FEP master Pod automatically. The end user application will experience a disconnection. When it re-establishes the connection, it will be connected to the new FEP master Pod. There is no need to reconfigure the application connection string.



Chapter 2 Know What it does

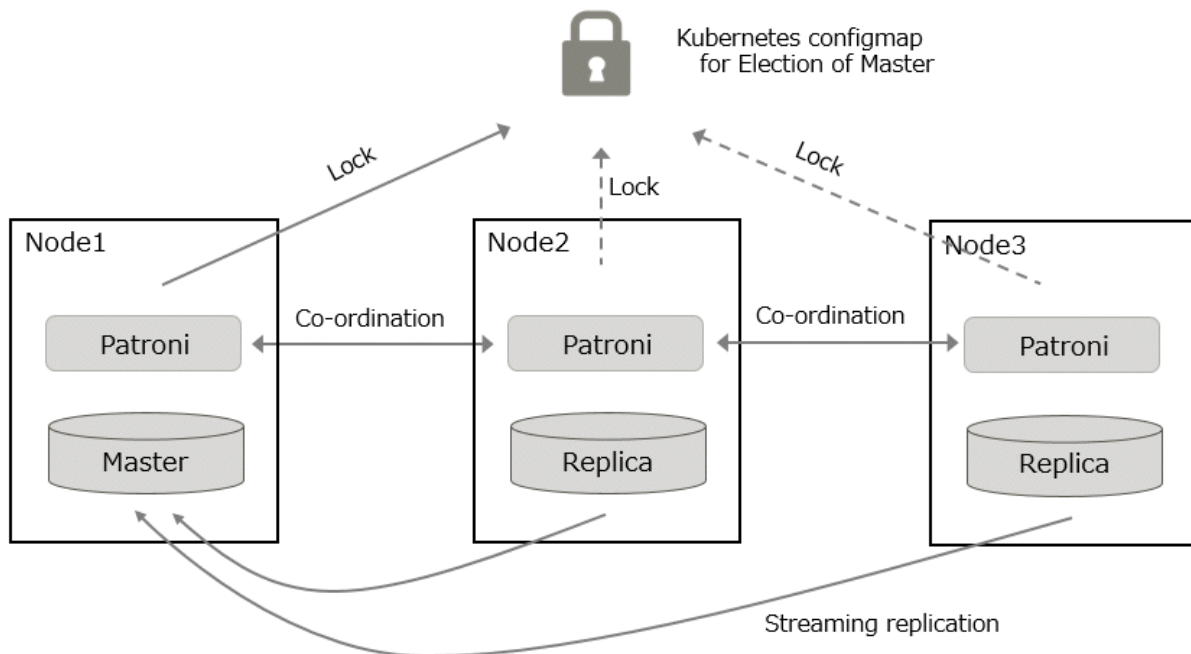
This chapter explains what you need to do.

2.1 Deployment

FEP operator is responsible for the lifecycle of FEPCluster. The operator will deploy a HA FEPCluster, together with all the associated containers such as backup container.

2.2 High Availability (Automatic failover and recovery)

The high availability and failover management of FEP is provided by Patroni. Both Patroni and FEP will be installed on the same container image. Patroni will then initialize and start an FEP instance. Patroni will then acquire a lock on a shared resource. In our case, it is a Kubernetes configmap. Whichever POD that can acquire the lock will become the Master. When subsequent FEP server container starts, Patroni will initialize that POD as a Replica with streaming replication.



If Patroni detects a failure in the cluster, either because the Postgres process crashed or the container where Postgres is running dies, Patroni will initiate a failover automatically.

2.3 Configuration Change

Traditionally, changing FEP configurations such as postgresql.conf, pg_hba.conf, TLS certificates and keystore passphrase will require a redeployment of FEP server container. That causes an outage in a Highly Available environment.

A new CRD FEPConfig is defined to encapsulate those configurations. The operator will monitor the CR with this CRD definition and perform action accordingly to minimize outages. For example, operator will reload FEP daemon, instead of redeploying the FEP server container when a reloadable postgresql.conf parameter is changed. If a parameter change requires restart of FEP (e.g. max_connections), the operator will update the configuration file but defer the restart. End user can follow a defined procedure to restart the cluster manually at a scheduled maintenance time.

2.4 Minor Version Upgrade

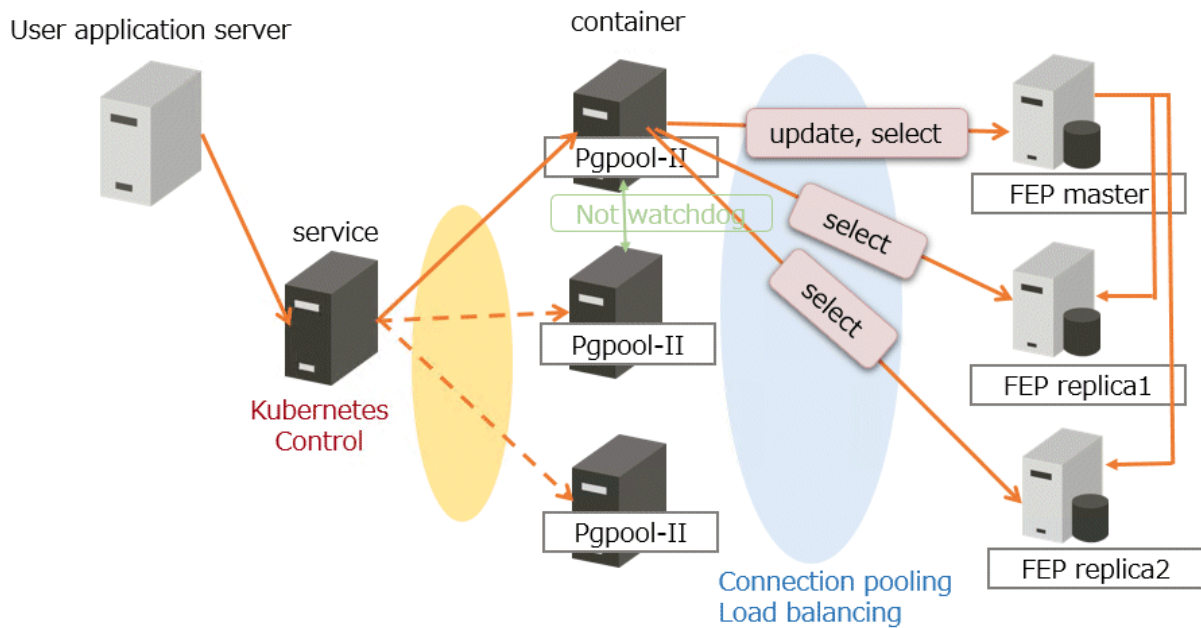
FEP version Minor upgrade is done by updating the Custom Resource with a new FEP image name. The POD will be redeployed with new image in a controlled manner. First, replica servers are upgraded, restarted and waited to be ready, one server at a time. When all replicas are upgraded, a controlled switchover is performed to pick a new master. Once that is done, the old master is upgraded as well.

2.5 Configurable Volume Per Cluster

To improve performance, may want to separate the volume 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 FEPCluster with multiple PVs and select a suitable storage class for the PV. For example, one can create a FEPCluster with data volume, wal volume on a storage class backed up by SSD and a log volume on a storage class backed up by HDD.

2.6 Deploying Pgpool-II and Connect to FEPCluster from Operator

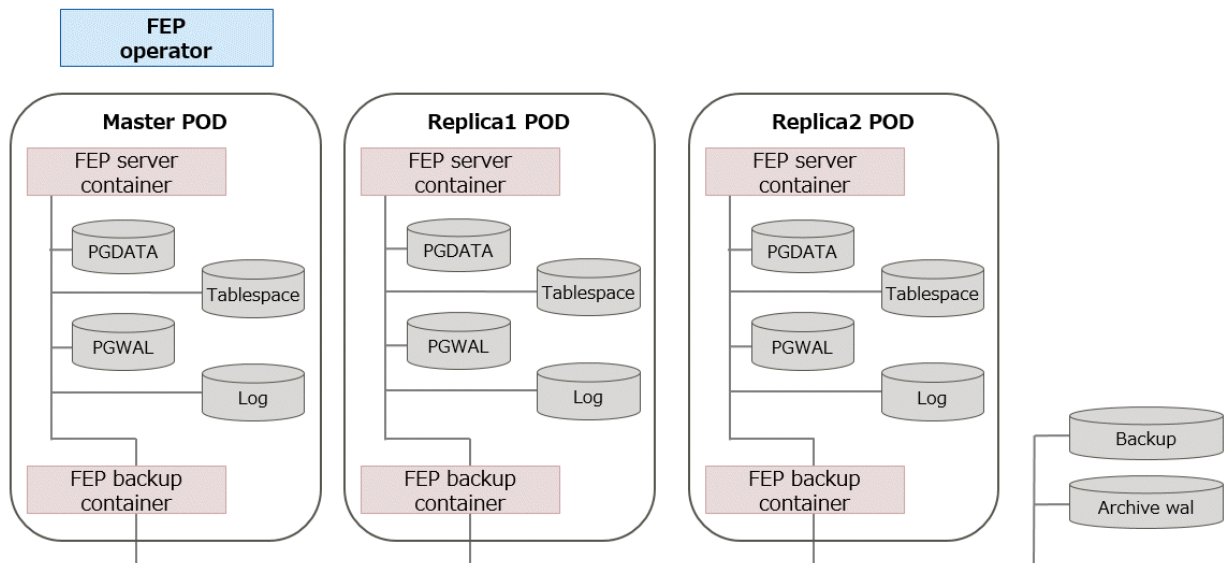
Users can deploy the FEP pgpool2 container and access the database via Pgpool-II to use load-balancing and connection pooling features. Multiple FEP pgpool2 containers can be deployed for load-share and high availability. Users can request a Kubernetes service to distribute their work across multiple FEP pgpool2 containers.



2.7 Scheduling Backup from Operator

The FEP backup container is deployed as a sidecar to each FEP server POD. The backup is performed at scheduled time set by the user (like crontab). The FEP backup container determines if the FEP server in the POD is a master or replica, and will perform the backup process only on the master POD. The volume storing backup and archived WAL files must be on a shared storage such as NFS or AWS S3.

Backup and WAL archiving is accomplished with pgBackRest.



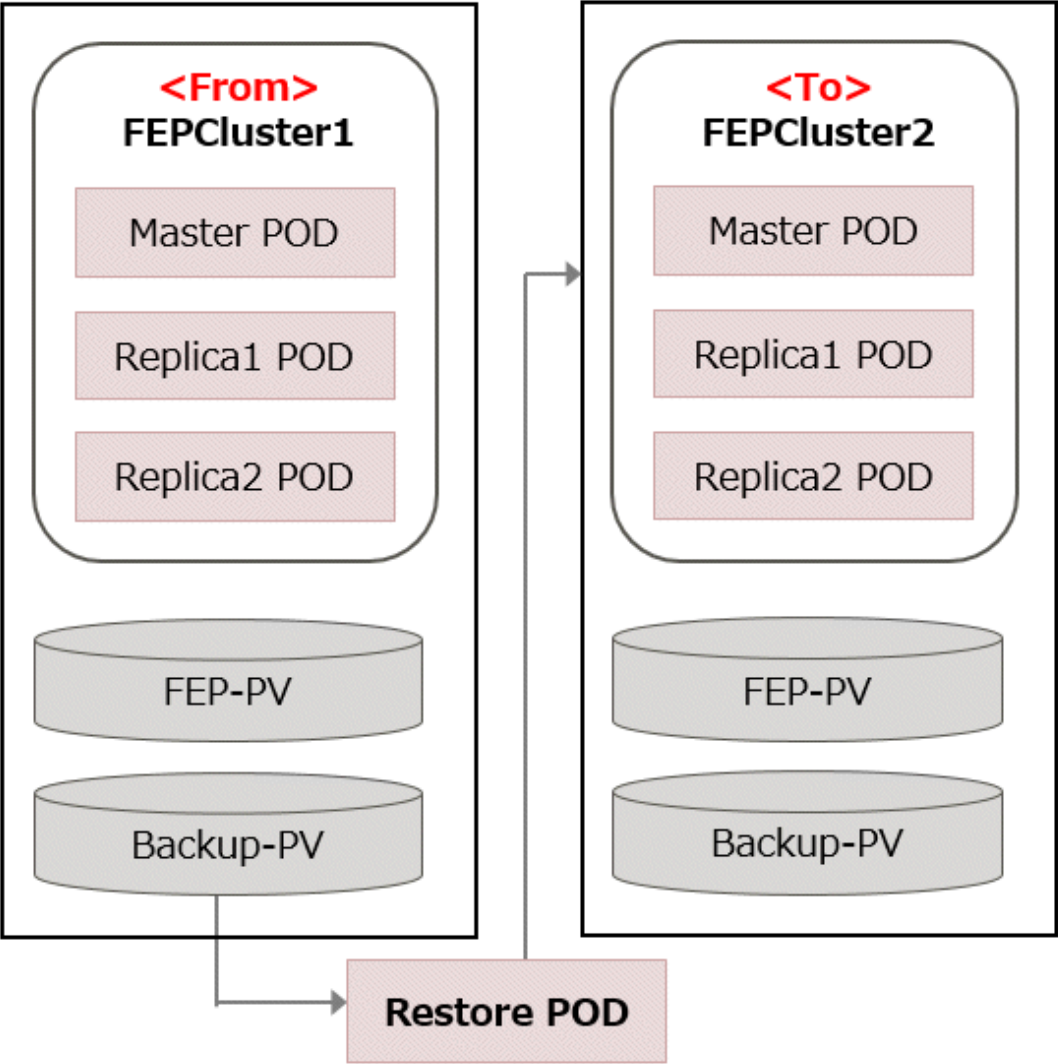
2.8 Perform PITR and Latest Backup Restore from Operator

There are two types of restore: one is to restore backup data to an existing FEP cluster, and the other is to create a new FEP cluster and restore backup data.

The former retains the attributes of the FEP 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 FEP cluster. After the data is restored to the master server, the FEP cluster is created by synchronizing the data to two replica servers.

FEP operator



Appendix A OSS Supported by FUJITSU Enterprise Postgres for Kubernetes

The OSS supported by FUJITSU Enterprise Postgres for Kubernetes is listed below.

OSS name	Version and level	Description	Reference
PostgreSQL	12.12	Database management system	PostgreSQL Documentation
orafce	3.18.1	Oracle-compatible SQL features	"Compatibility with Oracle Databases" in the FUJITSU Enterprise Postgres Application Development Guide
Pgpool-II	4.1.11	Failover, connection pooling, load balancing, etc.	"Pgpool-II" in the FUJITSU Enterprise Postgres Installation and Setup Guide for Server
pg_statsinfo	12.0	Collection and accumulation of statistics	"pg_statsinfo" in the FUJITSU Enterprise Postgres Installation and Setup Guide for Server
pg_hint_plan	12.1.3.7	Tuning (statistics management, query tuning)	- "pg_hint_plan" in the FUJITSU Enterprise Postgres Installation and Setup Guide for Server - "Optimizer Hints" in the FUJITSU Enterprise Postgres Application Development Guide
pg_dbms_stats	1.3.11		- "pg_dbms_stats" in the FUJITSU Enterprise Postgres Installation and Setup Guide for Server - "Locked Statistics" in the FUJITSU Enterprise Postgres Application Development Guide
pg_repack	1.4.6	Table reorganization	"pg_repack" in the FUJITSU Enterprise Postgres Installation and Setup Guide for Server
pg_rman	1.3.13	Backup and restore management	"pg_rman" in the FUJITSU Enterprise Postgres Installation and Setup Guide for Server
pgBadger	11.7	Log analysis	"pgBadger" in the FUJITSU Enterprise Postgres Installation and Setup Guide for Server
pg_bigm	1.2	Full-text search (multibyte)	"pg_bigm" in the FUJITSU Enterprise Postgres Installation and Setup Guide for Server
PostgreSQL JDBC driver	42.2.8	JDBC driver	"JDBC Driver" in the FUJITSU Enterprise Postgres Application Development Guide
psqlODBC	12.01.0000	ODBC driver	"ODBC Driver" in the FUJITSU Enterprise Postgres Application Development Guide
pgBackRest	2.30	Backup and restore management	"Scheduling Backup from Operator" in the User's Guide
patroni	2.0.2	Postgres cluster management	"High Availability" in the User's Guide
Python	3.7.5	PL/Python	