

Fujitsu Enterprise Postgres for Kubernetes

Quick Start Guide

J2UL-QK17-02ENZ0(00)

1. Prerequisites

- Registered OpenShift cluster with Red Hat Marketplace
cf. <https://marketplace.redhat.com/en-us/documentation/clusters#register-openshift-cluster-with-red-hat-marketplace>
- Buy or try the product 'Fujitsu Enterprise Postgres for Kubernetes ' from Red Hat Marketplace
cf. <https://marketplace.redhat.com/en-us/documentation/operators>

2. System Requirements

2.1. CPU

The following CPU architectures are supported.

No	CPU architecture
1	x86
2	s390x
3	ppc64le

2.2. Supported Platform

The following platform is supported.

No	Platform	Version
1	OpenShift Container Platform	4.14, 4.15, 4.16, 4.17

3. Operator installation from Red Hat Marketplace

1. For information on registering your cluster and creating a namespace, see [Red Hat Marketplace Docs](#). This must be done prior to installing the operator.
2. On the main menu, click **Workspace**, click **Software**, click on the **product box** of 'Fujitsu Enterprise Postgres for Kubernetes ', and then click Install Operator.
3. On the *Update Channel section*, select an option.
4. On the *Approval Strategy section*, select either *Automatic or Manual*. The approval strategy corresponds to how you want to process operator upgrades.
5. On the *Target Cluster section*:

- Click the checkbox next to the clusters where you want to install the Operator.
 - For each cluster you selected, under **Namespace Scope**, on the **Select Scope** list, select an option.
6. Click **Install**. It may take several minutes for installation to complete.
 7. Once installation is complete, the status will change from **installing** to **Up to date**.
 8. For further information, see the [Red Hat Marketplace Operator documentation](#)

Install Operator

Update channel

Operators are organized into packages and streams of updates called "channels". If an operator is available through multiple channels, you can choose which one you want to subscribe to. [Learn more](#)

stable

Approval strategy

Automatic updates keep the operator and any instances on the cluster up to date. Manual updates require approval and are done via OpenShift console or CLI. [Learn more](#)

Automatic
 Manual

Target clusters

Choose clusters where you want to install and manage this operator. Then select the Namespace scope for each cluster you are installing into. [Learn more](#)

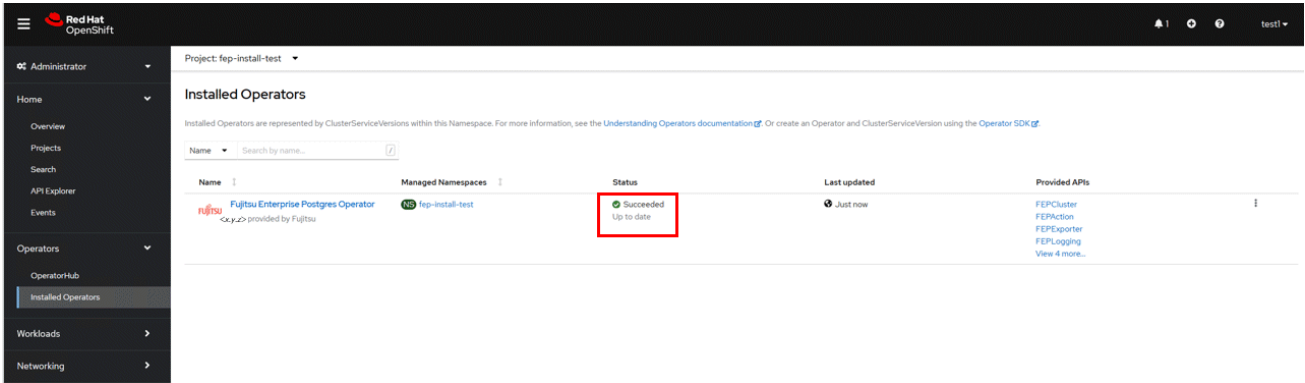
<input type="checkbox"/>	Name	Platform	Namespace Scope
<input type="checkbox"/>	rj-rhm-amy-test	IBM Cloud	Select Scope ▼

[Cancel](#)

4. Verification of operator installation

1. Once status changes to Up to date, click the vertical ellipsis and select Cluster Console.
2. Open the cluster where you installed the product
3. Go to **Operators > Installed Operators**
4. Select the Namespace or Project you installed on
5. Verify status for product is **Succeeded**

"<x.y.z>" in the screen example indicates the version level of the operator.



Installed operators status changes to "Succeeded"

5. Deploying FEPCluster using Operator

CR templates are published in the following repository on GitHub.

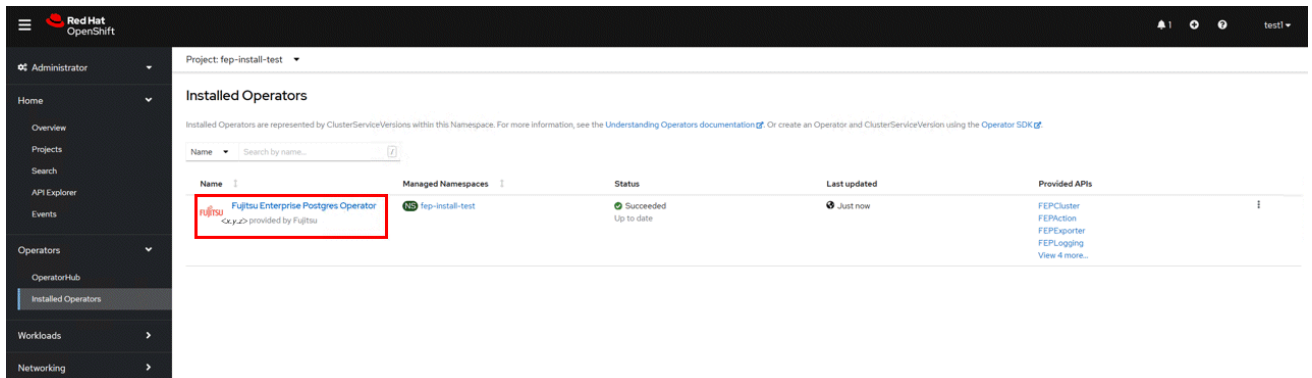
<https://github.com/fujitsu/fep-operator-examples>

This repository provides sample files that you can use to run Fujitsu Enterprise Postgres Operator.

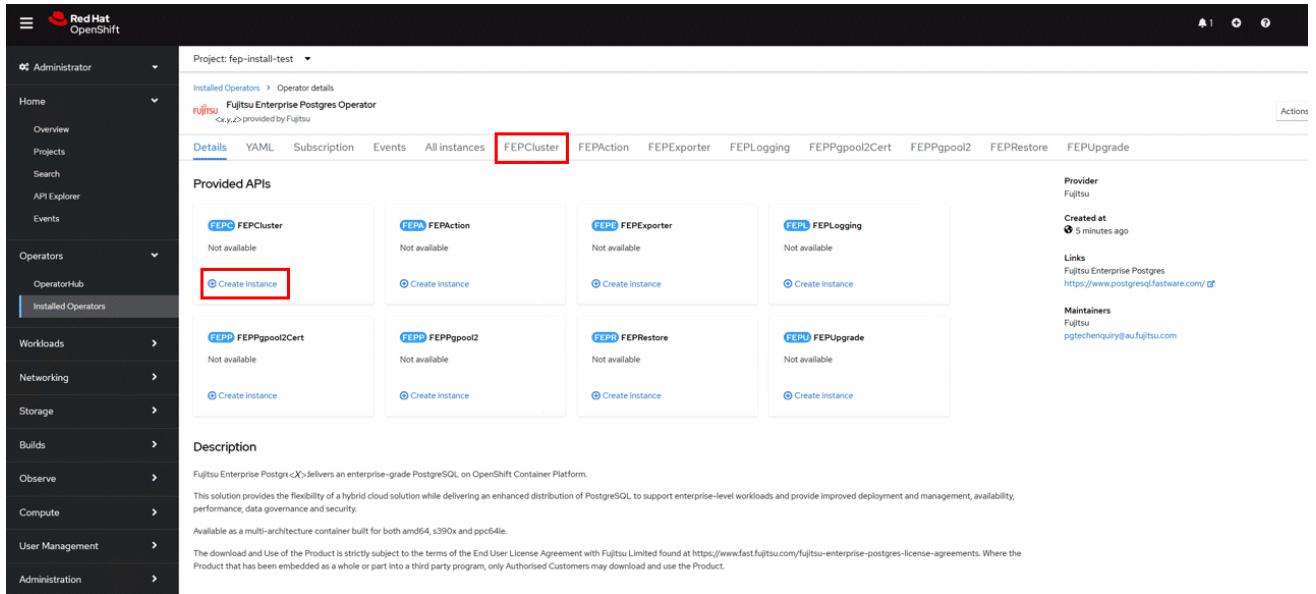
To deploy a FEPCluster in a given namespace, follow these steps:

"<x.y.z>" in the screen example indicates the version level of the operator. "<X>" indicates the product version of Fujitsu Enterprise Postgres.

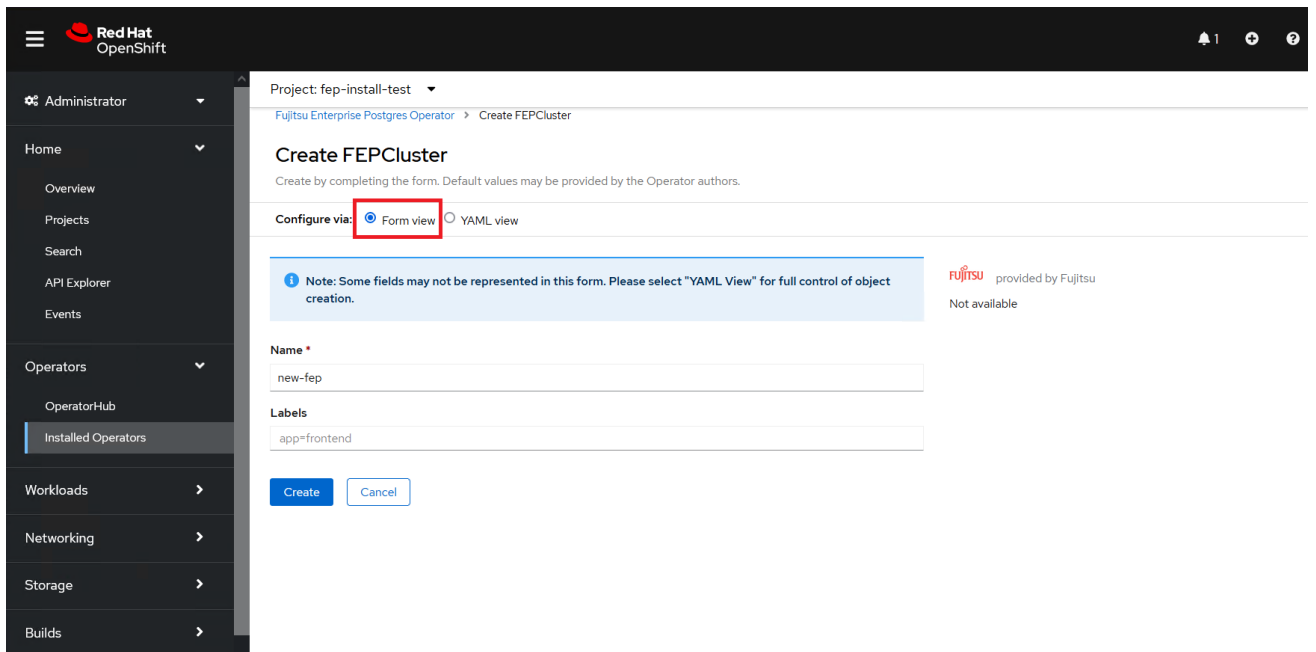
1. Under "Operators" menu item, click on "Installed Operators". You will see the installed FEP operator. Click on the name of operator.



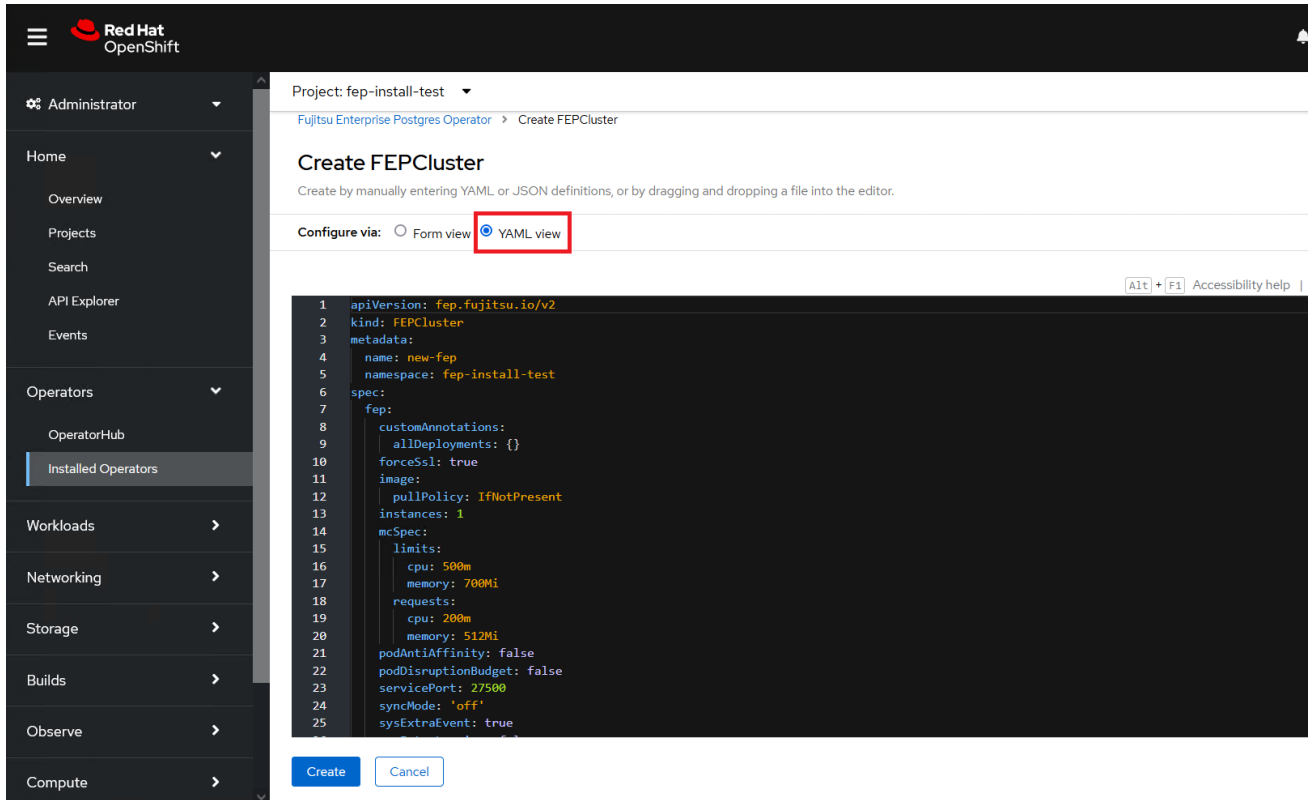
2. A page with all CRs that this operator supports will be displayed. FEPCluster is the main CR and all others are child CRs. We would create main CR and all other CRs will be created automatically by Operator.
To create Cluster CR, either
(1) Click on "Create Instance" under FEPCluster.
OR
(2) Click on "FEPCluster" on top and then click on "Create FEPCluster" on next page.



- This will bring to "Create FEPCluster" page. Here you have two options to configure. The first one is Form View. At the moment, in Form View, one can change only the name of cluster being deployed. Default name is "new-fep". This name must be unique within a namespace.



- In YAML View, starting value of CR is visible and one can choose to modify parameters before creating CR. Refer to the [Reference](#) for details of parameters. For example, add a configuration value for the customPgHba parameter according to your environment.



The FEPCluster custom resource allows you to define the container's CPU, Memory, disk size, etc. You can define each resource size individually, or you can use the following parameters to define the allocations for each resource in bulk.

Parameter	Description
spec.fep.databaseSize	Small, medium, and large define the following values for cpu/memory: small: 500m/700Mi medium: 2/4Gi large: 4/16Gi
spec.fepChildCrVal.storage.dataSize	Specifies the size of the data storage PV. Estimate and define the size of the backup storage area when the backup is enabled

We recommend that you use transparent data encryption to store your data. Tablespace PVs are mounted in/database/tablespaces/tbpace1. After building the database cluster, create tablespaces and tables as follows:

```
# Creating an Encrypted Tablespace
CREATE TABLESPACE secure_tablespace LOCATION '/database/tablespaces/tbpace1' WITH (tablespace_encryption_algorithm = 'AES256');

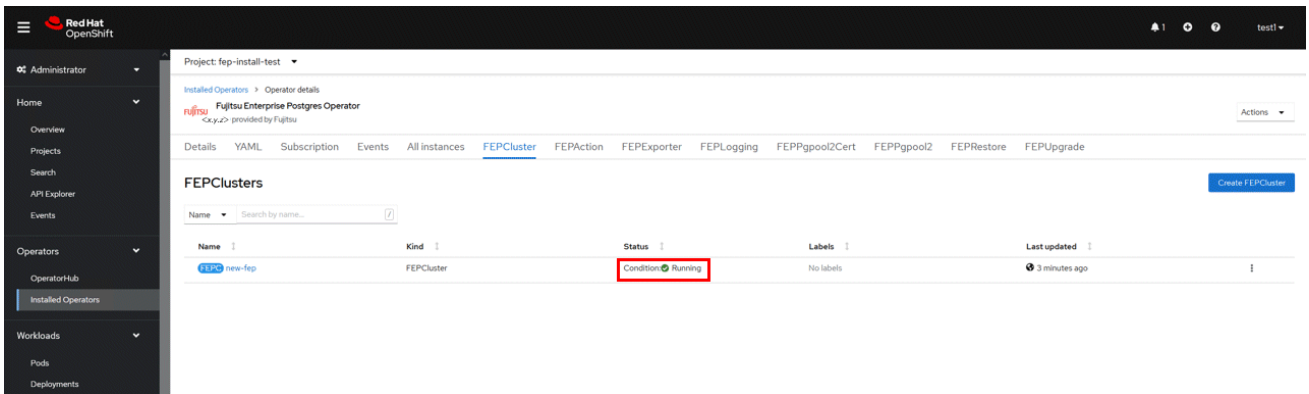
# Create Table in Encrypted Tablespace
CREATE TABLE secure_table (id int, pref text, city text, data text) TABLESPACE secure_tablespace;
```

Operator provides features such as backup, audit logging, and monitoring. These features can be enabled with the following parameters:

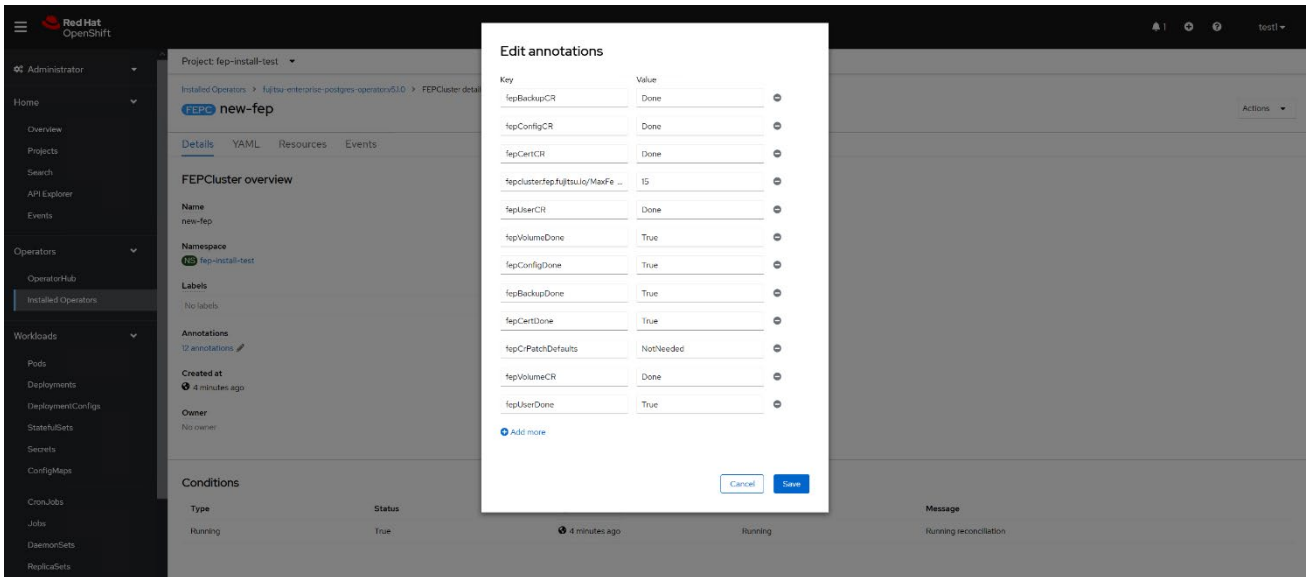
Parameter	Description
spec.fep.monitoring.enable	When true, monitoring feature is enabled.
spec.fep.pgAuditLog.enable	When true, audit log collection is enabled.
spec.fepChildCrVal.backup.type	local enables backup. The backup data is stored in the PV.
spec.fepChildCrVal.autoscale.scaleout.policy	cpu_utilization enables the ability to autoscale out replicas when CPU utilization exceeds a threshold.

Refer to the User's Guide for detailed settings and usage information for each function.

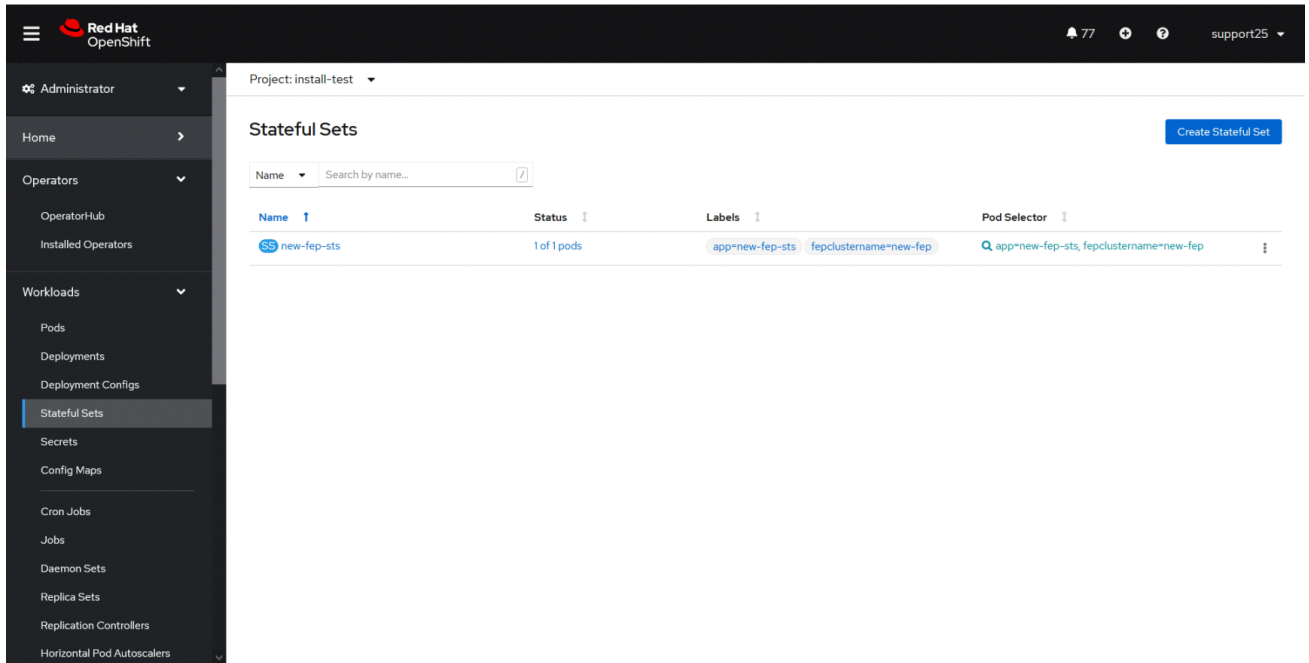
- When "Create" is clicked on either of two pages above, operator creates FEPCluster CR and there after one by one FEPBackup, FEPConfig, FEPVolume, FEPUser and FEPert child CRs are created automatically. The starting values for child CRs are taken from "fepChildCrVal" section of FEPCluster CR yaml file. Once child CRs are created, respective values are managed through child CRs only. If you want to change the value, modify the value in FEPCluster "fepChildCrVal" section. Operator reflects changes from FEPCluster parent CR to respective child CRs. Only allowable changes are reflected in child CRs. Child CRs are marked internal objects and hence will not be visible on OCP console. However, you can check child CRs using command line tools.



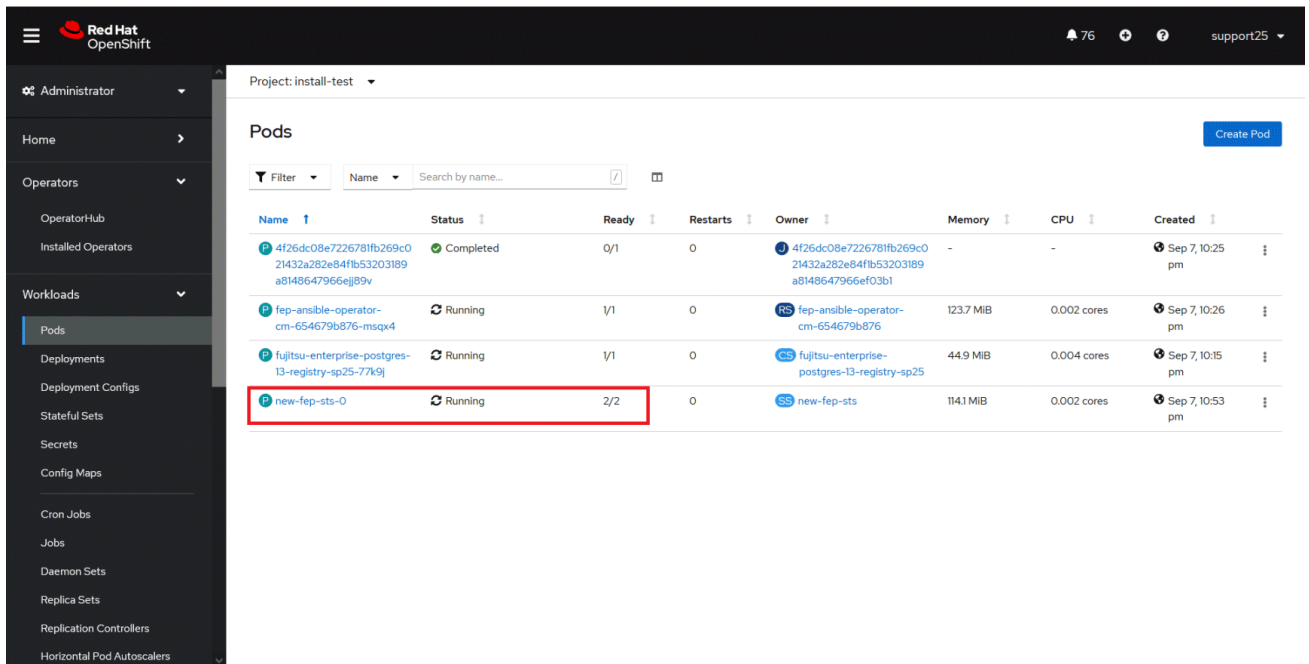
- In FEPCluster CR, annotations are added to indicate that child CRs are created successfully and have initialized properly. It may take some time to complete.



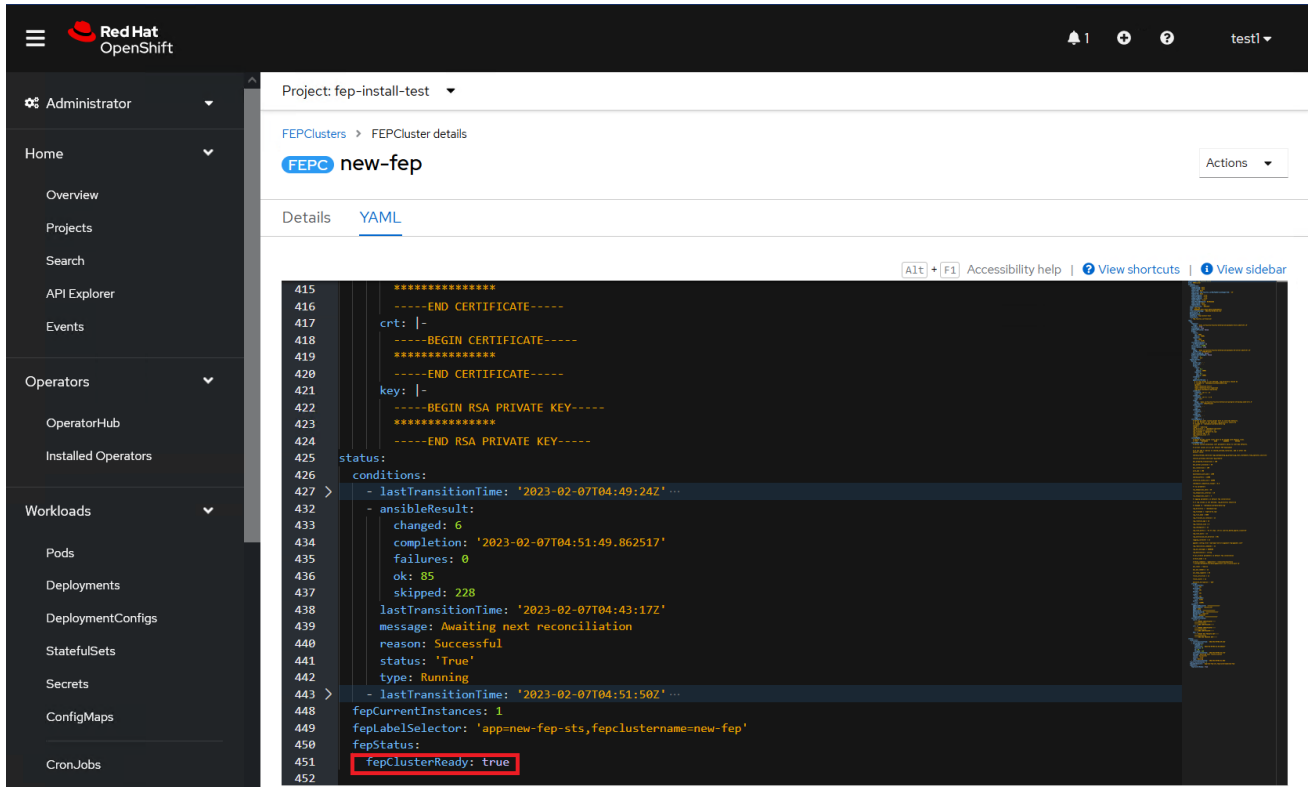
- Once all four child CRs are marked done in annotations, operator creates StatefulSet for the cluster.



8. StatefulSet will start one FEP instance at a time and will wait for each to be ready before starting next one.



9. Once all instances of FEP servers are started, operator marks a flag "fepClusterReady" in "fepStatus" section of CR to be **true**, indicating that FEPCluster is ready for use. Looking at YAML of FEPCluster CR, it would look like as below:



10. Operator also masks the sensitive fields like passwords, passphrase, certificates and keys in FEPC cluster `fepChildCrVal` and also in child CRs.

11. For further information, see [the Fujitsu Enterprise Postgres 17 for Kubernetes Manuals](#).