FUJITSU Enterprise Postgres for Kubernetes

Quick Start Guide

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

It should be noted that it provides supports to both the following CPU Architectures to meet the scope of work.

No	CPU architecture
1	x86
2	s390x

2.2. Supported Platform

It supports running on the following platforms.

No	Platform	Version
1	OpenShift Container Platform	4.6, 4.8, 4.9, 4.10

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 operator install.

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:

    o   Click the checkbox next to the clusters where you want to install the Operator.

o   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

	Name	Platform	Namespace Scope
☐	rj-rhm-amy-test	IBM Cloud	Select Scope ⌄

Cancel   Install

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
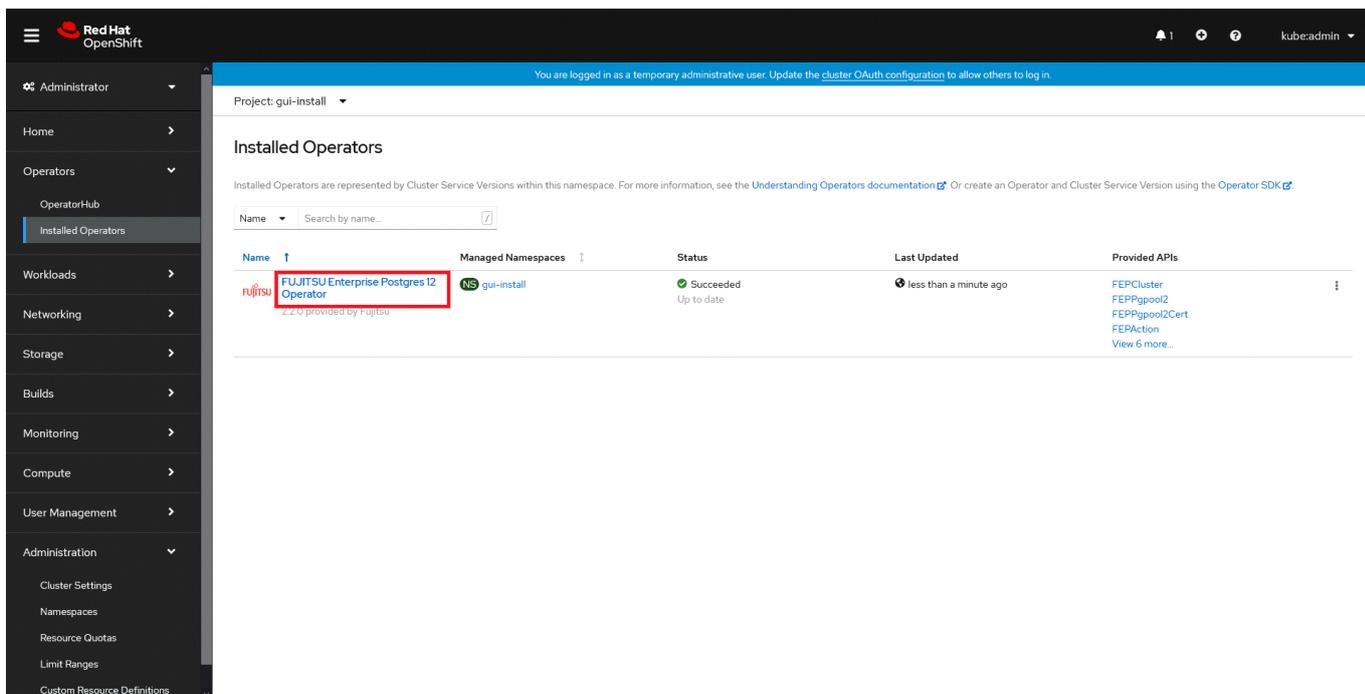
5. Verify status for product is **Succeeded**
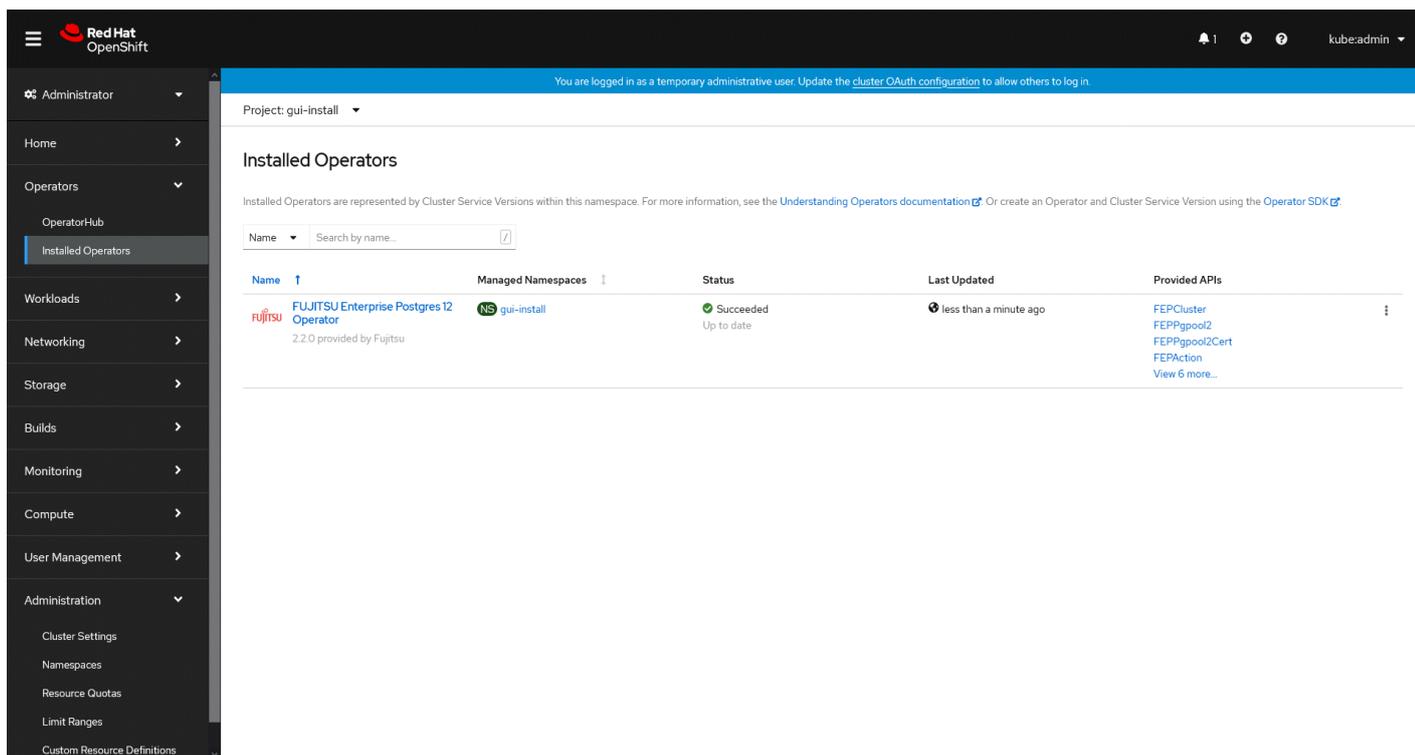
**Installed operators status changes to "Succeeded"**

5. Deploying FEPCluster using Operator

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

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


Copyright 2021-2022 FUJITSU LIMITED

2. It will display a page with all CRs this operator supports. FEPCluster is the main CR and all others are child CR.
   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.
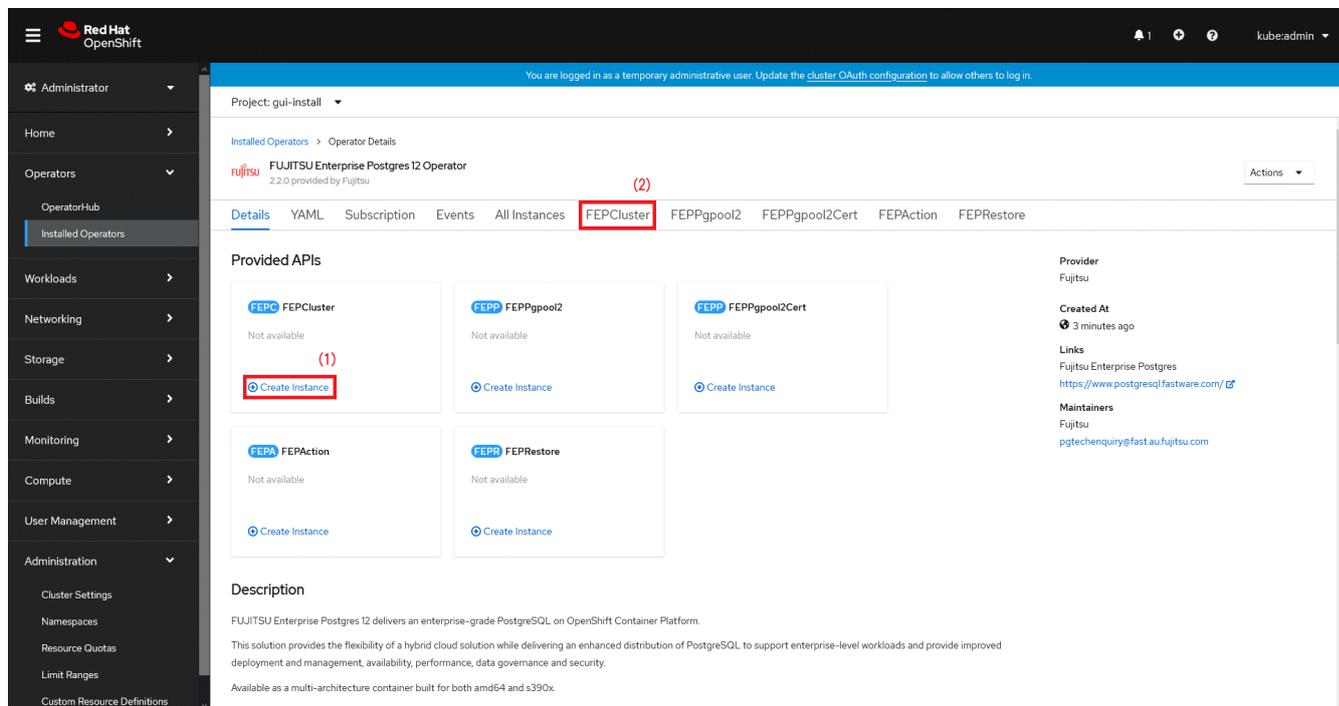
3. This will bring to "Create FEPCluster" page. Here you have two options to configure. The first one is Form View. At the moment, in Form View, one can change only the name of cluster being deployed. Default name is "new-fep".
   This name must be unique within a namespace.
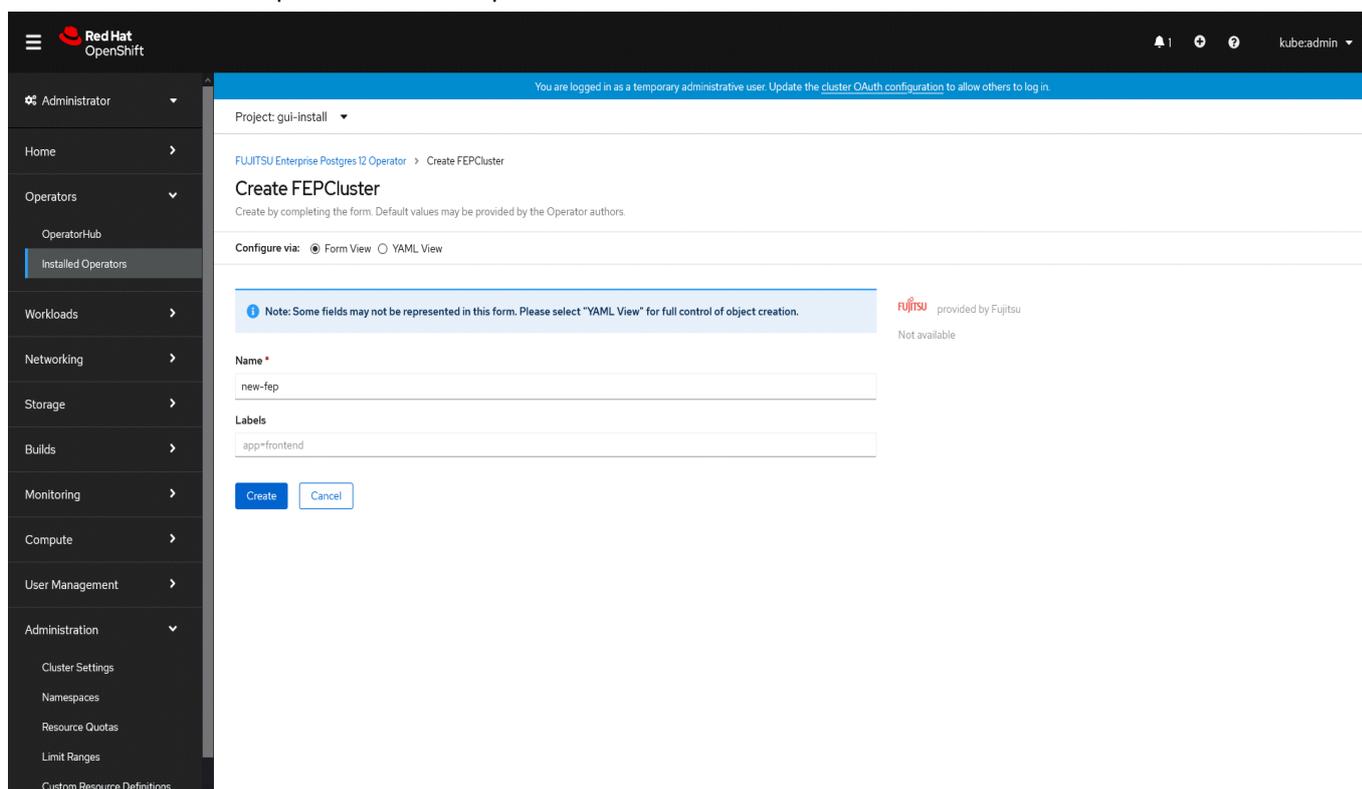
4. In YAML View, starting value of CR is visible and one can choose to modify parameters before creating CR.
    Refer to the [Reference](#) for details of parameters.For example, add a configuration value for the customPgHba
   parameter according to your environment.
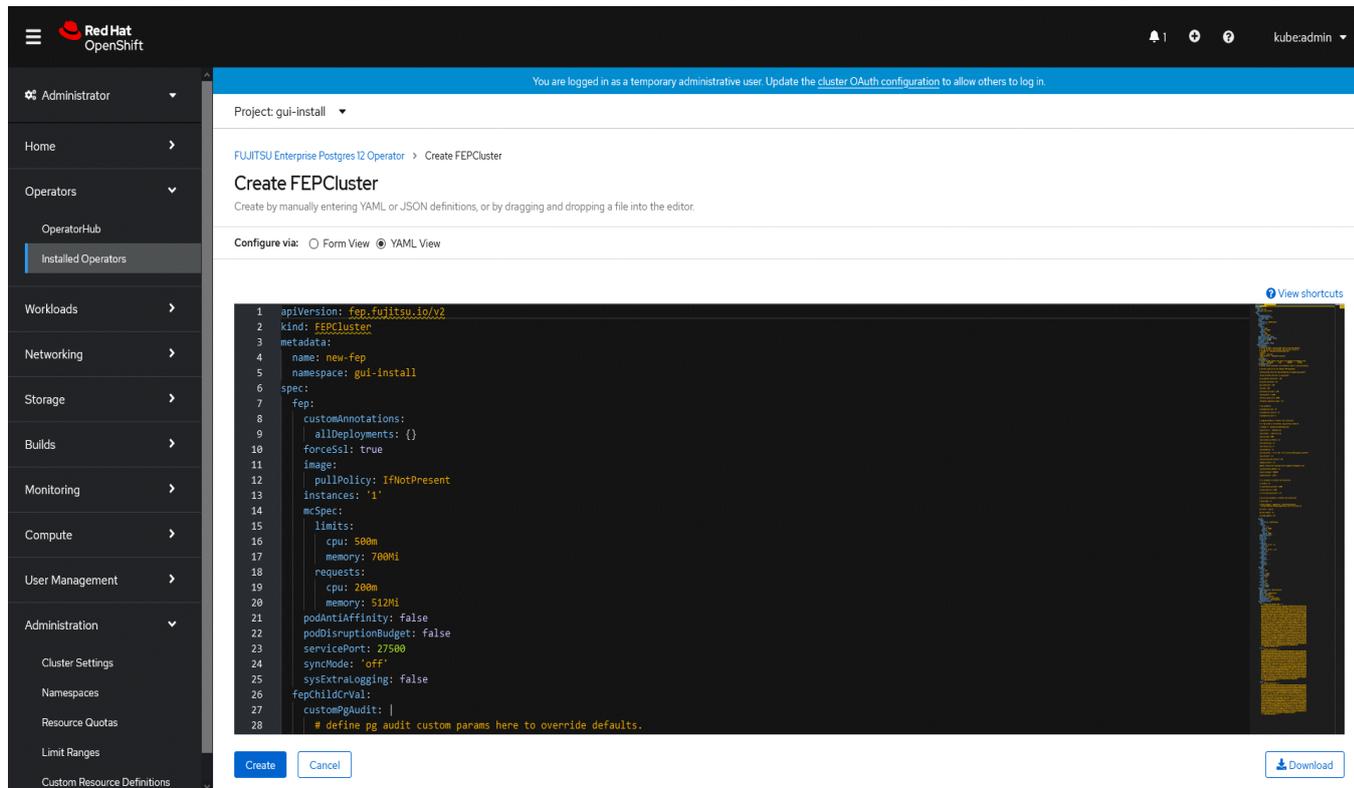
5. When "**Create**" is clicked on either of two pages above, operator creates FEPCluster CR and there after one by
    one FEPBackup, FEPConfig, FEPVolume, FEPUser and FEPCert child CRs are created automatically.
    The starting values for child CRs are taken from "fepChildCrVal" section of FEPCluster 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.
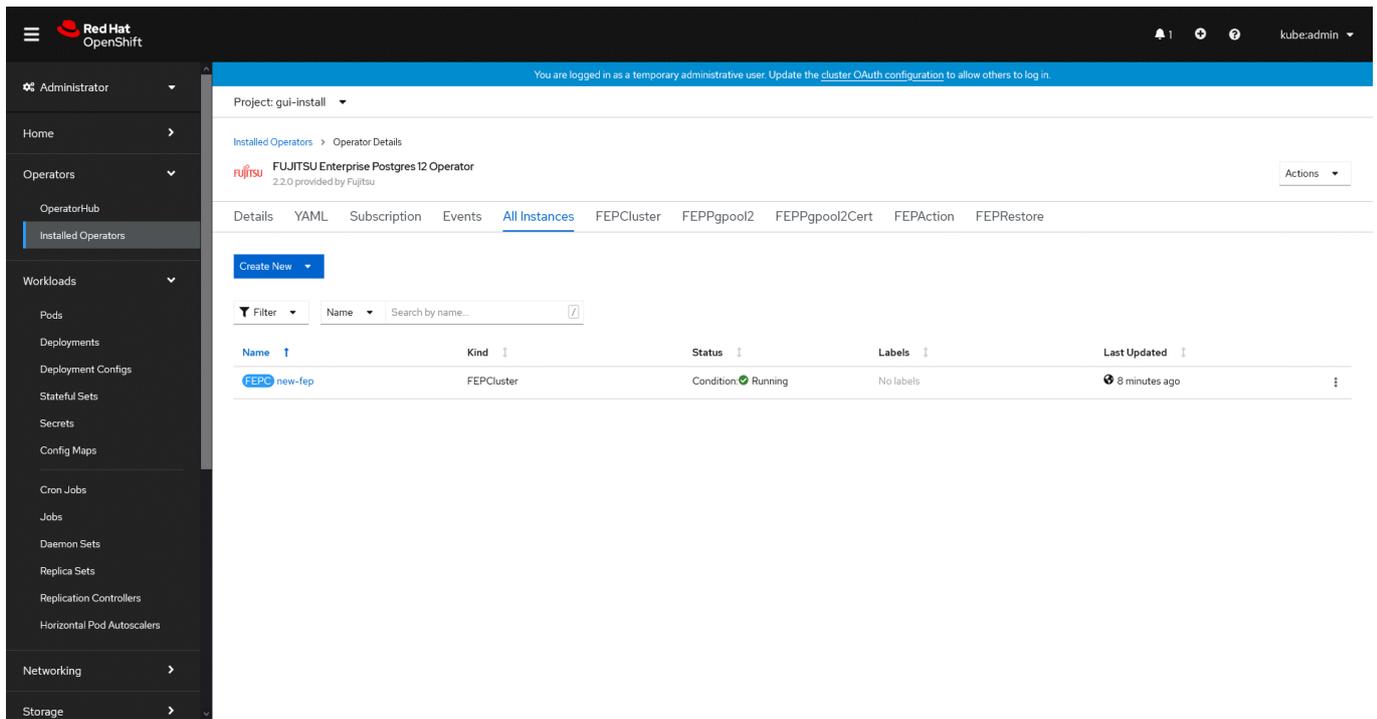
6. In FEPCluster CR, annotations are added to indicate that child CRs are created successfully and has initialized properly. It may take some time to complete.
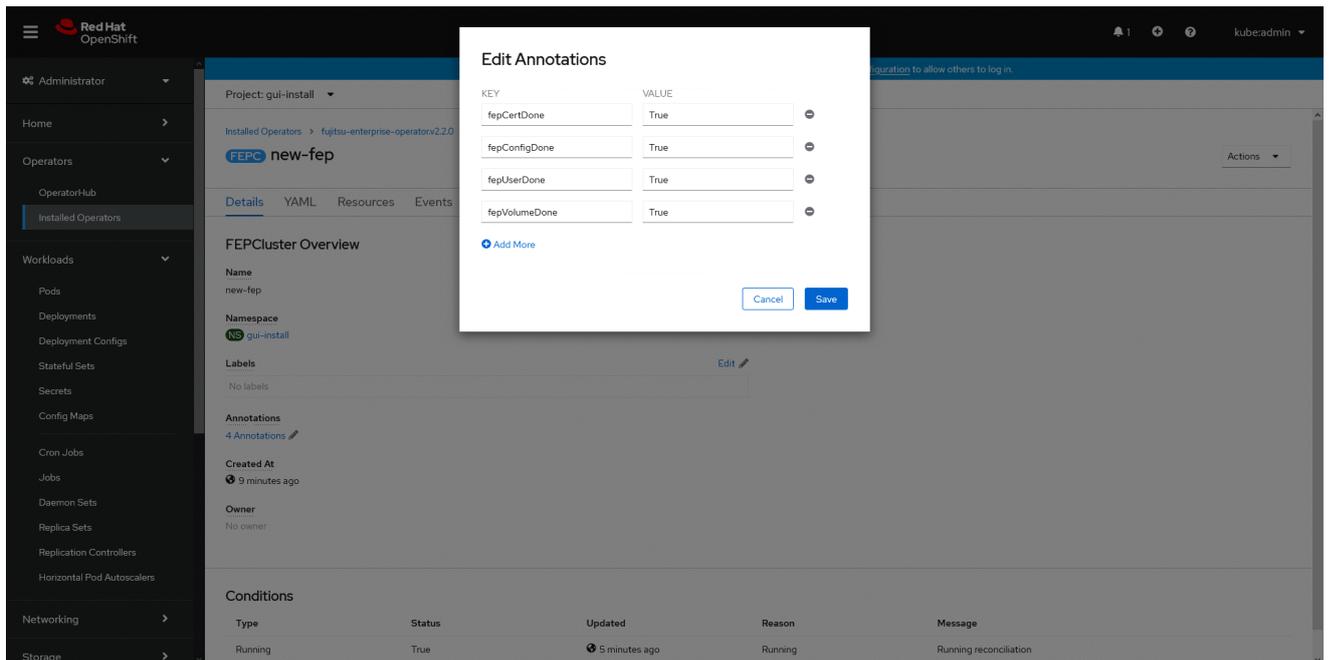

7. Once all four child CRs are marked done in annotations, operator creates StatefulSet for the cluster.
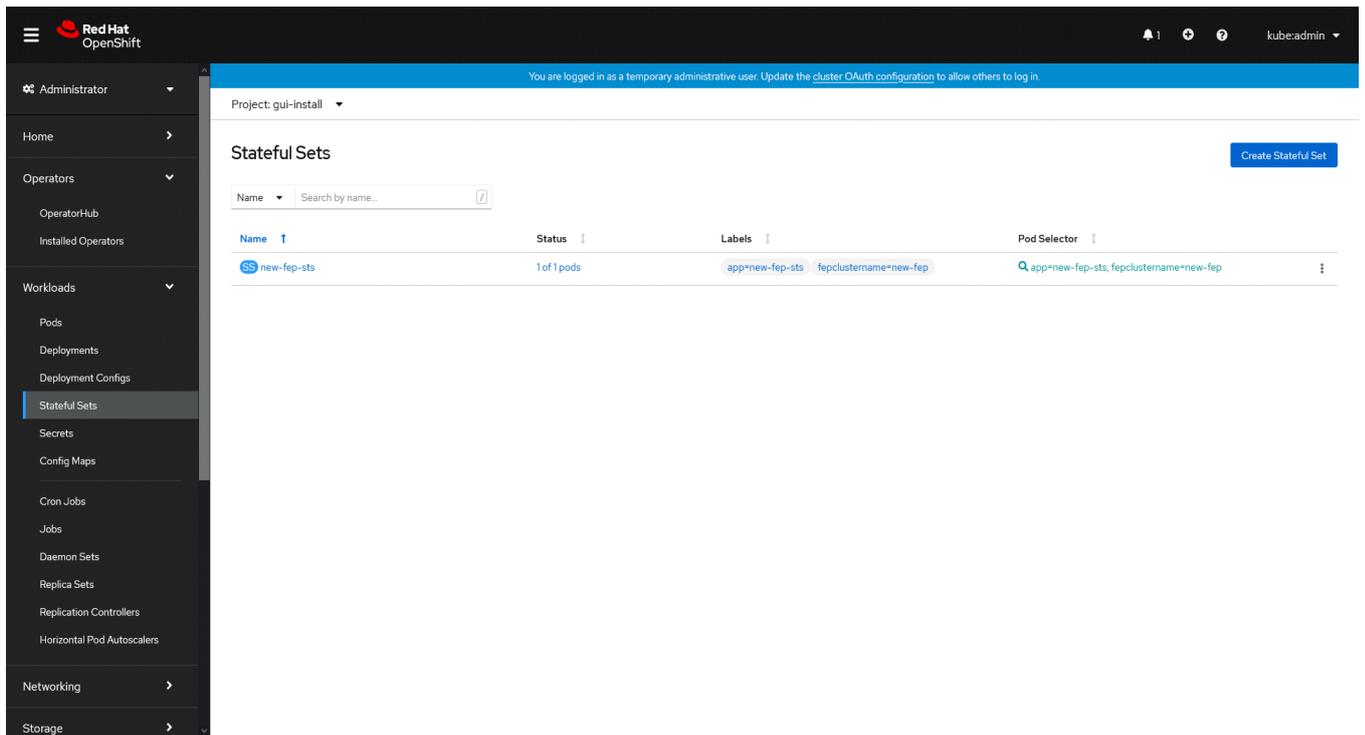
8. StatefulSet will start one FEP instance at one time and will wait it to be ready before starting next one.
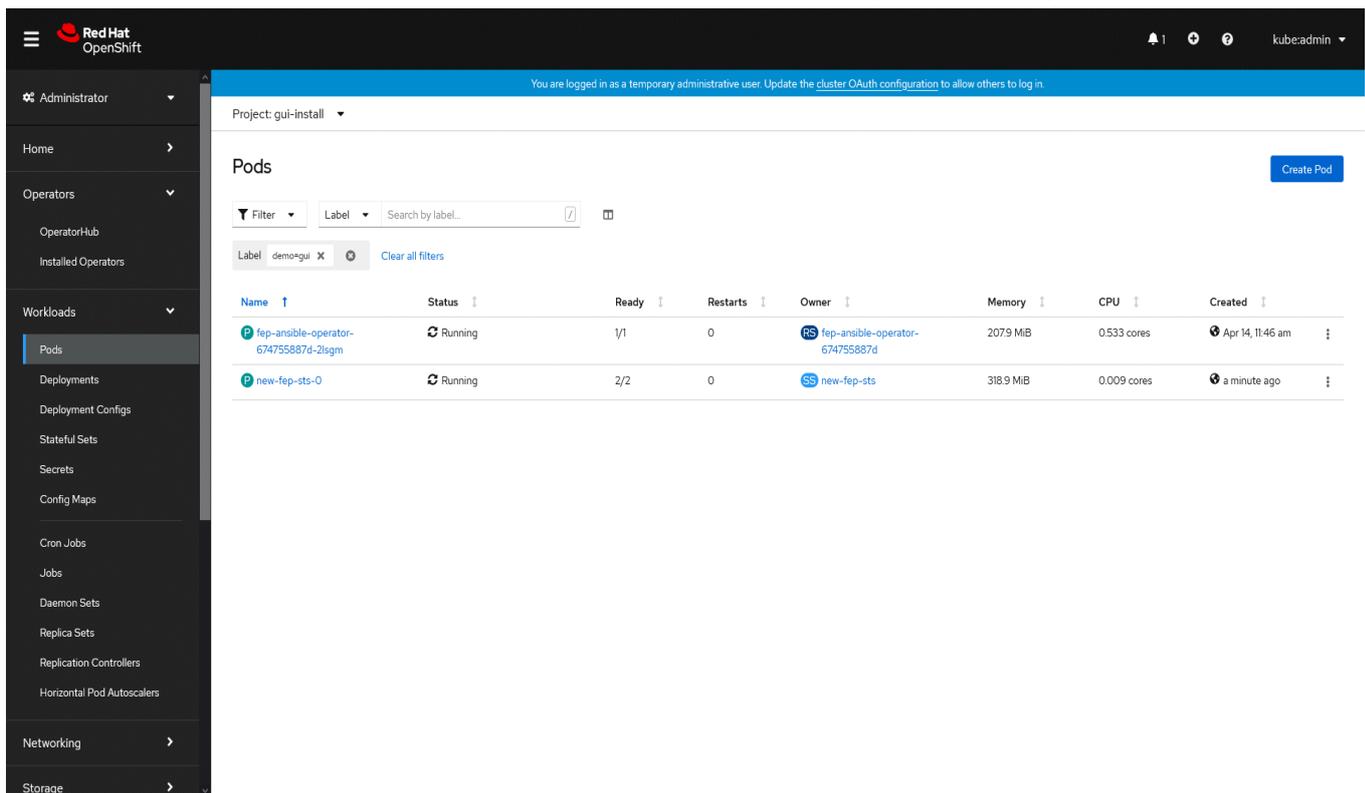
9. Once all instances of FEP servers are started, operator marks a flag "fepClusterReady" in "fepStatus" section of CR to be **true**, indicating that FEPCluster is ready for use. Looking at YAML of FEPCluster CR, it would look like as below:
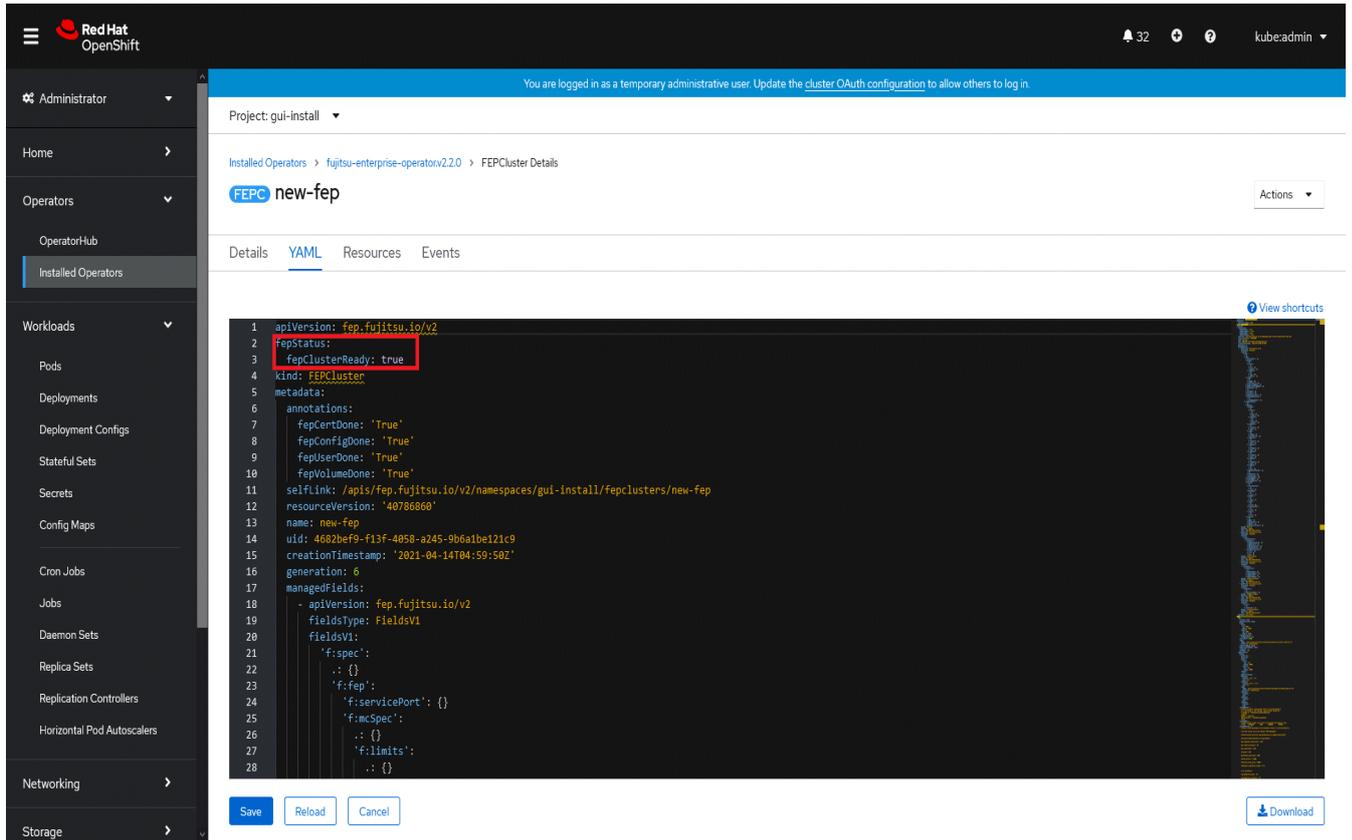
10. Operator also masks the sensitive fields like passwords, passphrase, certificates and keys in FEPCluster fepChildCrVal and also in child CRs.

11. For further information, see the FUJITSU Enterprise Postgres 12 for Kubernetes Manuals