Red Hat build of MicroShift 4.15

Running applications

Running applications in MicroShift

Running applications in MicroShift

Legal Notice

Abstract

This document provides details about how to run applications in MicroShift.

Table of Contents

CHAPTER 1. USING KUSTOMIZE MANIFESTS TO DEPLOY APPLICATIONS

You can use the **kustomize** configuration management tool with application manifests to deploy applications. Read through the following procedures for an example of how Kustomize works in MicroShift.

1.1. HOW KUSTOMIZE WORKS WITH MANIFESTS TO DEPLOY APPLICATIONS

The **kustomize** configuration management tool is integrated with MicroShift. You can use Kustomize and the OpenShift CLI (**oc**) together to apply customizations to your application manifests and deploy those applications to a MicroShift cluster.

- A **kustomization.yaml** file is a specification of resources plus customizations.

- Kustomize uses a **kustomization.yaml** file to load a resource, such as an application, then applies any changes you want to that application manifest and produces a copy of the manifest with the changes overlaid.

- Using a manifest copy with an overlay keeps the original configuration file for your application intact, while enabling you to deploy iterations and customizations of your applications efficiently.

- You can then deploy the application in your MicroShift cluster with an **oc** command.

1.1.1. How MicroShift uses manifests

At every start, MicroShift searches the following manifest directories for Kustomize manifest files:

- **/etc/microshift/manifests**

- **/etc/microshift/manifests.d/***

- **/usr/lib/microshift/**

- **/usr/lib/microshift/manifests.d/***

MicroShift automatically runs the equivalent of the **kubectl apply -k** command to apply the manifests to the cluster if any of the following file types exists in the searched directories:

- **kustomization.yaml**

- **kustomization.yml**

- **Kustomization**

This automatic loading from multiple directories means you can manage MicroShift workloads with the flexibility of having different workloads run independently of each other.

Table 1.1. MicroShift manifest directories

Location	Intent

Location	Intent
/etc/microshift/manifests	Read-write location for configuration management systems or development.
/etc/microshift/manifests.d/*	Read-write location for configuration management systems or development.
/usr/lib/microshift/manifests	Read-only location for embedding configuration manifests on OSTree-based systems.
/usr/lib/microshift/manifestsd./*	Read-only location for embedding configuration manifests on OSTree-based systems.

1.2. OVERRIDE THE LIST OF MANIFEST PATHS

You can override the list of default manifest paths by using a new single path, or by using a new glob pattern for multiple files. Use the following procedure to customize your manifest paths.

Procedure

1. Override the list of default paths by inserting your own values and running one of the following commands:

   a. Set **manifests.kustomizePaths** to **<"/opt/alternate/path">** in the configuration file for a single path.

   b. Set **kustomizePaths** to **,"/opt/alternative/path.d/*".** in the configuration file for a glob pattern.

   ```
   manifests:
     kustomizePaths:
       - <location>  ❶
   ```

   ❶ Set each location entry to an exact path by using **"/opt/alternate/path"** or a glob pattern by using **"/opt/alternative/path.d/*"**.

2. To disable loading manifests, set the configuration option to an empty list.

   ```
   manifests:
     kustomizePaths: []
   ```

   > **NOTE**
   >
   > The configuration file overrides the defaults entirely. If the **kustomizePaths** value is set, only the values in the configuration file are used. Setting the value to an empty list disables manifest loading.

1.3. USING MANIFESTS EXAMPLE

This example demonstrates automatic deployment of a BusyBox container using **kustomize** manifests in the **/etc/microshift/manifests** directory.

Procedure

1. Create the BusyBox manifest files by running the following commands:

   a. Define the directory location:

   ```
   $ MANIFEST_DIR=/etc/microshift/manifests
   ```

   b. Make the directory:

   ```
   $ sudo mkdir -p ${MANIFEST_DIR}
   ```

   c. Place the YAML file in the directory:

   ```
   sudo tee ${MANIFEST_DIR}/busybox.yaml &>/dev/null <<EOF
   apiVersion: v1
   kind: Namespace
   metadata:
     name: busybox
   ---
   apiVersion: apps/v1
   kind: Deployment
   metadata:
     name: busybox
     namespace: busybox-deployment
   spec:
     selector:
       matchLabels:
         app: busybox
     template:
       metadata:
         labels:
           app: busybox
       spec:
         containers:
         - name: busybox
           image: BUSYBOX_IMAGE
           command: [ "/bin/sh", "-c", "while true ; do date; sleep 3600; done;" ]
   EOF
   ```

2. Next, create the **kustomize** manifest files by running the following commands:

   a. Place the YAML file in the directory:

   ```
   sudo tee ${MANIFEST_DIR}/kustomization.yaml &>/dev/null <<EOF
   apiVersion: kustomize.config.k8s.io/v1beta1
   kind: Kustomization
   namespace: busybox
   resources:
     - busybox.yaml
   images:
   ```

```
    - name: BUSYBOX_IMAGE
      newName: busybox:1.35
  EOF
```

3. Restart MicroShift to apply the manifests by running the following command:

```
$ sudo systemctl restart microshift
```

4. Apply the manifests and start the **busybox** pod by running the following command:

```
$ oc get pods -n busybox
```

CHAPTER 2. OPTIONS FOR EMBEDDING MICROSHIFT APPLICATIONS IN A RHEL FOR EDGE IMAGE

You can embed microservices-based workloads and applications in a Red Hat Enterprise Linux for Edge (RHEL for Edge) image to run in a MicroShift cluster. Embedded applications can be installed directly on edge devices to run in air-gapped, disconnected, or offline environments.

2.1. ADDING APPLICATION RPMS TO AN RPM-OSTREE IMAGE

If you have an application that includes APIs, container images, and configuration files for deployment such as manifests, you can build application RPMs. You can then add the RPMs to your RHEL for Edge system image.

The following is an outline of the procedures to embed applications or workloads in a fully self-contained operating system image:

- Build your own RPM that includes your application manifest.

- Add the RPM to the blueprint you used to install Red Hat build of MicroShift.

- Add the workload container images to the same blueprint.

- Create a bootable ISO.

For a step-by-step tutorial about preparing and embedding applications in a RHEL for Edge image, use the following tutorial:

- Embedding applications tutorial

2.2. ADDING APPLICATION MANIFESTS TO AN IMAGE FOR OFFLINE USE

If you have a simple application that includes a few files for deployment such as manifests, you can add those manifests directly to a RHEL for Edge system image.

See the "Create a custom file blueprint customization" section of the following RHEL for Edge documentation for an example:

- Create a custom file blueprint customization

2.3. EMBEDDING APPLICATIONS FOR OFFLINE USE

If you have an application that includes more than a few files, you can embed the application for offline use. See the following procedure:

- Embedding applications for offline use

2.4. ADDITIONAL RESOURCES

- Embedding Red Hat build of MicroShift in an RPM-OSTree image

- Composing, installing, and managing RHEL for Edge images

- Preparing for image building

- Meet Red Hat Device Edge

- Composing a RHEL for Edge image using image builder command-line

- Image Builder system requirements

- Preparing for image building

- Meet Red Hat Device Edge

- Composing a RHEL for Edge image using image builder command-line

- Image Builder system requirements

CHAPTER 3. EMBEDDING APPLICATIONS FOR OFFLINE USE

You can embed microservices-based workloads and applications in a Red Hat Enterprise Linux for Edge (RHEL for Edge) image. Embedding means you can run a Red Hat build of MicroShift cluster in air-gapped, disconnected, or offline environments.

3.1. EMBEDDING WORKLOAD CONTAINER IMAGES FOR OFFLINE USE

To embed container images in devices at the edge that do not have any network connection, you must create a new container, mount the ISO, and then copy the contents into the file system.

Prerequisites

- You have root access to the host.

- Application RPMs have been added to a blueprint.

Procedure

1. Render the manifests, extract all of the container image references, and translate the application image to blueprint container sources by running the following command:

   ```
   $ oc kustomize ~/manifests | grep "image:" | grep -oE '[^ ]+$' | while read line; do echo -e "
   [[containers]]\nsource = \"${line}\"\n"; done >><my_blueprint>.toml
   ```

2. Push the updated blueprint to Image Builder by running the following command:

   ```
   $ sudo composer-cli blueprints push <my_blueprint>.toml
   ```

3. If your workload containers are located in a private repository, you must provide Image Builder with the necessary pull secrets:

   a. Set the **auth_file_path** in the **[containers]** section of the **osbuilder worker** configuration in the **/etc/osbuild-worker/osbuild-worker.toml** file to point to the pull secret.

   b. If needed, create a directory and file for the pull secret, for example:

      Example directory and file

      ```
      [containers]
      auth_file_path = "/<path>/pull-secret.json"  ❶
      ```

      ❶ Use the custom location previously set for copying and retrieving images.

4. Build the container image by running the following command:

   ```
   $ sudo composer-cli compose start-ostree <my_blueprint> edge-commit
   ```

5. Proceed with your preferred **rpm-ostree** image flow, such as waiting for the build to complete, exporting the image and integrating it into your **rpm-ostree** repository or creating a bootable ISO.

3.2. ADDITIONAL RESOURCES

- Options for embedding Red Hat build of MicroShift applications in a RHEL for Edge image

- Creating the RHEL for Edge image

- Add the blueprint to Image Builder and build the ISO

- Download the ISO and prepare it for use

- Upgrading RHEL for Edge systems

CHAPTER 4. EMBEDDING RED HAT BUILD OF MICROSHIFT APPLICATIONS TUTORIAL

The following tutorial gives a detailed example of how to embed applications in a RHEL for Edge image for use in a MicroShift cluster in various environments.

4.1. EMBED APPLICATION RPMS TUTORIAL

The following tutorial reviews the MicroShift installation steps and adds a description of the workflow for embedding applications. If you are already familiar with **rpm-ostree** systems such as Red Hat Enterprise Linux for Edge (RHEL for Edge) and MicroShift, you can go straight to the procedures.

4.1.1. Installation workflow review

Embedding applications requires a similar workflow to embedding MicroShift into a RHEL for Edge image.

- The following image shows how system artifacts such as RPMs, containers, and files are added to a blueprint and used by the image composer to create an ostree commit.

- The ostree commit then can follow either the ISO path or the repository path to edge devices.

- The ISO path can be used for disconnected environments, while the repository path is often used in places were the network is usually connected.

Embedding MicroShift workflow
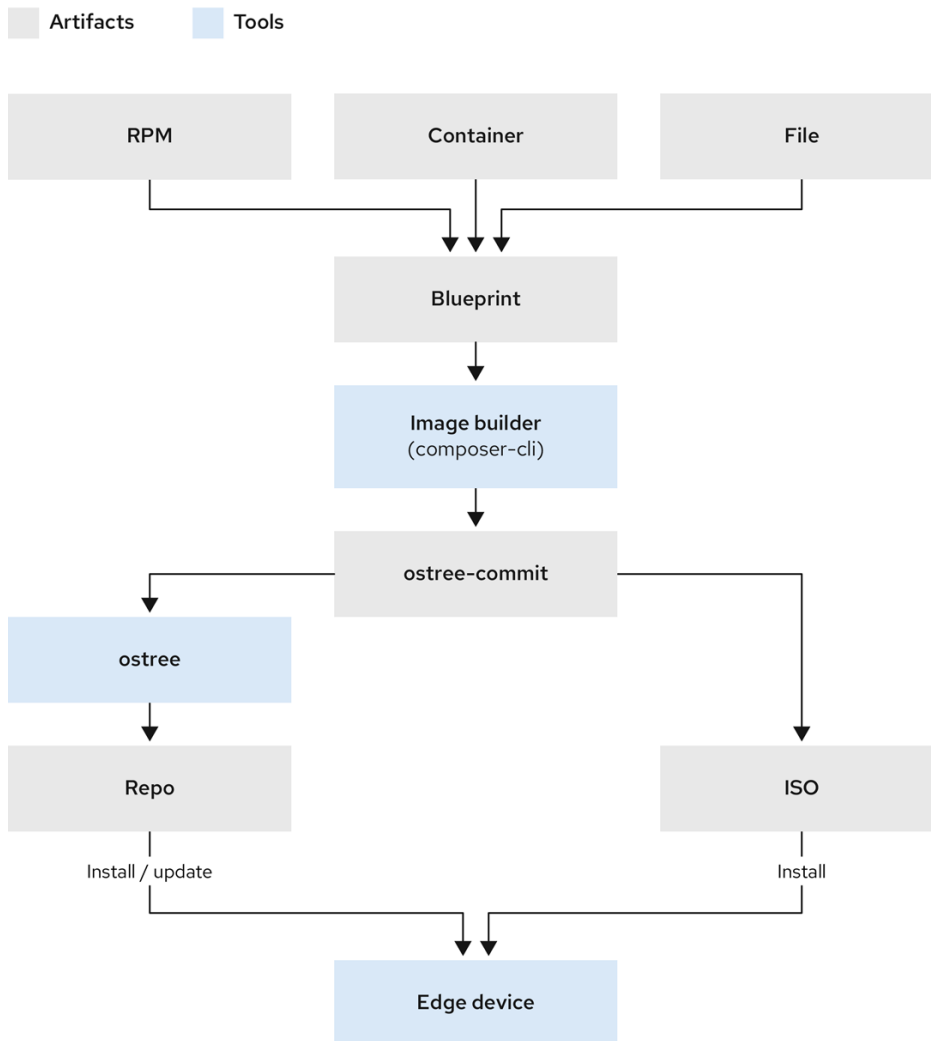
Reviewing these steps can help you understand the steps needed to embed an application:

1. To embed MicroShift on RHEL for Edge, you added the MicroShift repositories to Image Builder.

2. You created a blueprint that declared all the RPMs, container images, files and customizations you needed, including the addition of MicroShift.

3. You added the blueprint to Image Builder and ran a build with the Image Builder CLI tool (**composer-cli**). This step created **rpm-ostree** commits, which were used to create the container image. This image contained RHEL for Edge.

4. You added the installer blueprint to Image Builder to create an **rpm-ostree** image (ISO) to boot from. This build contained both RHEL for Edge and MicroShift.

5. You downloaded the ISO with MicroShift embedded, prepared it for use, provisioned it, then installed it onto your edge devices.

4.1.2. Embed application RPMs workflow

After you have set up a build host that meets the Image Builder requirements, you can add your application in the form of a directory of manifests to the image. After those steps, the simplest way to embed your application or workload into a new ISO is to create your own RPMs that include the manifests. Your application RPMs contain all of the configuration files describing your deployment.

The following "Embedding applications workflow" image shows how Kubernetes application manifests and RPM spec files are combined in a single application RPM build. This build becomes the RPM artifact included in the workflow for embedding MicroShift in an ostree commit.

Embedding applications workflow

The following procedures use the **rpmbuild** tool to create a specification file and local repository. The specification file defines how the package is built, moving your application manifests to the correct location inside the RPM package for MicroShift to pick them up. That RPM package is then embedded in the ISO.

4.1.3. Preparing to make application RPMs

To build your own RPMs, choose a tool of your choice, such as the **rpmbuild** tool, and initialize the RPM build tree in your home directory. The following is an example procedure. As long as your RPMs are accessible to Image Builder, you can use the method you prefer to build the application RPMs.

Prerequisites

- You have set up a Red Hat Enterprise Linux for Edge (RHEL for Edge) 9.2 build host that meets the Image Builder system requirements.

- You have root access to the host.

Procedure

1. Install the **rpmbuild** tool and create the yum repository for it by running the following command:

```
$ sudo dnf install rpmdevtools rpmlint yum-utils createrepo
```

2. Create the file tree you need to build RPM packages by running the following command:

```
$ rpmdev-setuptree
```

Verification

- List the directories to confirm creation by running the following command:

```
$ ls ~/rpmbuild/
```

Example output

```
BUILD RPMS SOURCES SPECS SRPMS
```

4.1.4. Building the RPM package for the application manifests

To build your own RPMs, you must create a spec file that adds the application manifests to the RPM package. The following is an example procedure. As long as the application RPMs and other elements needed for image building are accessible to Image Builder, you can use the method that you prefer.

Prerequisites

- You have set up a Red Hat Enterprise Linux for Edge (RHEL for Edge) 9.2 build host that meets the Image Builder system requirements.

- You have root access to the host.

- The file tree required to build RPM packages was created.

Procedure

1. In the **~/rpmbuild/SPECS** directory, create a file such as **<application_workload_manifests.spec>** using the following template:

   Example spec file

   ```
   Name: <application_workload_manifests>
   Version: 0.0.1
   Release: 1%{?dist}
   Summary: Adds workload manifests to microshift
   BuildArch: noarch
   License: GPL
   Source0: %{name}-%{version}.tar.gz
   #Requires: microshift
   %description
   Adds workload manifests to microshift
   %prep
   %autosetup
   %install ❶
   rm -rf $RPM_BUILD_ROOT
   mkdir -p $RPM_BUILD_ROOT/%{_prefix}/lib/microshift/manifests
   cp -pr ~/manifests $RPM_BUILD_ROOT/%{_prefix}/lib/microshift/
   %clean
   ```

```
rm -rf $RPM_BUILD_ROOT

%files
%{_prefix}/lib/microshift/manifests/**
%changelog
* <DDD MM DD YYYY username@domain - V major.minor.patch>
- <your_change_log_comment>
```

1. The **%install** section creates the target directory inside the RPM package, **/usr/lib/microshift/manifests/** and copies the manifests from the source home directory, ~/**manifests**.

> **IMPORTANT**
>
> All of the required YAML files must be in the source home directory ~/**manifests**, including a **kustomize.yaml** file if you are using kustomize.

2. Build your RPM package in the ~/**rpmbuild/RPMS** directory by running the following command:

   ```
   $ rpmbuild -bb ~/rpmbuild/SPECS/<application_workload_manifests.spec>
   ```

4.1.5. Adding application RPMs to a blueprint

To add application RPMs to a blueprint, you must create a local repository that Image Builder can use to create the ISO. With this procedure, the required container images for your workload can be pulled over the network.

Prerequisites

- You have root access to the host.

- Workload or application RPMs exist in the ~/**rpmbuild/RPMS** directory.

Procedure

1. Create a local RPM repository by running the following command:

   ```
   $ createrepo ~/rpmbuild/RPMS/
   ```

2. Give Image Builder access to the RPM repository by running the following command:

   ```
   $ sudo chmod a+rx ~
   ```

   > **NOTE**
   >
   > You must ensure that Image Builder has all of the necessary permissions to access all of the files needed for image building, or the build cannot proceed.

3. Create the blueprint file, **repo-local-rpmbuild.toml** using the following template:

   ```
   id = "local-rpm-build"
   ```

```
name = "RPMs build locally"
type = "yum-baseurl"
url = "file://<path>/rpmbuild/RPMS"  1
check_gpg = false
check_ssl = false
system = false
```

1   Specify part of the path to create a location that you choose. Use this path in the later commands to set up the repository and copy the RPMs.

4. Add the repository as a source for Image Builder by running the following command:

```
$ sudo composer-cli sources add repo-local-rpmbuild.toml
```

5. Add the RPM to your blueprint, by adding the following lines:

```
…
[[packages]]
name = "<application_workload_manifests>"  1
version = "*"
…
```

1   Add the name of your workload here.

6. Push the updated blueprint to Image Builder by running the following command:

```
$ sudo composer-cli blueprints push repo-local-rpmbuild.toml
```

7. At this point, you can either run Image Builder to create the ISO, or embed the container images for offline use.

    a. To create the ISO, start Image Builder by running the following command:

    ```
    $ sudo composer-cli compose start-ostree repo-local-rpmbuild edge-commit
    ```

In this scenario, the container images are pulled over the network by the edge device during startup.

Additional resources

- Composing a RHEL for Edge image using the Image Builder CLI

- Network-based deployments workflow

4.2. ADDITIONAL RESOURCES

- Embedding applications for offline use

- Embedding Red Hat build of MicroShift in an RPM-OSTree image

- Composing, installing, and managing RHEL for Edge images

- Preparing for image building

- Meet Red Hat Device Edge with Red Hat build of MicroShift

- How to create a Linux RPM package

- Composing a RHEL for Edge image using image builder command-line

- Image Builder system requirements

- Meet Red Hat Device Edge with Red Hat build of MicroShift

- How to create a Linux RPM package

- Composing a RHEL for Edge image using image builder command-line

CHAPTER 5. GREENBOOT WORKLOAD HEALTH CHECK SCRIPTS

Greenboot health check scripts are helpful on edge devices where direct serviceability is either limited or non-existent. You can create health check scripts to assess the health of your workloads and applications. These additional health check scripts are useful components of software problem checks and automatic system rollbacks.

A MicroShift health check script is included in the **microshift-greenboot** RPM. You can also create your own health check scripts based on the workloads you are running. For example, you can write one that verifies that a service has started.

5.1. HOW WORKLOAD HEALTH CHECK SCRIPTS WORK

The workload or application health check script described in this tutorial uses the MicroShift health check functions that are available in the **/usr/share/microshift/functions/greenboot.sh** file. This enables you to reuse procedures already implemented for the MicroShift core services.

The script starts by running checks that the basic functions of the workload are operating as expected. To run the script successfully:

- Execute the script from a root user account.

- Enable the MicroShift service.

The health check performs the following actions:

- Gets a wait timeout of the current boot cycle for the **wait_for** function.

- Calls the **namespace_images_downloaded** function to wait until pod images are available.

- Calls the **namespace_pods_ready** function to wait until pods are ready.

- Calls the **namespace_pods_not_restarting** function to verify pods are not restarting.

> **NOTE**
>
> Restarting pods can indicate a crash loop.

5.2. INCLUDED GREENBOOT HEALTH CHECKS

Health check scripts are available in **/usr/lib/greenboot/check**, a read-only directory in RPM-OSTree systems. The following health checks are included with the **greenboot-default-health-checks** framework.

- Check if repository URLs are still DNS solvable:
  This script is under **/usr/lib/greenboot/check/required.d/01_repository_dns_check.sh** and ensures that DNS queries to repository URLs are still available.

- Check if update platforms are still reachable:
  This script is under **/usr/lib/greenboot/check/wanted.d/01_update_platform_check.sh** and tries to connect and get a 2XX or 3XX HTTP code from the update platforms defined in **/etc/ostree/remotes.d**.

- Check if the current boot has been triggered by the hardware watchdog:
  This script is under **/usr/lib/greenboot/check/required.d/02_watchdog.sh** and checks whether the current boot has been watchdog-triggered or not.

  - If the watchdog-triggered reboot occurs within the grace period, the current boot is marked as red. Greenboot does not trigger a rollback to the previous deployment.

  - If the watchdog-triggered reboot occurs after the grace period, the current boot is not marked as red. Greenboot does not trigger a rollback to the previous deployment.

  - A 24-hour grace period is enabled by default. This grace period can be either disabled by modifying **GREENBOOT_WATCHDOG_CHECK_ENABLED** in **/etc/greenboot/greenboot.conf to false**, or configured by changing the **GREENBOOT_WATCHDOG_GRACE_PERIOD=number_of_hours** variable value in **/etc/greenboot/greenboot.conf**.

5.3. HOW TO CREATE A HEALTH CHECK SCRIPT FOR YOUR APPLICATION

You can create workload or application health check scripts in the text editor of your choice using the example in this documentation. Save the scripts in the **/etc/greenboot/check/required.d** directory. When a script in the **/etc/greenboot/check/required.d** directory exits with an error, Greenboot triggers a reboot in an attempt to heal the system.

> **NOTE**
>
> Any script in the **/etc/greenboot/check/required.d** directory triggers a reboot if it exits with an error.

If your health check logic requires any post-check steps, you can also create additional scripts and save them in the relevant greenboot directories. For example:

- You can also place shell scripts you want to run after a boot has been declared successful in **/etc/greenboot/green.d**.

- You can place shell scripts you want to run after a boot has been declared failed in **/etc/greenboot/red.d**. For example, if you have steps to heal the system before restarting, you can create scripts for your use case and place them in the **/etc/greenboot/red.d** directory.

5.3.1. About the workload health check script example

The following example uses the MicroShift health check script as a template. You can use this example with the provided libraries as a guide for creating basic health check scripts for your applications.

5.3.1.1. Basic prerequisites for creating a health check script

- The workload must be installed.

- You must have root access.

5.3.1.2. Example and functional requirements

You can start with the following example health check script. Modify it for your use case. In your workload health check script, you must complete the following minimum steps:

- Set the environment variables.

- Define the user workload namespaces.

- List the expected pod count.

> **IMPORTANT**
>
> Choose a name prefix for your application that ensures it runs after the **40_microshift_running_check.sh** script, which implements the Red Hat build of MicroShift health check procedure for its core services.

Example workload health check script

```bash
# #!/bin/bash
set -e

SCRIPT_NAME=$(basename $0)
PODS_NS_LIST=(<user_workload_namespace1> <user_workload_namespace2>)
PODS_CT_LIST=(<user_workload_namespace1_pod_count>
<user_workload_namespace2_pod_count>)
# Update these two lines with at least one namespace and the pod counts that are specific to your
workloads. Use the kubernetes <namespace> where your workload is deployed.

# Set Greenboot to read and execute the workload health check functions library.
source /usr/share/microshift/functions/greenboot.sh

# Set the exit handler to log the exit status.
trap 'script_exit' EXIT

# Set the script exit handler to log a `FAILURE` or `FINISHED` message depending on the exit status
of the last command.
# args: None
# return: None
function script_exit() {
    [ "$?" -ne 0 ] && status=FAILURE || status=FINISHED
    echo $status
}

# Set the system to automatically stop the script if the user running it is not 'root'.
if [ $(id -u) -ne 0 ] ; then
    echo "The '${SCRIPT_NAME}' script must be run with the 'root' user privileges"
    exit 1
fi

echo "STARTED"

# Set the script to stop without reporting an error if the MicroShift service is not running.
if [ $(systemctl is-enabled microshift.service 2>/dev/null) != "enabled" ] ; then
    echo "MicroShift service is not enabled. Exiting..."
    exit 0
fi

# Set the wait timeout for the current check based on the boot counter.
WAIT_TIMEOUT_SECS=$(get_wait_timeout)
```

```
# Set the script to wait for the pod images to be downloaded.
for i in ${!PODS_NS_LIST[@]}; do
    CHECK_PODS_NS=${PODS_NS_LIST[$i]}

    echo "Waiting ${WAIT_TIMEOUT_SECS}s for pod image(s) from the ${CHECK_PODS_NS}
namespace to be downloaded"
    wait_for ${WAIT_TIMEOUT_SECS} namespace_images_downloaded
done

# Set the script to wait for pods to enter ready state.
for i in ${!PODS_NS_LIST[@]}; do
    CHECK_PODS_NS=${PODS_NS_LIST[$i]}
    CHECK_PODS_CT=${PODS_CT_LIST[$i]}

    echo "Waiting ${WAIT_TIMEOUT_SECS}s for ${CHECK_PODS_CT} pod(s) from the
${CHECK_PODS_NS} namespace to be in 'Ready' state"
    wait_for ${WAIT_TIMEOUT_SECS} namespace_pods_ready
done

# Verify that pods are not restarting by running, which could indicate a crash loop.
for i in ${!PODS_NS_LIST[@]}; do
    CHECK_PODS_NS=${PODS_NS_LIST[$i]}

    echo "Checking pod restart count in the ${CHECK_PODS_NS} namespace"
    namespace_pods_not_restarting ${CHECK_PODS_NS}
done
```

5.4. TESTING A WORKLOAD HEALTH CHECK SCRIPT

Prerequisites

- You have root access.

- You have installed a workload.

- You have created a health check script for the workload.

- The Red Hat build of MicroShift service is enabled.

Procedure

1. To test that Greenboot is running a health check script file, reboot the host by running the following command:

   ```
   $ sudo reboot
   ```

2. Examine the output of Greenboot health checks by running the following command:

   ```
   $ sudo journalctl -o cat -u greenboot-healthcheck.service
   ```

   > **NOTE**
   >
   > MicroShift core service health checks run before the workload health checks.

Example output

```
GRUB boot variables:
boot_success=0
boot_indeterminate=0
Greenboot variables:
GREENBOOT_WATCHDOG_CHECK_ENABLED=true

...

...
FINISHED
Script '40_microshift_running_check.sh' SUCCESS
Running Wanted Health Check Scripts...
Finished greenboot Health Checks Runner.
```

5.5. ADDITIONAL RESOURCES

- The Greenboot health check

- Auto applying manifests

CHAPTER 6. AUTOMATING APPLICATION MANAGEMENT WITH THE GITOPS CONTROLLER

GitOps with Argo CD for MicroShift is a lightweight, optional add-on controller derived from the Red Hat OpenShift GitOps Operator. GitOps for MicroShift uses the command-line interface (CLI) of Argo CD to interact with the GitOps controller that acts as the declarative GitOps engine. You can consistently configure and deploy Kubernetes-based infrastructure and applications across clusters and development lifecycles.

6.1. WHAT YOU CAN DO WITH THE GITOPS AGENT

By using the GitOps with Argo CD agent with MicroShift, you can utilize the following principles:

- Implement application lifecycle management.

  - Create and manage your clusters and application configuration files using the core principles of developing and maintaining software in a Git repository.

  - You can update the single repository and GitOps automates the deployment of new applications or updates to existing ones.

  - For example, if you have 1,000 edge devices, each using MicroShift and a local GitOps agent, you can easily add or update an application on all 1,000 devices with just one change in your central Git repository.

- The Git repository contains a declarative description of the infrastructure you need in your specified environment and contains an automated process to make your environment match the described state.

- You can also use the Git repository as an audit trail of changes so that you can create processes based on Git flows such as review and approval for merging pull requests that implement configuration changes.

6.2. CREATING GITOPS APPLICATIONS ON MICROSHIFT

You can create a custom YAML configuration to deploy and manage applications in your MicroShift service. To install the necessary packages to run GitOps applications, follow the documentation in "Installing the GitOps Argo CD manifests from an RPM package".

Prerequisites

- You installed the **microshift-gitops** packages and the Argo CD pods are running in the **openshift-gitops** namespace.

Procedure

1. Create a YAML file and add your customized configurations for the application:

   Example YAML for a cert-manager application

   ```
   kind: AppProject
   apiVersion: argoproj.io/v1alpha1
   metadata:
     name: default
   ```

```
  namespace: openshift-gitops
spec:
  clusterResourceWhitelist:
  - group: '*'
    kind: '*'
  destinations:
  - namespace: '*'
    server: '*'
  sourceRepos:
  - '*'
---
apiVersion: argoproj.io/v1alpha1
kind: Application
metadata:
  name: cert-manager
  namespace: openshift-gitops
spec:
  destination:
    namespace: cert-manager
    server: https://kubernetes.default.svc
  project: default
  source:
    path: cert-manager
    repoURL: https://github.com/anandf/microshift-install
  syncPolicy:
    automated: {}
    syncOptions:
    - CreateNamespace=true
    - ServerSideApply=true
```

Example YAML for a **spring-petclinic** application

```
kind: AppProject
apiVersion: argoproj.io/v1alpha1
metadata:
  name: default
  namespace: openshift-gitops
spec:
  clusterResourceWhitelist:
  - group: '*'
    kind: '*'
  destinations:
  - namespace: '*'
    server: '*'
  sourceRepos:
  - '*'
---
kind: Application
apiVersion: argoproj.io/v1alpha1
metadata:
  name: spring-petclinic
  namespace: openshift-gitops
spec:
  destination:
    namespace: spring-petclinic
```

```
      server: https://kubernetes.default.svc
    project: default
    source:
      directory:
        recurse: true
      path: app
      repoURL: https://github.com/siamaksade/openshift-gitops-getting-started
    syncPolicy:
      automated: {}
      syncOptions:
      - CreateNamespace=true
      - ServerSideApply=true
```

2. To deploy the applications defined in the YAML file, run the following command:

```
$ oc apply -f <filename>.yaml
```

Verification

- To verify your application is deployed and synced, run the following command:

```
$ oc get applications -A
```

It might take a few minutes for the application to show the **Healthy** status.

Example output

```
NAMESPACE        NAME           SYNC STATUS   HEALTH STATUS
openshift-gitops  cert-manager     Synced        Healthy
openshift-gitops  spring-petclinic  Synced        Healthy
```

Additional resources

- Installing the GitOps Argo CD manifests from an RPM package

6.3. LIMITATIONS OF USING THE GITOPS AGENT WITH MICROSHIFT

GitOps with Argo CD for MicroShift has the following differences from the Red Hat OpenShift GitOps Operator:

- The **gitops-operator** component is not used with MicroShift.

- To maintain the small resource use of MicroShift, the Argo CD web console is not available. You can use the Argo CD CLI or use a pull-based approach.

- Because MicroShift is single-node, there is no multi-cluster support. Each instance of MicroShift is paired with a local GitOps agent.

- The **oc adm must-gather** command is not available in MicroShift.

6.4. TROUBLESHOOTING GITOPS

If you have problems with your GitOps controller, you can use either the OpenShift CLI (**oc**) tool or run an sos report.

6.4.1. Debugging GitOps with oc adm inspect

You can debug GitOps by using the OpenShift CLI (**oc**).

Prerequisites

- The **oc** command line tool is installed.

Procedure

1. Run the **oc adm inspect** command when in the GitOps namespace:

   ```
   $ oc adm inspect ns/openshift-gitops
   ```

   Example output

   ```
   Gathering data for ns/openshift-gitops...
   W0501 20:34:35.978508 57625 util.go:118] the server doesn't have a resource type
   egressfirewalls, skipping the inspection
   W0501 20:34:35.980881 57625 util.go:118] the server doesn't have a resource type
   egressqoses, skipping the inspection
   W0501 20:34:36.040664 57625 util.go:118] the server doesn't have a resource type
   servicemonitors, skipping the inspection
   Wrote inspect data to inspect.local.2673575938140296280.
   ```

Next steps

- If **oc adm inspect** did not provide the information you need, you can run an sos report.

6.4.2. Gathering data from an sos report

Prerequisites

- You must have the **sos** package installed.

Procedure

1. Log into the failing host as a root user.

2. Perform the debug report creation procedure by running the following command:

   ```
   $ microshift-sos-report
   ```

   Example output

   ```
   sosreport (version 4.5.1)

   This command will collect diagnostic and configuration information from
   this Red Hat Enterprise Linux system and installed applications.
   ```

An archive containing the collected information will be generated in /var/tmp/sos.o0sznf_8 and may be provided to a Red Hat support representative.

Any information provided to Red Hat will be treated in accordance with the published support policies at:

> Distribution Website : https://www.redhat.com/
> Commercial Support   : https://www.access.redhat.com/

The generated archive may contain data considered sensitive and its content should be reviewed by the originating organization before being passed to any third party.

No changes will be made to system configuration.

 Setting up archive ...
 Setting up plugins ...
 Running plugins. Please wait ...

 Starting 1/2   microshift   [Running: microshift]
 Starting 2/2   microshift_ovn [Running: microshift microshift_ovn]
 Finishing plugins   [Running: microshift]

 Finished running plugins

Found 1 total reports to obfuscate, processing up to 4 concurrently

sosreport-microshift-rhel9-2023-03-31-axjbyxw :   Beginning obfuscation...
sosreport-microshift-rhel9-2023-03-31-axjbyxw :   Obfuscation completed

Successfully obfuscated 1 report(s)

Creating compressed archive...

A mapping of obfuscated elements is available at
 /var/tmp/sosreport-microshift-rhel9-2023-03-31-axjbyxw-private_map

Your sosreport has been generated and saved in:
 /var/tmp/sosreport-microshift-rhel9-2023-03-31-axjbyxw-obfuscated.tar.xz

 Size 444.14KiB
 Owner root
 sha256 922e5ff2db25014585b7c6c749d2c44c8492756d619df5e9838ce863f83d4269

Please send this file to your support representative.

6.5. ADDITIONAL RESOURCES

- Using sos reports

- Red Hat OpenShift GitOps

- Generating an sos report for technical support (Red Hat Enterprise Linux)

- Generating an sos report for technical support (Red Hat Enterprise Linux)

CHAPTER 7. POD SECURITY AUTHENTICATION AND AUTHORIZATION

7.1. UNDERSTANDING AND MANAGING POD SECURITY ADMISSION

Pod security admission is an implementation of the Kubernetes pod security standards . Use pod security admission to restrict the behavior of pods.

7.2. SECURITY CONTEXT CONSTRAINT SYNCHRONIZATION WITH POD SECURITY STANDARDS

MicroShift includes Kubernetes pod security admission .

In addition to the global pod security admission control configuration, a controller exists that applies pod security admission control **warn** and **audit** labels to namespaces according to the security context constraint (SCC) permissions of the service accounts that are in a given namespace.

> **IMPORTANT**
>
> Namespaces that are defined as part of the cluster payload have pod security admission synchronization disabled permanently. You can enable pod security admission synchronization on other namespaces as necessary. If an Operator is installed in a user-created **openshift-*** namespace, synchronization is turned on by default after a cluster service version (CSV) is created in the namespace.

The controller examines **ServiceAccount** object permissions to use security context constraints in each namespace. Security context constraints (SCCs) are mapped to pod security profiles based on their field values; the controller uses these translated profiles. Pod security admission **warn** and **audit** labels are set to the most privileged pod security profile found in the namespace to prevent warnings and audit logging as pods are created.

Namespace labeling is based on consideration of namespace-local service account privileges.

Applying pods directly might use the SCC privileges of the user who runs the pod. However, user privileges are not considered during automatic labeling.

7.2.1. Viewing security context constraints in a namespace

You can view the security context constraints (SCC) permissions in a given namespace.

Prerequisites

- You have installed the OpenShift CLI (**oc**).

Procedure

- To view the security context constraints in your namespace, run the following command:

  ```
  oc get --show-labels namespace <namespace>
  ```

7.3. CONTROLLING POD SECURITY ADMISSION SYNCHRONIZATION

You can enable automatic pod security admission synchronization for most namespaces.

System defaults are not enforced when the **security.openshift.io/scc.podSecurityLabelSync** field is empty or set to **false**. You must set the label to **true** for synchronization to occur.

> **IMPORTANT**
>
> Namespaces that are defined as part of the cluster payload have pod security admission synchronization disabled permanently. These namespaces include:
>
> - **default**
>
> - **kube-node-lease**
>
> - **kube-system**
>
> - **kube-public**
>
> - **openshift**
>
> - All system-created namespaces that are prefixed with **openshift-**, except for **openshift-operators** By default, all namespaces that have an **openshift-** prefix are not synchronized. You can enable synchronization for any user-created **openshift-*** namespaces. You cannot enable synchronization for any system-created **openshift-*** namespaces, except for **openshift-operators**.
>
> If an Operator is installed in a user-created **openshift-*** namespace, synchronization is turned on by default after a cluster service version (CSV) is created in the namespace. The synchronized label inherits the permissions of the service accounts in the namespace.

Procedure

- To enable pod security admission label synchronization in a namespace, set the value of the **security.openshift.io/scc.podSecurityLabelSync** label to **true**.
  Run the following command:

  ```
  $ oc label namespace <namespace> security.openshift.io/scc.podSecurityLabelSync=true
  ```

> **NOTE**
>
> You can use the --overwrite flag to reverse the effects of the pod security label synchronization in a namespace.

CHAPTER 8. OPERATORS

8.1. USING OPERATORS WITH MICROSHIFT

You can use Operators with MicroShift to create applications that monitor the running services in your cluster. Operators can manage applications and their resources, such as deploying a database or message bus. As customized software running inside your cluster, Operators can be used to implement and automate common operations.

Operators offer a more localized configuration experience and integrate with Kubernetes APIs and CLI tools such as **kubectl** and **oc**. Operators are designed specifically for your applications. Operators enable you to configure components instead of modifying a global configuration file.

MicroShift applications are generally expected to be deployed in static environments. However, Operators are available if helpful in your use case. To determine the compatibility of an Operator with MicroShift, check the Operator documentation.

8.1.1. How to use Operators with MicroShift clusters

There are two ways to use Operators for your MicroShift clusters:

8.1.1.1. Manifests for Operators

Operators can be installed and managed directly by using manifests. You can use the **kustomize** configuration management tool with MicroShift to deploy an application. Use the same steps to install Operators with manifests.

- See Using Kustomize manifests to deploy applications and Using manifests example for details.

8.1.1.2. Operator Lifecycle Manager for Operators

You can also install add-on Operators to a MicroShift cluster using Operator Lifecycle Manager (OLM). OLM can be used to manage both custom Operators and Operators that are widely available. Building catalogs is required to use OLM with MicroShift.

- For details, see Using Operator Lifecycle Manager with MicroShift .

8.2. USING OPERATOR LIFECYCLE MANAGER WITH MICROSHIFT

The Operator Lifecycle Manager (OLM) package manager is used in MicroShift for installing and running optional add-on Operators.

8.2.1. Considerations for using OLM with MicroShift

- Cluster Operators as applied in OpenShift Container Platform are not used in MicroShift.

- You must create your own catalogs for the add-on Operators you want to use with your applications. Catalogs are not provided by default.

  - Each catalog must have an accessible **CatalogSource** added to a cluster, so that the OLM catalog Operator can use the catalog for content.

- You must use the CLI to conduct OLM activities with MicroShift. The console and OperatorHub GUIs are not available.

- ○ Use the Operator Package Manager **opm** CLI with network-connected clusters, or for building catalogs for custom Operators that use an internal registry.

- ○ To mirror your catalogs and Operators for disconnected or offline clusters, install the oc-mirror OpenShift CLI plugin.

> **IMPORTANT**
>
> Before using an Operator, verify with the provider that the Operator is supported on Red Hat build of MicroShift.

8.2.2. Determining your OLM installation type

You can install the OLM package manager for use with MicroShift 4.15 or newer versions. There are different ways to install OLM for MicroShift clusters, depending on your use case.

- You can install the **microshift-olm** RPM at the same time you install the MicroShift RPM on Red Hat Enterprise Linux (RHEL).

- You can install the **microshift-olm** on an existing MicroShift 4.15. Restart the MicroShift service after installing OLM for the changes to apply. See Installing the Operator Lifecycle Manager (OLM) from an RPM package.

- You can embed OLM in a Red Hat Enterprise Linux for Edge (RHEL for Edge) image. See Adding the Operator Lifecycle Manager (OLM) service to a blueprint .

8.2.3. Namespace use in MicroShift

The **microshift-olm** RPM creates the three default namespaces: one for running OLM, and two for catalog and Operator installation. You can create additional namespaces as needed for your use case.

8.2.3.1. Default namespaces

The following table lists the default namespaces and a brief description of how each namespace works.

Table 8.1. Default namespaces created by OLM for MicroShift

Default Namespace	Details
openshift-operator-lifecycle-manager	The OLM package manager runs in this namespace.
openshift-marketplace	The global namespace. Empty by default. To make the catalog source to be available globally to users in all namespaces, set the **openshift-marketplace** namespace in the catalog-source YAML.
openshift-operators	The default namespace where Operators run in MicroShift. Operators that reference catalogs in the **openshift-operators** namespace must have the AllNamespaces watch scope.

8.2.3.2. Custom namespaces

If you want to use a catalog and Operator together in a single namespace, then you must create a custom namespace. After you create the namespace, you must create the catalog in that namespace. All Operators running in the custom namespace must have the same single-namespace watch scope.

8.2.4. About building Operator catalogs

To use Operator Lifecycle Manager (OLM) with MicroShift, you must build custom Operator catalogs that you can then manage with OLM. The standard catalogs that are included with OpenShift Container Platform are not included with MicroShift.

8.2.4.1. File-based Operator catalogs

You can create catalogs for your custom Operators or filter catalogs of widely available Operators. You can combine both methods to create the catalogs needed for your specific use case. To run MicroShift with your own Operators and OLM, make a catalog by using the file-based catalog structure.

- For details, see Managing custom catalogs and Example catalog.

- See also **opm** CLI reference.

> **IMPORTANT**
>
> - When adding a catalog source to a cluster , set the **securityContextConfig** value to **restricted** in the **catalogSource.yaml** file. Ensure that your catalog can run with **restricted** permissions.

Additional resources

- **opm** CLI reference

- About Operator catalogs

- To create file-based catalogs by using the **opm** CLI, see Managing custom catalogs

8.2.5. How to deploy Operators using OLM

After you create and deploy your custom catalog, you must create a Subscription custom resource (CR) that can access the catalog and install the Operators you choose. Where Operators run depends on the namespace in which you create the Subscription CR.

> **IMPORTANT**
>
> Operators in OLM have a watch scope. For example, some Operators only support watching their own namespace, while others support watching every namespace in the cluster. All Operators installed in a given namespace must have the same watch scope.

8.2.5.1. Connectivity and OLM Operator deployment

Operators can be deployed anywhere a catalog is running.

- For clusters that are connected to the internet, mirroring images is not required. Images can be pulled over the network.

- For restricted networks in which MicroShift has access to an internal network only, images must be mirrored to an internal registry.

- For use cases in which MicroShift clusters are completely offline, all images must be embedded into an **osbuild** blueprint.

Additional resources

- [Operator group membership](#)

8.2.5.2. Adding OLM-based Operators to a networked cluster using the global namespace

To deploy different operators to different namespaces, use this procedure. For MicroShift clusters that have network connectivity, Operator Lifecycle Manager (OLM) can access sources hosted on remote registries. The following procedure lists the basic steps of using configuration files to install an Operator that uses the global namespace.

> **NOTE**
>
> To use an Operator installed in a different namespace, or in more than one namespace, make sure that the catalog source and the Subscription CR that references the Operator are running in the **openshift-marketplace** namespace.

Prerequisites

- The OpenShift CLI (**oc**) is installed.

- Operator Lifecycle Manager (OLM) is installed.

- You have created a custom catalog in the global namespace.

Procedure

1. Confirm that OLM is running by using the following command:

   ```
   $ oc -n openshift-operator-lifecycle-manager get pod -l app=olm-operator
   ```

   Example output

   ```
   NAME                        READY   STATUS    RESTARTS   AGE
   olm-operator-85b5c6786-n6kbc   1/1     Running   0          2m24s
   ```

2. Confirm that the OLM catalog Operator is running by using the following command:

   ```
   $ oc -n openshift-operator-lifecycle-manager get pod -l app=catalog-operator
   ```

   Example output

   ```
   NAME                        READY   STATUS    RESTARTS   AGE
   catalog-operator-5fc7f857b6-tj8cf   1/1     Running   0          2m33s
   ```

> **NOTE**
>
> The following steps assume you are using the global namespace, **openshift-marketplace**. The catalog must run in the same namespace as the Operator. The Operator must support the **AllNamespaces** mode.

1. Create the **CatalogSource** object by using the following example YAML:

   Example catalog source YAML

   ```
   apiVersion: operators.coreos.com/v1alpha1
   kind: CatalogSource
   metadata:
     name: operatorhubio-catalog
     namespace: openshift-marketplace  ❶
   spec:
     sourceType: grpc
     image: quay.io/operatorhubio/catalog:latest
     displayName: Community Operators  ❷
     publisher: OperatorHub.io
     grpcPodConfig:
       securityContextConfig: restricted  ❸
     updateStrategy:
       registryPoll:
         interval: 60m
   ```

   ❶ The global namespace. Setting the **metadata.namespace** to **openshift-marketplace** enables the catalog to run in all namespaces. Subscriptions in any namespace can reference catalogs created in the **openshift-marketplace** namespace.

   ❷ Community Operators are not installed by default with OLM for MicroShift. Listed here for example only.

   ❸ The value of **securityContextConfig** must be set to **restricted** for MicroShift.

2. Apply the **CatalogSource** configuration by running the following command:

   ```
   $ oc apply -f <my-catalog-source.yaml>  ❶
   ```

   ❶ Replace **<my-catalog-source.yaml>** with your catalog source configuration file name. In this example, **catalogsource.yaml** is used.

   Example output

   ```
   catalogsource.operators.coreos.com/operatorhubio-catalog created
   ```

3. To verify that the catalog source is applied, check for the **READY** state by using the following command:

   ```
   $ oc describe catalogsources.operators.coreos.com -n openshift-marketplace operatorhubio-catalog
   ```

Example output

```
Name:         operatorhubio-catalog
Namespace:    openshift-marketplace
Labels:       <none>
Annotations: <none>
API Version: operators.coreos.com/v1alpha1
Kind:         CatalogSource
Metadata:
  Creation Timestamp:  2024-01-31T09:55:31Z
  Generation:          1
  Resource Version:    1212
  UID:                 4edc1a96-83cd-4de9-ac8c-c269ca895f3e
Spec:
  Display Name:  Community Operators
  Grpc Pod Config:
    Security Context Config:  restricted
  Image:                quay.io/operatorhubio/catalog:latest
  Publisher:            OperatorHub.io
  Source Type:          grpc
  Update Strategy:
    Registry Poll:
      Interval: 60m
Status:
  Connection State:
    Address:            operatorhubio-catalog.openshift-marketplace.svc:50051
    Last Connect:       2024-01-31T09:55:57Z
    Last Observed State:  READY    1
  Registry Service:
    Created At:         2024-01-31T09:55:31Z
    Port:           50051
    Protocol:         grpc
    Service Name:       operatorhubio-catalog
    Service Namespace:  openshift-marketplace
Events:               <none>
```

1　　The status is reported as **READY**.

4. Confirm that the catalog source is running by using the following command:

```
$ oc get pods -n openshift-marketplace -l olm.catalogSource=operatorhubio-catalog
```

Example output

```
NAME                     READY  STATUS   RESTARTS  AGE
operatorhubio-catalog-x24nh  1/1    Running  0         59s
```

5. Create a Subscription CR configuration file by using the following example YAML:

Example Subscription custom resource YAML

```
apiVersion: operators.coreos.com/v1alpha1
kind: Subscription
```

```
metadata:
  name: my-cert-manager
  namespace: openshift-operators
spec:
  channel: stable
  name: cert-manager
  source: operatorhubio-catalog
  sourceNamespace: openshift-marketplace ❶
```

❶ The global namespace. Setting the **sourceNamespace** value to **openshift-marketplace** enables Operators to run in multiple namespaces if the catalog also runs in the **openshift-marketplace** namespace.

6. Apply the Subscription CR configuration by running the following command:

```
$ oc apply -f <my-subscription-cr.yaml> ❶
```

❶ Replace **<my-subscription-cr.yaml>** with your Subscription CR filename. In this example, **sub.yaml** is used.

Example output

```
subscription.operators.coreos.com/my-cert-manager created
```

7. You can create a configuration file for the specific Operand you want to use and apply it now.

Verification

1. Verify that your Operator is running by using the following command:

```
$ oc get pods -n openshift-operators ❶
```

❶ The namespace from the Subscription CR is used.

> **NOTE**
>
> Allow a minute or two for the Operator start.

Example output

```
NAME                                READY  STATUS   RESTARTS  AGE
cert-manager-7df8994ddb-4vrkr       1/1    Running  0         19s
cert-manager-cainjector-5746db8fd7-69442  1/1   Running  0     18s
cert-manager-webhook-f858bf58b-748nt  1/1   Running  0       18s
```

8.2.5.3. Adding OLM-based Operators to a networked cluster in a specific namespace

Use this procedure if you want to specify a namespace for an Operator, for example, **olm-microshift**. In this example, the catalog is scoped and available in the global **openshift-marketplace** namespace. The Operator uses content from the global namespace, but runs only in the **olm-microshift** namespace. For

MicroShift clusters that have network connectivity, Operator Lifecycle Manager (OLM) can access sources hosted on remote registries.

> **IMPORTANT**
>
> All of the Operators installed in a specific namespace must have the same watch scope. In this case, the watch scope is **OwnNamespace**.

Prerequisites

- The OpenShift CLI (**oc**) is installed.

- Operator Lifecycle Manager (OLM) is installed.

- You have created a custom catalog that is running in the global namespace.

Procedure

1. Confirm that OLM is running by using the following command:

   ```
   $ oc -n openshift-operator-lifecycle-manager get pod -l app=olm-operator
   ```

   Example output

   ```
   NAME                     READY  STATUS   RESTARTS  AGE
   olm-operator-85b5c6786-n6kbc  1/1    Running  0         16m
   ```

2. Confirm that the OLM catalog Operator is running by using the following command:

   ```
   $ oc -n openshift-operator-lifecycle-manager get pod -l app=catalog-operator
   ```

   Example output

   ```
   NAME                     READY  STATUS   RESTARTS  AGE
   catalog-operator-5fc7f857b6-tj8cf  1/1    Running  0         16m
   ```

3. Create a namespace by using the following example YAML:

   Example namespace YAML

   ```
   apiVersion: v1
   kind: Namespace
   metadata:
     name: olm-microshift
   ```

4. Apply the namespace configuration using the following command:

   ```
   $ oc apply -f _<ns.yaml>_  ❶
   ```

   ❶ Replace *<ns.yaml>* with the name of your namespace configuration file. In this example, **olm-microshift** is used.

Example output

> namespace/olm-microshift created

5. Create the Operator group YAML by using the following example YAML:

Example Operator group YAML

```
kind: OperatorGroup
apiVersion: operators.coreos.com/v1
metadata:
  name: og
  namespace: olm-microshift
spec: 1
  targetNamespaces:
  - olm-microshift
```

[1] For Operators using the global namespace, omit the **spec.targetNamespaces** field and values.

6. Apply the Operator group configuration by running the following command:

> $ oc apply -f _<og.yaml>_ 1

[1] Replace *<og.yaml>* with the name of your operator group configuration file.

Example output

> operatorgroup.operators.coreos.com/og created

7. Create the **CatalogSource** object by using the following example YAML:

Example catalog source YAML

```
apiVersion: operators.coreos.com/v1alpha1
kind: CatalogSource
metadata:
  name: operatorhubio-catalog
  namespace: openshift-marketplace 1
spec:
  sourceType: grpc
  image: quay.io/operatorhubio/catalog:latest
  displayName: Community Operators 2
  publisher: OperatorHub.io
  grpcPodConfig:
    securityContextConfig: restricted 3
  updateStrategy:
    registryPoll:
      interval: 60m
```

1 The global namespace. Setting the **metadata.namespace** to **openshift-marketplace** enables the catalog to run in all namespaces. Subscriptions CRs in any namespace can

2 Community Operators are not installed by default with OLM for MicroShift. Listed here for example only.

3 The value of **securityContextConfig** must be set to **restricted** for MicroShift.

8. Apply the **CatalogSource** configuration by running the following command:

```
$ oc apply -f _<my-catalog-source.yaml>_ ❶
```

1 Replace *<my-catalog-source.yaml>* with your catalog source configuration file name.

9. To verify that the catalog source is applied, check for the **READY** state by using the following command:

```
$ oc describe catalogsources.operators.coreos.com -n openshift-marketplace operatorhubio-catalog
```

Example output

```
Name:         operatorhubio-catalog
Namespace:    openshift-marketplace
Labels:       <none>
Annotations:  <none>
API Version:  operators.coreos.com/v1alpha1
Kind:         CatalogSource
Metadata:
  Creation Timestamp:  2024-01-31T10:09:46Z
  Generation:        1
  Resource Version:    2811
  UID:               60ce4a36-86d3-4921-b9fc-84d67c28df48
Spec:
  Display Name:  Community Operators
  Grpc Pod Config:
    Security Context Config:  restricted
  Image:                 quay.io/operatorhubio/catalog:latest
  Publisher:              OperatorHub.io
  Source Type:            grpc
  Update Strategy:
    Registry Poll:
      Interval:  60m
Status:
  Connection State:
    Address:            operatorhubio-catalog.openshift-marketplace.svc:50051
    Last Connect:       2024-01-31T10:10:04Z
    Last Observed State:  READY ❶
  Registry Service:
    Created At:       2024-01-31T10:09:46Z
    Port:             50051
    Protocol:         grpc
```

```
    Service Name:       operatorhubio-catalog
    Service Namespace:  openshift-marketplace
Events:                 <none>
```

1   The status is reported as **READY**.

10. Confirm that the catalog source is running by using the following command:

```
$ oc get pods -n openshift-marketplace -l olm.catalogSource=operatorhubio-catalog
```

Example output

```
NAME                         READY  STATUS   RESTARTS  AGE
operatorhubio-catalog-j7sc8  1/1    Running  0         43s
```

11. Create a Subscription CR configuration file by using the following example YAML:

Example Subscription custom resource YAML

```
apiVersion: operators.coreos.com/v1alpha1
kind: Subscription
metadata:
  name: my-gitlab-operator-kubernetes
  namespace: olm-microshift 1
spec:
  channel: stable
  name: gitlab-operator-kubernetes
  source: operatorhubio-catalog
  sourceNamespace: openshift-marketplace 2
```

1   The specific namespace. Operators reference the global namespace for content, but run in the **olm-microshift** namespace.

2   The global namespace. Subscriptions CRs in any namespace can reference catalogs created in the **openshift-marketplace** namespace.

12. Apply the Subscription CR configuration by running the following command:

```
$ oc apply -f _<my-subscription-cr.yaml>_
```

Example output

```
subscription.operators.coreos.com/my-gitlab-operator-kubernetes
```

13. You can create a configuration file for the specific Operand you want to use and apply it now.

Verification

1. Verify that your Operator is running by using the following command:

```
$ oc get pods -n olm-microshift 1
```

1 The namespace from the Subscription CR is used.

> **NOTE**
>
> Allow a minute or two for the Operator start.

Example output

```
NAME                                    READY  STATUS   RESTARTS  AGE
gitlab-controller-manager-69bb6df7d6-g7ntx  2/2    Running  0         3m24s
```

Additional resources

- [Updating installed Operators](#)

- [Deleting Operators from a cluster using the CLI](#)

8.3. CREATING CUSTOM CATALOGS USING THE OC-MIRROR PLUGIN

You can create custom catalogs with widely available Operators and mirror them by using the oc-mirror OpenShift CLI (oc) plugin.

8.3.1. Using Red Hat-provided Operator catalogs and mirror registries

You can filter and prune catalogs to get specific Operators and mirror them by using the oc-mirror OpenShift CLI (oc) plugin. You can also use Operators in disconnected settings or embedded in Red Hat Enterprise Linux for Edge (RHEL for Edge) images. To read more details about how to configure your systems for mirroring, use the links in the following "Additional resources" section. If you are ready to deploy Operators from Red Hat-provided Operator catalogs, mirror them, or to embed them in RHEL for Edge images, start with the following section, "Inspecting catalog contents by using the oc-mirror plugin."

Additional resources

- [Using Operator Lifecycle Manager on restricted networks](#)

- [Configuring hosts for mirror registry access](#)

- [Configuring network settings for fully disconnected hosts](#)

- [Getting the mirror registry container image list](#)

- [Embedding in a RHEL for Edge image for offline use](#)

8.3.2. About the oc-mirror plugin for creating a mirror registry

You can use the oc-mirror OpenShift CLI (oc) plugin with MicroShift to filter and prune Operator catalogs. You can then mirror the filtered catalog contents to a mirror registry or use the container images in disconnected or offline deployments with RHEL for Edge.

> **NOTE**
>
> MicroShift uses the generally available version (1) of the oc-mirror plugin. Do not use the following procedures with the Technical Preview version (2) of oc-mirror plugin.

You can mirror the container images required by the desired Operators locally or to a container mirror registry that supports Docker v2-2, such as Red Hat Quay. The procedure to mirror content from Red Hat-hosted registries connected to the internet to a disconnected image registry is the same, independent of the registry you choose. After you mirror the contents of your catalog, configure each cluster to retrieve this content from your mirror registry.

8.3.2.1. Connectivity considerations when populating a mirror registry

When you populate your registry, you can use one of following connectivity scenarios:

Connected mirroring

If you have a host that can access both the internet and your mirror registry, but not your cluster node, you can directly mirror the content from that machine.

Disconnected mirroring

If you do not have a host that can access both the internet and your mirror registry, you must mirror the images to a file system and then bring that host or removable media into your disconnected environment.

> **IMPORTANT**
>
> A container registry must be reachable by every machine in the clusters that you provision. Installing, updating, and other operations, such as relocating workloads, might fail if the registry is unreachable.

To avoid problems caused by an unreachable registry, use the following standard practices:

- Run mirror registries in a highly available way.

- Ensure that the mirror registry at least matches the production availability of your clusters.

Additional resources

- Installing the oc mirror plugin

8.3.2.2. Inspecting catalog contents by using the oc-mirror plugin

Use the following example procedure to select a catalog and list Operators from available OpenShift Container Platform content to add to your oc-mirror plugin image set configuration file.

> **NOTE**
>
> If you use your own catalogs and Operators, you can push the images directly to your internal registry.

Prerequisites

- The OpenShift CLI (**oc**) is installed.

- Operator Lifecycle Manager (OLM) is installed.

- The oc-mirror OpenShift CLI (oc) plugin is installed.

Procedure

1. Get a list of available Red Hat-provided Operator catalogs to filter by running the following command:

   $ oc mirror list operators --version 4.15 --catalogs

2. Get a list of Operators in the Red Hat Operators catalog by running the following command:

   $ oc mirror list operators <--catalog=<catalog_source>> **1**

   1 Specifies your catalog source, such as **registry.redhat.io/redhat/redhat-operator-index:v4.15** or **quay.io/operatorhubio/catalog:latest**.

3. Select an Operator. For this example, **amq-broker-rhel8** is selected.

4. Optional: To inspect the channels and versions of the Operator you want to filter, enter the following commands:

   a. Get a list of channels by running the following command:

      $ oc mirror list operators --catalog=registry.redhat.io/redhat/redhat-operator-index:v4.15 --package=amq-broker-rhel8

   b. Get a list of versions within a channel by running the following command:

      $ oc mirror list operators --catalog=registry.redhat.io/redhat/redhat-operator-index:v4.15 --package=amq-broker-rhel8 --channel=7.11.x

Next steps

- Create and edit an image set configuration file using the information gathered in this procedure.

- Mirror the images from the transformed image set configuration file to a mirror registry or disk.

8.3.2.3. Creating an image set configuration file

You must create an image set configuration file to mirror catalog contents with the oc-mirror plugin. The image set configuration file defines which Operators to mirror along with other configuration settings for the oc-mirror plugin. After generating a default image set file, you must edit the contents so that remaining entries are compatible with both MicroShift and the Operator you plan to use.

You must specify a storage backend in the image set configuration file. This storage backend can be a local directory or a registry that supports Docker v2-2. The oc-mirror plugin stores metadata in this storage backend during image set creation.

> **IMPORTANT**
>
> Do not delete or modify the metadata that is generated by the oc-mirror plugin. You must use the same storage backend every time you run the oc-mirror plugin for the same mirror registry.

Prerequisites

- You have created a container image registry credentials file. See Configuring credentials that allow images to be mirrored.

Procedure

1. Use the **oc mirror init** command to create a template for the image set configuration and save it to a file called **imageset-config.yaml**:

   ```
   $ oc mirror init <--registry <storage_backend> > imageset-config.yaml ❶
   ```

   ❶ Specifies the location of your storage backend, such as **example.com/mirror/oc-mirror-metadata**.

   Example default image set configuration file

   ```
   kind: ImageSetConfiguration
   apiVersion: mirror.openshift.io/v1alpha2
   storageConfig:
     registry:
       imageURL: registry.example.com/oc-mirror
       skipTLS: false
   mirror:
    platform: ❶
      channels:
      - name: stable-4.15
        type: ocp
   operators:
   - catalog: registry.redhat.io/redhat/redhat-operator-index:v4.15
     packages:
     - name: serverless-operator
       channels:
       - name: stable
   additionalImages: ❷
   - name: registry.redhat.io/ubi8/ubi:latest
   helm: {} ❸
   ```

   ❶ The **platform** field and related fields are not supported by MicroShift and must be deleted.

   ❷ Specify any additional images to include in the image set. If you do not need to specify additional images, delete this field.

   ❸ Helm is not supported by MicroShift, and must be deleted.

2. Edit the values of your image set configuration file to meet the requirements of both MicroShift and the Operator you want to mirror, like the following example:

Example edited MicroShift image set configuration file

```
kind: ImageSetConfiguration
apiVersion: mirror.openshift.io/v1alpha2
storageConfig: 1
  registry:
    imageURL: <storage_backend> 2
    skipTLS: false
mirror:
  operators:
  - catalog: registry.redhat.io/redhat/redhat-operator-index:v4.15 3
    packages:
    - name: amq-broker-rhel8 4
      channels:
      - name: 7.11.x 5
```

1 Set the backend location where the image set metadata is saved. This location can be a registry or local directory. It is required to specify **storageConfig** values.

2 Set the registry URL for the storage backend, such as **<example.com/mirror/oc-mirror-metadata>**.

3 Set the Operator catalog to retrieve images from.

4 Specify the Operator packages to include in the image set. Remove this field to retrieve all packages in the catalog.

5 Specify only certain channels of the Operator packages to include in the image set. You must always include the default channel for the Operator package even if you do not use the bundles in that channel. You can find the default channel by running the following command: **oc mirror list operators --catalog=<catalog_name> --package=<package_name>**.

3. Save the updated file.

Next steps

- Use the oc-mirror plugin to mirror an image set directly to a target mirror registry.

- Configure CRI-O.

- Apply the catalog sources to your clusters.

8.3.2.3.1. Image set configuration parameters

The oc-mirror plugin requires an image set configuration file that defines what images to mirror. The following table lists the available parameters for the **ImageSetConfiguration** resource.

Table 8.2. **ImageSetConfiguration** parameters

Parameter	Description	Values

Parameter	Description	Values
apiVersion	The API version for the **ImageSetConfiguration** content.	String. For example: **mirror.openshift.io/v1alpha2**.
mirror	The configuration of the image set.	Object
mirror.additionalImages	The additional images configuration of the image set.	Array of objects. For example: additionalImages: - name: registry.redhat.io/ubi8/ubi:latest
mirror.additionalImages.name	The tag or digest of the image to mirror.	String. For example: **registry.redhat.io/ubi8/ubi:latest**
mirror.blockedImages	The full tag, digest, or pattern of images to block from mirroring.	Array of strings. For example: **docker.io/library/alpine**

mirror.operators	The Operators configuration of the image set.	Array of objects. For example:
		
		operators:
		  - catalog: registry.redhat.io/redhat/redhat-operator-index:v4.15
		    packages:
		      - name: elasticsearch-operator
		
		minVersion: '2.4.0'

Parameter	Description	Values
mirror.operators.catalog	The Operator catalog to include in the image set.	String. For example: **registry.redhat.io/redhat/redhat-operator-index:v4.15**.
mirror.operators.full	When **true**, downloads the full catalog, Operator package, or Operator channel.	Boolean. The default value is **false**.

mirror.operators.packages	The Operator packages configuration.	Array of objects. For example:
		
		operators:
		 - catalog: registry.redhat.io/redhat/redhat-operator-index:v4.15
		   packages:
		    - name: elasticsearch-operator
		
		minVersion: '5.2.3-31'

mirror.operators.packages.name	The Operator package name to include in the image set	String. For example: **elasticsearch-operator**.
mirror.operators.packages.channels	The Operator package channel configuration.	Object
mirror.operators.packages.channels.name	The Operator channel name, unique within a package, to include in the image set.	String. For example: **fast** or **stable-v4.15**.
mirror.operators.packages.channels.maxVersion	The highest version of the Operator mirror across all channels in which it exists. See the following note for further information.	String. For example: **5.2.3-31**

Parameter	Description	Values
mirror.operators.packages.channels.minBundle	The name of the minimum bundle to include, plus all bundles in the update graph to the channel head. Set this field only if the named bundle has no semantic version metadata.	String. For example: **bundleName**
mirror.operators.packages.channels.minVersion	The lowest version of the Operator to mirror across all channels in which it exists. See the following note for further information.	String. For example: **5.2.3-31**
mirror.operators.packages.maxVersion	The highest version of the Operator to mirror across all channels in which it exists. See the following note for further information.	String. For example: **5.2.3-31**.
mirror.operators.packages.minVersion	The lowest version of the Operator to mirror across all channels in which it exists. See the following note for further information.	String. For example: **5.2.3-31**.
mirror.operators.skipDependencies	If **true**, dependencies of bundles are not included.	Boolean. The default value is **false**.
mirror.operators.targetCatalog	An alternative name and optional namespace hierarchy to mirror the referenced catalog as.	String. For example: **my-namespace/my-operator-catalog**
mirror.operators.targetName	An alternative name to mirror the referenced catalog as. The **targetName** parameter is deprecated. Use the **targetCatalog** parameter instead.	String. For example: **my-operator-catalog**
mirror.operators.targetTag	An alternative tag to append to the **targetName** or **targetCatalog**.	String. For example: **v1**
storageConfig	The back-end configuration of the image set.	Object
storageConfig.local	The local back-end configuration of the image set.	Object

Parameter	Description	Values
storageConfig.local.path	The path of the directory to contain the image set metadata.	String. For example: **./path/to/dir/**.
storageConfig.registry	The registry back-end configuration of the image set.	Object
storageConfig.registry.imageURL	The back-end registry URI. Can optionally include a namespace reference in the URI.	String. For example: **quay.io/myuser/imageset:metadata**.
storageConfig.registry.skipTLS	Optionally skip TLS verification of the referenced back-end registry.	Boolean. The default value is **false**.

NOTE

Using the **minVersion** and **maxVersion** properties to filter for a specific Operator version range can result in a multiple channel heads error. The error message states that there are **multiple channel heads**. This is because when the filter is applied, the update graph of the Operator is truncated.

Operator Lifecycle Manager requires that every Operator channel contains versions that form an update graph with exactly one end point, that is, the latest version of the Operator. When the filter range is applied, that graph can turn into two or more separate graphs or a graph that has more than one end point.

To avoid this error, do not filter out the latest version of an Operator. If you still run into the error, depending on the Operator, either the **maxVersion** property must be increased or the **minVersion** property must be decreased. Because every Operator graph can be different, you might need to adjust these values until the error resolves.

Additional resources

- Imageset configuration examples

8.3.2.4. Mirroring from mirror to mirror

You can use the oc-mirror plugin to mirror an image set directly to a target mirror registry that is accessible during image set creation.

You are required to specify a storage backend in the image set configuration file. This storage backend can be a local directory or a Docker v2 registry. The oc-mirror plugin stores metadata in this storage backend during image set creation.

IMPORTANT

Do not delete or modify the metadata that is generated by the oc-mirror plugin. You must use the same storage backend every time you run the oc-mirror plugin for the same mirror registry.

Prerequisites

- You have access to the internet to get the necessary container images.

- You have installed the OpenShift CLI (**oc**).

- You have installed the **oc-mirror** CLI plugin.

- You have created the image set configuration file.

Procedure

- Run the **oc mirror** command to mirror the images from the specified image set configuration to a specified registry:

  ```
  $ oc mirror --config=./<imageset-config.yaml> \ ①
      docker://registry.example:5000          ②
  ```

  ① Specify the image set configuration file that you created. For example, **imageset-config.yaml**.

  ② Specify the registry to mirror the image set file to. The registry must start with **docker:**//. If you specify a top-level namespace for the mirror registry, you must also use this same namespace on subsequent executions.

Example output

```
Rendering catalog image "registry.example.com/redhat/redhat-operator-index:v{ocp-version}" with file-based catalog
```

Verification

1. Navigate into the **oc-mirror-workspace/** directory that was generated.

2. Navigate into the results directory, for example, **results-1639608409/**.

3. Verify that YAML files are present for the **ImageContentSourcePolicy** and **CatalogSource** resources.

IMPORTANT

The **ImageContentSourcePolicy** YAML file is used as reference content for manual configuration of CRI-O in MicroShift. You cannot apply the resource directly into a MicroShift cluster.

Next steps

- Convert the **ImageContentSourcePolicy** YAML content for use in manually configuring CRI-O.

- If required, mirror the images from mirror to disk for disconnected or offline use.

- Configure your cluster to use the resources generated by oc-mirror.

Troubleshooting

- [Unable to retrieve source image](#) .

Additional resources

- [Mirroring an image set in a partially disconnected environment](#)

- [Mirroring an image set in a fully disconnected environment](#)

8.3.2.5. Configuring CRI-O for using a registry mirror for Operators

You must transform the **imageContentSourcePolicy.yaml** file created with the oc-mirror plugin into a format that is compatible with the CRI-O container runtime configuration used by MicroShift.

Prerequisites

- The OpenShift CLI (**oc**) is installed.

- Operator Lifecycle Manager (OLM) is installed.

- The oc-mirror OpenShift CLI (oc) plugin is installed.

- The **yq** binary is installed.

- **ImageContentSourcePolicy** and **CatalogSource** YAML files are available in the **oc-mirror-workspace/results-*** directory.

Procedure

1. Confirm the contents of the **imageContentSourcePolicy.yaml** file by running the following command:

   ```
   $ cat oc-mirror-workspace/<results-directory>/imageContentSourcePolicy.yaml ❶
   ```

   ❶ Specify the **results** directory name, such as **<results-1707148826>**.

   ## Example output

   ```
   apiVersion: operator.openshift.io/v1alpha1
   kind: ImageContentSourcePolicy
   metadata:
     labels:
       operators.openshift.org/catalog: "true"
     name: operator-0
   spec:
     repositoryDigestMirrors:
   ```

```
- mirrors:
  - registry.<example.com>/amq7
  source: registry.redhat.io/amq7
```

2. Transform the **imageContentSourcePolicy.yaml** into a format ready for CRI-O configuration by running the following command:

```
yq '.spec.repositoryDigestMirrors[] as $item ireduce([]; . + [{"mirror": $item.mirrors[], "source":
($item | .source)}]) | .[] |
  "[[registry]]
     prefix = \"" + .source + "\"
     location = \"" + .mirror + "\"
     mirror-by-digest-only = true
     insecure = true
     "' ./icsp.yaml
```

Example output

```
[[registry]]
     prefix = "registry.redhat.io/amq7"
     location = "registry.example.com/amq7"
     mirror-by-digest-only = true
     insecure = true
```

3. Add the output to the CRI-O configuration file in the **/etc/containers/registries.conf.d/** directory:

Example crio-config.yaml mirror configuration file

```
[[registry]]
     prefix = "registry.redhat.io/amq7"
     location = "registry.example.com/amq7"
     mirror-by-digest-only = true
     insecure = true

[[registry]]
     prefix = ""
     location = "quay.io"
     mirror-by-digest-only = true
[[registry.mirror]]
     location = "<registry_host>:<port>"    1
     insecure = false
```

1 Specify the host name and port of your mirror registry server, for example **microshift-quay:8443**.

4. Apply the CRI-O configuration changes by restarting MicroShift with the following command:

```
$ sudo systemctl restart crio
```

8.3.2.6. Installing a custom catalog created with the oc-mirror plugin

After you mirror your image set to the mirror registry, you must apply the generated **CatalogSource** custom resource (CR) into the cluster. The **CatalogSource** CR is used by Operator Lifecycle Manager (OLM) to retrieve information about the available Operators in the mirror registry. You must then create and apply a subscription CR to subscribe to your custom catalog.

Prerequisites

- You mirrored the image set to your registry mirror.

- You added image reference information to the CRI-O container runtime configuration.

Procedure

1. Apply the catalog source configuration file from the results directory to create the catalog source object by running the following command:

   ```
   $ oc apply -f ./oc-mirror-workspace/results-1708508014/catalogSource-cs-redhat-operator-index.yaml
   ```

   Example catalog source configuration file

   ```
   apiVersion: operators.coreos.com/v1alpha1
   kind: CatalogSource
   metadata:
     name: redhat-catalog
     namespace: openshift-marketplace ❶
   spec:
     sourceType: grpc
     image: registry.example.com/redhat/redhat-operator-index:v4.15
     updateStrategy:
       registryPoll:
         interval: 60m
   ```

   ❶ Specifies the global namespace. Setting the **metadata.namespace** to **openshift-marketplace** enables the catalog to reference catalogs in all namespaces. Subscriptions in any namespace can reference catalogs created in the **openshift-marketplace** namespace.

   Example output

   ```
   catalogsource.operators.coreos.com/cs-redhat-operator-index created
   ```

2. Verify that the **CatalogSource** resources were successfully installed by running the following command:

   ```
   $ oc get catalogsource --all-namespaces
   ```

3. Verify that the catalog source is running by using the following command:

   ```
   $ oc get pods -n openshift-marketplace
   ```

   Example output

```
NAME                        READY   STATUS    RESTARTS   AGE
cs-redhat-operator-index-4227b   2/2     Running   0          2m5s
```

4. Create a **Subscription** CR, similar to the following example:

Example Subscription CR

```
apiVersion: operators.coreos.com/v1alpha1
kind: Subscription
metadata:
  name: amq-broker
  namespace: openshift-operators
spec:
  channel: 7.11.x
  name: amq-broker-rhel8
  source: cs-redhat-operator-index
  sourceNamespace: openshift-marketplace
```

5. Apply the Subscription CR configuration by running the following command:

```
$ oc apply -f ./<my-subscription-cr.yaml> 1
```

1 Specify the name of your subscription, such as **my-subscription-cr.yaml**.

Example output

```
subscription.operators.coreos.com/amq-broker created
```

8.4. ADDING OLM-BASED OPERATORS TO A DISCONNECTED CLUSTER

You can use OLM-based Operators in disconnected situations by embedding them in a Red Hat Enterprise Linux for Edge (RHEL for Edge) image.

8.4.1. About adding OLM-based Operators to a disconnected cluster

For Operators that are installed on disconnected clusters, Operator Lifecycle Manager (OLM) by default cannot access sources hosted on remote registries because those remote sources require full internet connectivity. Therefore, you must mirror the remote registries to a highly available container registry.

The following steps are required to use OLM-based Operators in disconnected situations:

- Include OLM in the container image list for your mirror registry.

- Configure the system to use your mirror registry by updating your CRI-O configuration directly. **ImageContentSourcePolicy** is not supported in MicroShift.

- Add a **CatalogSource** object to the cluster so that the OLM catalog Operator can use the local catalog on the mirror registry.

- Ensure that MicroShift is installed to run in a disconnected capacity.

- Ensure that the network settings are configured to run in disconnected mode.

After enabling OLM in a disconnected cluster, you can continue to use your internet-connected workstation to keep your local catalog sources updated as newer versions of Operators are released.

Additional resources

- [Creating the RHEL for Edge image](#)

- [Embedding in a RHEL for Edge image for offline use](#)

- [Configuring network settings for fully disconnected hosts](#)

8.4.1.1. Performing a dry run

You can use oc-mirror to perform a dry run, without actually mirroring any images. This allows you to review the list of images that would be mirrored, as well as any images that would be pruned from the mirror registry. A dry run also allows you to catch any errors with your image set configuration early or use the generated list of images with other tools to carry out the mirroring operation.

Prerequisites

- You have access to the internet to obtain the necessary container images.

- You have installed the OpenShift CLI (**oc**).

- You have installed the **oc-mirror** CLI plugin.

- You have created the image set configuration file.

Procedure

1. Run the **oc mirror** command with the **--dry-run** flag to perform a dry run:

   ```
   $ oc mirror --config=./imageset-config.yaml \    ❶
     docker://registry.example:5000            \    ❷
     --dry-run                                       ❸
   ```

   ❶ Pass in the image set configuration file that was created. This procedure assumes that it is named **imageset-config.yaml**.

   ❷ Specify the mirror registry. Nothing is mirrored to this registry as long as you use the **--dry-run** flag.

   ❸ Use the **--dry-run** flag to generate the dry run artifacts and not an actual image set file.

   ### Example output

   ```
   Checking push permissions for registry.example:5000
   Creating directory: oc-mirror-workspace/src/publish
   Creating directory: oc-mirror-workspace/src/v2
   Creating directory: oc-mirror-workspace/src/charts
   Creating directory: oc-mirror-workspace/src/release-signatures
   No metadata detected, creating new workspace
   ```

> wrote mirroring manifests to oc-mirror-workspace/operators.1658342351/manifests-redhat-operator-index
>
> ...
>
> info: Planning completed in 31.48s
> info: Dry run complete
> Writing image mapping to oc-mirror-workspace/mapping.txt

2. Navigate into the workspace directory that was generated:

   > $ cd oc-mirror-workspace/

3. Review the **mapping.txt** file that was generated.
   This file contains a list of all images that would be mirrored.

4. Review the **pruning-plan.json** file that was generated.
   This file contains a list of all images that would be pruned from the mirror registry when the image set is published.

   > **NOTE**
   >
   > The **pruning-plan.json** file is only generated if your oc-mirror command points to your mirror registry and there are images to be pruned.

8.4.1.2. Getting catalogs and Operator container image references to use with RHEL for Edge in disconnected environments

After performing a dry run with the oc-mirror plugin to review the list of images that you want to mirror, you must get all of the container image references, then format the output for adding to an Image Builder blueprint.

> **NOTE**
>
> For catalogs made for proprietary Operators, you can format image references for the Image Builder blueprint without using the following procedure.

Prerequisites

- You have a catalog index for the Operators you want to use.

- You have installed the **jq** CLI tool.

- You are familiar with Image Builder blueprint files.

- You have an Image Builder blueprint TOML file.

Procedure

1. Parse the catalog **index.json** file to get the image references that you need to include in the Image Builder blueprint. You can use either the unfiltered catalog or you can filter out images that cannot be mirrored:

a. Parse the unfiltered catalog **index.json** file to get the image references by running the following command:

```
jq -r --slurp '.[] | select(.relatedImages != null) | "[[containers]]\nsource = \"" +
.relatedImages[].image + "\"\n"'   ./oc-mirror-
workspace/src/catalogs/registry.redhat.io/redhat/redhat-operator-
index/v4.15/index/index.json
```

b. If you want to filter out images that cannot be mirrored, filter and parse the catalog **index.json** file by running the following command:

```
$ jq -r --slurp '.[] | select(.relatedImages != null) | .relatedImages[] | select(.name |
contains("ppc") or contains("s390x") | not) | "[[containers]]\\nsource = \\"" + .image +
"\\"\\n"' ./oc-mirror-workspace/src/catalogs/registry.redhat.io/redhat/redhat-operator-
index/v4.15/index/index.json
```

> **NOTE**
>
> This step uses the AMQ Broker Operator as an example. You can add other criteria to the **jq** command for further filtering as required by your use case.

Example image-reference output

```
[[containers]]
source = "registry.redhat.io/amq7/amq-broker-init-
rhel8@sha256:0b2126cfb6054fdf428c1f43b69e36e93a09a49ce15350e9273c98cc08c6598
b"

[[containers]]
source = "registry.redhat.io/amq7/amq-broker-init-
rhel8@sha256:0dde839c2dce7cb684094bf26523c8e16677de03149a0fff468b8c3f106e1f4f
"

...
...

[[containers]]
source = "registry.redhat.io/amq7/amq-broker-
rhel8@sha256:e8fa2a00e576ecb95561ffbdbf87b1c82d479c8791ab2c6ce741dd0d0b496d
15"

[[containers]]
source = "registry.redhat.io/amq7/amq-broker-
rhel8@sha256:ff6fefad518a6c997d4c5a6e475ba89640260167f0bc27715daf3cc30116fad1
"

…
EOF
```

> **IMPORTANT**
>
> For mirrored and disconnected use cases, ensure that all of the sources filtered from your catalog **index.json** file are digests. If any of the sources use tags instead of digests, the Operator installation fails. Tags require an internet connection.

2. View the **imageset-config.yaml** to get the catalog image reference for the **CatalogSource** custom resource (CR) by running the following command:

```
$ cat imageset-config.yaml
```

Example output

```
kind: ImageSetConfiguration
apiVersion: mirror.openshift.io/v1alpha2
storageConfig:
  registry:
    imageURL: registry.example.com/microshift-mirror
mirror:
  operators:
  - catalog: registry.redhat.io/redhat/redhat-operator-index:v4.15 ❶
    packages:
    - name: amq-broker-rhel8
      channels:
      - name: 7.11.x
```

❶ Use the value in the **mirror.catalog** catalog image reference for the follwing **jq** command to get the image digest. In this example, *<registry.redhat.io/redhat/redhat-operator-index:v4.15>.*

3. Get the SHA of the catalog index image by running the following command:

```
$ skopeo inspect docker://<registry.redhat.io/redhat/redhat-operator-index:v4.15> | jq
`.Digest` ❶
```

❶ Use the value in the **mirror.catalog** catalog image reference for the **jq** command to get the image digest. In this example, *<registry.redhat.io/redhat/redhat-operator-index:v4.15>.*

Example output

```
"sha256:7a76c0880a839035eb6e896d54ebd63668bb37b82040692141ba39ab4c539bc6"
```

4. To get ready to add the image references to your Image Builder blueprint file, format the catalog image reference by using the following example:

```
[[containers]]
source = "registry.redhat.io/redhat/redhat-operator-
index@sha256:7a76c0880a839035eb6e896d54ebd63668bb37b82040692141ba39ab4c539bc
6"
```

5. Add the image references from all the previous steps to the Image Builder blueprint.

Generated Image Builder blueprint example snippet

```
name = "microshift_blueprint"
description = "MicroShift 4.15.1 on x86_64 platform"
version = "0.0.1"
```

```
modules = []
groups = []

[[packages]]
name = "microshift"
version = "4.15.1"
...
...

[customizations.services]
enabled = ["microshift"]

[customizations.firewall]
ports = ["22:tcp", "80:tcp", "443:tcp", "5353:udp", "6443:tcp", "30000-32767:tcp", "30000-
32767:udp"]
...
...

[[containers]]
source = "quay.io/openshift-release-dev/ocp-v4.0-art-
dev@sha256:f41e79c17e8b41f1b0a5a32c3e2dd7cd15b8274554d3f1ba12b2598a347475f4"

[[containers]]
source = "quay.io/openshift-release-dev/ocp-v4.0-art-
dev@sha256:dbc65f1fba7d92b36cf7514cd130fe83a9bd211005ddb23a8dc479e0eea645fd"
...
...

[[containers]]
source = "registry.redhat.io/redhat/redhat-operator-
index@sha256:7a76c0880a839035eb6e896d54ebd63668bb37b82040692141ba39ab4c539bc
6"
...
...

[[containers]]
source = "registry.redhat.io/amq7/amq-broker-init-
rhel8@sha256:0dde839c2dce7cb684094bf26523c8e16677de03149a0fff468b8c3f106e1f4f"
...
...

[[containers]]
source = "registry.redhat.io/amq7/amq-broker-
rhel8@sha256:e8fa2a00e576ecb95561ffbdbf87b1c82d479c8791ab2c6ce741dd0d0b496d15"


[[containers]]
source = "registry.redhat.io/amq7/amq-broker-
rhel8@sha256:ff6fefad518a6c997d4c5a6e475ba89640260167f0bc27715daf3cc30116fad1"
...
EOF
```

1 References for all non-optional MicroShift RPM packages using the same version compatible with the **microshift-release-info** RPM.

2   References for automatically enabling MicroShift on system startup and applying default networking settings.

3   References for all non-optional MicroShift container images necessary for a disconnected deployment.

4   References for the catalog index.

8.4.1.3. Applying catalogs and Operators in a disconnected-deployment RHEL for Edge image

After you have created a RHEL for Edge image for a disconnected environment and configured MicroShift networking settings for disconnected use, you can configure the namespace and create catalog and Operator custom resources (CR) for running your Operators.

Prerequisites

- You have a RHEL for Edge image.

- Networking is configured for disconnected use.

- You completed the oc-mirror plugin dry run procedure.

Procedure

1. Create a **CatalogSource** custom resource (CR), similar to the following example:

   Example my-catalog-source-cr.yaml file

   ```
   apiVersion: operators.coreos.com/v1alpha1
   kind: CatalogSource
   metadata:
     name: cs-redhat-operator-index
     namespace: openshift-marketplace 1
   spec:
     image: registry.example.com/redhat/redhat-operator-index:v4.15
     sourceType: grpc
     displayName:
     publisher:
     updateStrategy:
       registryPoll:
         interval: 60m
   ```

1   The global namespace. Setting the **metadata.namespace** to **openshift-marketplace** enables the catalog to run in all namespaces. Subscriptions in any namespace can reference catalogs created in the **openshift-marketplace** namespace.

> **NOTE**
>
> The default pod security admission definition for **openshift-marketplace** is **baseline**, therefore a catalog source custom resource (CR) created in that namespace does not require a **spec.grpcPodConfig.securityContextConfig** value to be set. You can set a **legacy** or **restricted** value if required for the namespace and Operators you want to use.

2. Add the SHA of the catalog index commit to the Catalog Source (CR), similar to the following example:

Example namespace spec.image configuration

```
apiVersion: operators.coreos.com/v1alpha1
kind: CatalogSource
metadata:
  name: cs-redhat-operator-index
  namespace: openshift-marketplace
spec:
  image: registry.example.com/redhat/redhat-operator-index@sha256:7a76c0880a839035eb6e896d54ebd63668bb37b82040692141ba39ab4c539bc6 ❶
  sourceType: grpc
  displayName:
  publisher:
  updateStrategy:
    registryPoll:
      interval: 60m
```

❶ The SHA of the image commit. Use the same SHA you added to the Image Builder blueprint.

> **IMPORTANT**
>
> You must use the SHA instead of a tag in your catalog CR or the pod fails to start.

3. Apply the YAML file from the oc-mirror plugin dry run results directory to the cluster by running the following command:

```
$ oc apply -f ./oc-mirror-workspace/results-1708508014/catalogSource-cs-redhat-operator-index.yaml
```

Example output

```
catalogsource.operators.coreos.com/cs-redhat-operator-index created
```

4. Verify that the **CatalogSource** resources were successfully installed by running the following command:

```
$ oc get catalogsource --all-namespaces
```

5. Verify that the catalog source is running by using the following command:

```
$ oc get pods -n openshift-marketplace
```

Example output

```
NAME                          READY   STATUS    RESTARTS   AGE
cs-redhat-operator-index-4227b   2/2     Running   0          2m5s
```

6. Create a **Subscription** CR, similar to the following example:

Example my-subscription-cr.yaml file

```
apiVersion: operators.coreos.com/v1alpha1
kind: Subscription
metadata:
  name: amq-broker
  namespace: openshift-operators
spec:
  channel: 7.11.x
  name: amq-broker-rhel8
  source: cs-redhat-operator-index
  sourceNamespace: openshift-marketplace
```

7. Apply the **Subscription** CR by running the following command:

```
$ oc apply -f ./<my-subscription-cr.yaml>   ❶
```

❶ Specify the name of your **Subscription** CR, such as **my-subscription-cr.yaml**.

Example output

```
subscription.operators.coreos.com/amq-broker created
```