



Red Hat Enterprise Linux 8

Composing, installing, and managing RHEL for Edge images

Creating, deploying, and managing Edge systems with Red Hat Enterprise Linux 8

Red Hat Enterprise Linux 8 Composing, installing, and managing RHEL for Edge images

Creating, deploying, and managing Edge systems with Red Hat Enterprise Linux 8

Legal Notice

Copyright © 2024 Red Hat, Inc.

The text of and illustrations in this document are licensed by Red Hat under a Creative Commons Attribution–Share Alike 3.0 Unported license ("CC-BY-SA"). An explanation of CC-BY-SA is available at

<http://creativecommons.org/licenses/by-sa/3.0/>

. In accordance with CC-BY-SA, if you distribute this document or an adaptation of it, you must provide the URL for the original version.

Red Hat, as the licensor of this document, waives the right to enforce, and agrees not to assert, Section 4d of CC-BY-SA to the fullest extent permitted by applicable law.

Red Hat, Red Hat Enterprise Linux, the Shadowman logo, the Red Hat logo, JBoss, OpenShift, Fedora, the Infinity logo, and RHCE are trademarks of Red Hat, Inc., registered in the United States and other countries.

Linux[®] is the registered trademark of Linus Torvalds in the United States and other countries.

Java[®] is a registered trademark of Oracle and/or its affiliates.

XFS[®] is a trademark of Silicon Graphics International Corp. or its subsidiaries in the United States and/or other countries.

MySQL[®] is a registered trademark of MySQL AB in the United States, the European Union and other countries.

Node.js[®] is an official trademark of Joyent. Red Hat is not formally related to or endorsed by the official Joyent Node.js open source or commercial project.

The OpenStack[®] Word Mark and OpenStack logo are either registered trademarks/service marks or trademarks/service marks of the OpenStack Foundation, in the United States and other countries and are used with the OpenStack Foundation's permission. We are not affiliated with, endorsed or sponsored by the OpenStack Foundation, or the OpenStack community.

All other trademarks are the property of their respective owners.

Abstract

Use the RHEL image builder tool to compose customized RHEL (rpm-ostree) images optimized for Edge. Then, remotely install, and securely manage and scale deployments of the images on Edge servers.

Table of Contents

MAKING OPEN SOURCE MORE INCLUSIVE	5
PROVIDING FEEDBACK ON RED HAT DOCUMENTATION	6
CHAPTER 1. INTRODUCING RHEL FOR EDGE IMAGES	7
1.1. RHEL FOR EDGE–SUPPORTED ARCHITECTURE	8
1.2. HOW TO COMPOSE AND DEPLOY A RHEL FOR EDGE IMAGE	8
1.3. NON-NETWORK-BASED DEPLOYMENTS	10
1.4. NETWORK-BASED DEPLOYMENTS	11
1.5. DIFFERENCE BETWEEN RHEL RPM IMAGES AND RHEL FOR EDGE IMAGES	12
CHAPTER 2. SETTING UP RHEL IMAGE BUILDER	14
2.1. IMAGE BUILDER SYSTEM REQUIREMENTS	14
2.2. INSTALLING RHEL IMAGE BUILDER	14
CHAPTER 3. CONFIGURING RHEL IMAGE BUILDER REPOSITORIES	16
3.1. ADDING CUSTOM THIRD-PARTY REPOSITORIES TO RHEL IMAGE BUILDER	16
3.2. ADDING THIRD-PARTY REPOSITORIES WITH SPECIFIC DISTRIBUTIONS TO RHEL IMAGE BUILDER	17
3.3. CHECKING REPOSITORIES METADATA WITH GPG	17
3.4. RHEL IMAGE BUILDER OFFICIAL REPOSITORY OVERRIDES	19
3.5. OVERRIDING A SYSTEM REPOSITORY	19
3.6. OVERRIDING A SYSTEM REPOSITORY THAT REQUIRES SUBSCRIPTIONS	21
CHAPTER 4. COMPOSING A RHEL FOR EDGE IMAGE USING IMAGE BUILDER IN RHEL WEB CONSOLE	23
4.1. ACCESSING RHEL IMAGE BUILDER IN THE RHEL WEB CONSOLE	23
4.2. CREATING A BLUEPRINT FOR A RHEL FOR EDGE IMAGE USING IMAGE BUILDER IN THE WEB CONSOLE	24
4.3. CREATING A RHEL FOR EDGE COMMIT IMAGE BY USING IMAGE BUILDER IN WEB CONSOLE	26
4.4. CREATING A RHEL FOR EDGE CONTAINER IMAGE BY USING RHEL IMAGE BUILDER IN RHEL WEB CONSOLE	27
4.5. CREATING A RHEL FOR EDGE INSTALLER IMAGE BY USING IMAGE BUILDER IN RHEL WEB CONSOLE	28
4.6. DOWNLOADING A RHEL FOR EDGE IMAGE	29
4.7. ADDITIONAL RESOURCES	30
CHAPTER 5. COMPOSING A RHEL FOR EDGE IMAGE USING IMAGE BUILDER COMMAND-LINE	31
5.1. NETWORK-BASED DEPLOYMENTS WORKFLOW	31
5.1.1. Creating a RHEL for Edge Commit image blueprint using image builder command-line interface	31
5.1.2. Creating a RHEL for Edge Commit image using image builder command-line interface	33
5.1.3. Creating a RHEL for Edge image update with a ref commit by using RHEL image builder CLI	34
5.1.4. Downloading a RHEL for Edge image using the image builder command-line interface	36
5.2. NON-NETWORK-BASED DEPLOYMENTS WORKFLOW	36
5.2.1. Creating a RHEL for Edge Container blueprint by using image builder CLI	36
5.2.2. Creating a RHEL for Edge Installer blueprint using image builder CLI	38
5.2.3. Creating a RHEL for Edge Container image by using image builder CLI	39
5.2.4. Creating a RHEL for Edge Installer image using command-line interface for non-network-based deployments	40
5.2.5. Downloading a RHEL for Edge Installer image using the image builder CLI	41
5.3. SUPPORTED IMAGE CUSTOMIZATIONS	42
5.3.1. Selecting a distribution	42
5.3.2. Selecting a package group	43
5.3.3. Setting the image hostname	43
5.3.4. Specifying additional users	43

5.3.5. Specifying additional groups	44
5.3.6. Setting SSH key for existing users	44
5.3.7. Appending a kernel argument	45
5.3.8. Setting time zone and NTP	45
5.3.9. Customizing the locale settings	46
5.3.10. Customizing firewall	46
5.3.11. Enabling or disabling services	47
5.3.12. Specifying a custom filesystem configuration	47
5.4. PACKAGES INSTALLED BY RHEL IMAGE BUILDER	51
5.5. ADDITIONAL RESOURCES	55
CHAPTER 6. BUILDING SIMPLIFIED INSTALLER IMAGES TO PROVISION A RHEL FOR EDGE IMAGE ...	56
6.1. SIMPLIFIED INSTALLER IMAGE BUILD AND DEPLOYMENT	56
6.2. CREATING A BLUEPRINT FOR A SIMPLIFIED IMAGE USING RHEL IMAGE BUILDER CLI	57
6.3. CREATING A RHEL FOR EDGE SIMPLIFIED INSTALLER IMAGE USING IMAGE BUILDER CLI	58
6.4. DOWNLOADING A SIMPLIFIED RHEL FOR EDGE IMAGE USING THE IMAGE BUILDER COMMAND-LINE INTERFACE	60
6.5. CREATING A BLUEPRINT FOR A SIMPLIFIED IMAGE RHEL USING IMAGE BUILDER GUI	60
6.6. CREATING A RHEL FOR EDGE SIMPLIFIED INSTALLER IMAGE USING IMAGE BUILDER GUI	63
6.7. DOWNLOADING A SIMPLIFIED RHEL FOR EDGE IMAGE USING THE IMAGE BUILDER GUI	64
6.8. SETTING UP AN UEFI HTTP BOOT SERVER	64
6.9. DEPLOYING THE SIMPLIFIED ISO IMAGE IN A VIRTUAL MACHINE	65
6.10. DEPLOYING THE SIMPLIFIED ISO IMAGE FROM A USB FLASH DRIVE	66
CHAPTER 7. AUTOMATICALLY PROVISIONING AND ONBOARDING RHEL FOR EDGE DEVICES WITH FDO .	68
7.1. THE FIDO DEVICE ONBOARDING (FDO) PROCESS	68
7.2. AUTOMATICALLY PROVISIONING AND ONBOARDING RHEL FOR EDGE DEVICES	71
7.3. GENERATING KEY AND CERTIFICATES	72
7.4. INSTALLING AND RUNNING THE MANUFACTURING SERVER	73
7.5. INSTALLING, CONFIGURING, AND RUNNING THE RENDEZVOUS SERVER	76
7.6. INSTALLING, CONFIGURING, AND RUNNING THE OWNER SERVER	77
7.7. AUTOMATICALLY ONBOARDING A RHEL FOR EDGE DEVICE BY USING FDO AUTHENTICATION	79
CHAPTER 8. DEPLOYING A RHEL FOR EDGE IMAGE IN A NETWORK-BASED ENVIRONMENT	81
8.1. EXTRACTING THE RHEL FOR EDGE IMAGE COMMIT	81
8.2. SETTING UP A WEB SERVER TO INSTALL RHEL FOR EDGE IMAGES	83
8.3. PERFORMING AN ATTENDED INSTALLATION TO AN EDGE DEVICE BY USING KICKSTART	85
8.4. PERFORMING AN UNATTENDED INSTALLATION TO AN EDGE DEVICE BY USING KICKSTART	87
CHAPTER 9. DEPLOYING A RHEL FOR EDGE IMAGE IN A NON-NETWORK-BASED ENVIRONMENT ...	90
9.1. CREATING A RHEL FOR EDGE CONTAINER IMAGE FOR NON-NETWORK-BASED DEPLOYMENTS	90
9.2. CREATING A RHEL FOR EDGE INSTALLER IMAGE FOR NON-NETWORK-BASED DEPLOYMENTS	91
9.3. INSTALLING THE RHEL FOR EDGE IMAGE FOR NON-NETWORK-BASED DEPLOYMENTS	92
CHAPTER 10. MANAGING RHEL FOR EDGE IMAGES	95
10.1. EDITING A RHEL FOR EDGE IMAGE BLUEPRINT BY USING IMAGE BUILDER	95
10.1.1. Adding a component to RHEL for Edge blueprint using image builder in RHEL web console	95
10.1.2. Removing a component from a blueprint using RHEL image builder in the web console	96
10.1.3. Editing a RHEL for Edge image blueprint using command-line interface	96
10.2. UPDATING RHEL FOR EDGE IMAGES	97
10.2.1. How RHEL for Edge image updates are deployed	97
10.2.2. Building a commit update	99
10.2.3. Deploying RHEL for Edge image updates manually	100

10.2.4. Deploying RHEL for Edge image updates manually using the command-line	103
10.2.5. Deploying RHEL for Edge image updates manually for non-network-base deployments	105
10.3. DEPLOYING RHEL FOR EDGE AUTOMATIC IMAGE UPDATES	107
10.3.1. Updating the RHEL for Edge image update policy	108
10.3.2. Enabling RHEL for Edge automatic download and staging of updates	108
10.4. ROLLING BACK RHEL FOR EDGE IMAGES	109
10.4.1. Introducing the greenboot checks	110
10.4.2. RHEL for Edge images roll back with greenboot	110
10.4.3. Greenboot health check status	112
10.4.4. Checking greenboot health checks statuses	113
10.4.5. Manually rolling back RHEL for Edge images	113
10.4.6. Rolling back RHEL for Edge images using an automated process	114
CHAPTER 11. CREATING AND MANAGING OSTREE IMAGE UPDATES	116
11.1. BASIC CONCEPTS FOR OSTREE	116
11.2. CREATING OSTREE REPOSITORIES	117
11.3. MANAGING A CENTRALIZED OSTREE MIRROR	117
APPENDIX A. TERMINOLOGY AND COMMANDS	122
A.1. OSTREE AND RPM-OSTREE TERMINOLOGY	122
A.2. OSTREE COMMANDS	122
A.3. RPM-OSTREE COMMANDS	123
A.4. FDO AUTOMATIC ONBOARDING TERMINOLOGY	124
A.5. FDO AUTOMATIC ONBOARDING TECHNOLOGIES	125

MAKING OPEN SOURCE MORE INCLUSIVE

Red Hat is committed to replacing problematic language in our code, documentation, and web properties. We are beginning with these four terms: master, slave, blacklist, and whitelist. Because of the enormity of this endeavor, these changes will be implemented gradually over several upcoming releases. For more details, see [our CTO Chris Wright's message](#).

PROVIDING FEEDBACK ON RED HAT DOCUMENTATION

We appreciate your feedback on our documentation. Let us know how we can improve it.

Submitting feedback through Jira (account required)

1. Log in to the [Jira](#) website.
2. Click **Create** in the top navigation bar.
3. Enter a descriptive title in the **Summary** field.
4. Enter your suggestion for improvement in the **Description** field. Include links to the relevant parts of the documentation.
5. Click **Create** at the bottom of the dialogue.

CHAPTER 1. INTRODUCING RHEL FOR EDGE IMAGES

A RHEL for Edge image is an **rpm-ostree** image that includes system packages to remotely install RHEL on Edge servers.

The system packages include:

- **Base OS** package
- Podman as the container engine
- Additional RPM content

Differently from RHEL images, RHEL for Edge is an immutable operating system, that is, it contains a **read-only** root directory with the following characteristics:

- The packages are isolated from root directory
- Package installs create layers that make it easy to rollback to previous versions
- Efficient updates to disconnected environments
- Supports multiple operating system branches and repositories
- Has a hybrid **rpm-ostree** package system

You can deploy a RHEL for Edge image on Bare Metal, Appliance, and Edge servers.

You can compose customized RHEL for Edge images using the RHEL image builder tool. You can also create RHEL for Edge images by accessing the [edge management](#) application in the Red Hat Hybrid Cloud Console platform and configure automated management.

The edge management application simplifies the way you can provision and register your images. To learn more about the edge management, see the [Create RHEL for Edge images and configure automated management](#) documentation.



WARNING

Using RHEL for Edge customized images that were created using the **RHEL image builder** on-premise version artifacts is not supported in the [edge management](#) application. See [Edge management supportability](#).

With a RHEL for Edge image, you can achieve the following:

Atomic upgrades

State of each update is known / no changes are seen until reboot

Custom health checks and intelligent rollbacks

Resiliency in case of failed upgrades

Container-focused workflow

Separate core OS updates

Optimized OTA payloads

Transfer only delta updates

1.1. RHEL FOR EDGE–SUPPORTED ARCHITECTURE

Currently, you can deploy RHEL for Edge images on AMD and Intel 64-bit systems.



NOTE

RHEL for Edge does not support ARM systems in RHEL 8.

1.2. HOW TO COMPOSE AND DEPLOY A RHEL FOR EDGE IMAGE

Composing and deploying a RHEL for Edge image involves two phases:

1. Composing a RHEL **rpm-ostree** image using the RHEL image builder tool. You can access RHEL image builder through a command-line interface in the **composer-cli** tool, or use a graphical user interface in the RHEL web console.
2. Deploying the image using RHEL installer.

While composing a RHEL for Edge image, you can select any of the following image types. Composing the different RHEL for Edge images might or might not require network access. See the table:

Table 1.1. RHEL for Edge images type

Image type	Description	Suitable for network-based deployments	Suitable for non-network-based deployments
RHEL for Edge Commit (.tar)	The edge-commit image is not directly bootable, even though it contains a full operating system. To boot the edge-commit image type, you must deploy it.	Yes	No
RHEL for Edge Container (.tar)	The edge-container creates an OSTree commit and embeds it into an OCI container with a web server. When the edge-commit image starts, the web server serves the commit as an OSTree repository.	No	Yes

Image type	Description	Suitable for network-based deployments	Suitable for non-network-based deployments
RHEL for Edge Installer (.iso)	The edge-installer image type pulls the commit from the running container and creates an installable boot ISO with a Kickstart file configured to use the embedded OSTree commit.	No	Yes
RHEL for Edge Raw Image (.raw.xz)	The edge-raw-image compressed raw images consist of a file that contains a partition layout with an existing deployed OSTree commit in it. You can flash the RHEL Raw Images on a hard disk or boot on a virtual machine.	Yes	Yes
RHEL for Edge Simplified Installer (.iso)	The edge-simplified-installer image uses the Ignition tool to inject the user configuration into the images at an early stage of the boot process.	Yes	Yes
RHEL for Edge AMI (.ami)	The edge-ami image uses the Ignition tool to inject the user configuration into the images at an early stage of the boot process. You can upload the .ami image to AWS and boot an EC2 instance in AWS.	Yes	Yes

Image type	Description	Suitable for network-based deployments	Suitable for non-network-based deployments
RHEL for Edge VMDK (.vmdk)	The edge-vshpere image uses the Ignition tool to inject the user configuration into the images at an early stage of the boot process. You can load the image on vSphere and boot the image in a VM vSphere	Yes	Yes

The image types vary in terms of their contents, and are therefore suitable for different types of deployment environments.

Additional resources

- [Performing a standard RHEL 8 installation](#) .

1.3. NON-NETWORK-BASED DEPLOYMENTS

Use RHEL image builder to create flexible RHEL **rpm-ostree** images to suit your requirements, and then use Anaconda to deploy them in your environment.

You can access RHEL image builder through a command-line interface in the **composer-cli** tool, or use a graphical user interface in the RHEL web console.

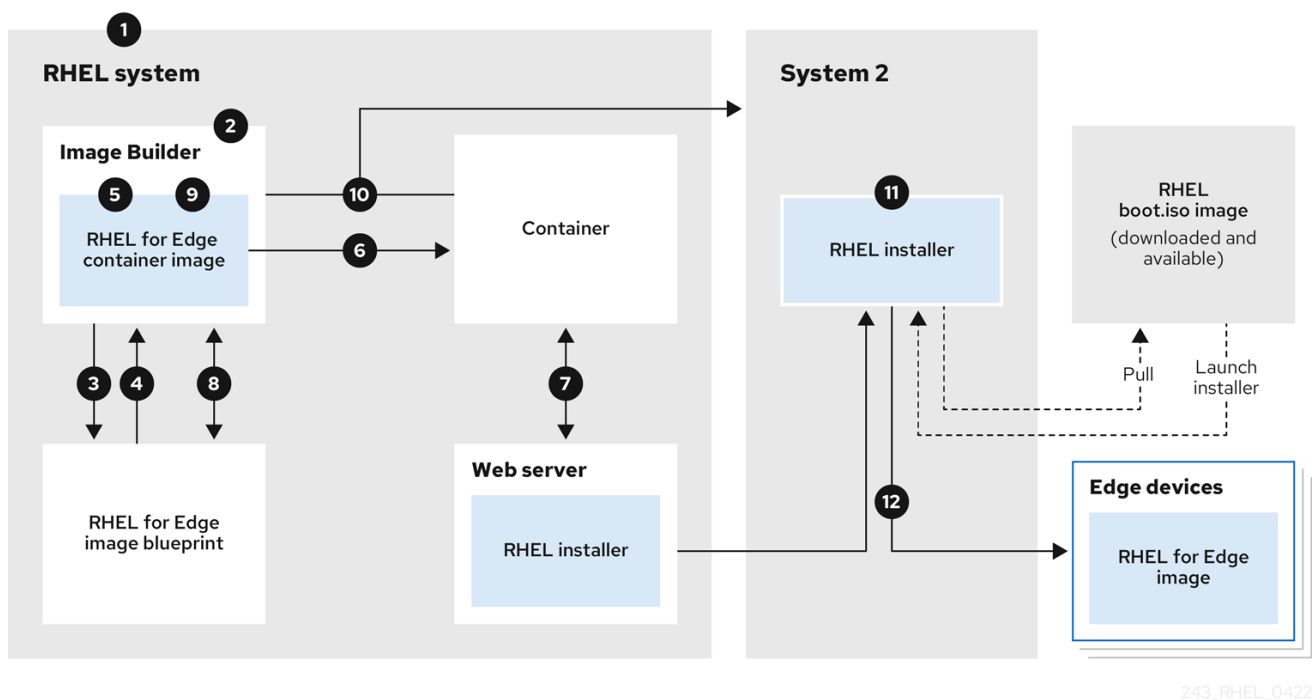
Composing and deploying a RHEL for Edge image in non-network-based deployments involves the following high-level steps:

1. Install and register a RHEL system
2. Install RHEL image builder
3. Using RHEL image builder, create a blueprint with customizations for RHEL for Edge Container image
4. Import the RHEL for Edge blueprint in RHEL image builder
5. Create a RHEL for Edge image embed in an OCI container with a webserver ready to deploy the commit as an OSTree repository
6. Download the RHEL for Edge Container image file
7. Deploy the container serving a repository with the RHEL for Edge Container commit
8. Using RHEL image builder, create another blueprint for RHEL for Edge Installer image
9. Create a RHEL for Edge Installer image configured to pull the commit from the running container embedded with RHEL for Edge Container image
10. Download the RHEL for Edge Installer image

11. Run the installation

The following diagram represents the RHEL for Edge image non-network deployment workflow:

Figure 1.1. Deploying RHEL for Edge in non-network environment



1.4. NETWORK-BASED DEPLOYMENTS

Use RHEL image builder to create flexible RHEL **rpm-ostree** images to suit your requirements, and then use Anaconda to deploy them in your environment. RHEL image builder automatically identifies the details of your deployment setup and generates the image output as an **edge-commit** as a **.tar** file.

You can access RHEL image builder through a command-line interface in the **composer-cli** tool, or use a graphical user interface in the RHEL web console.

You can compose and deploy the RHEL for Edge image by performing the following high-level steps:

For an attended installation

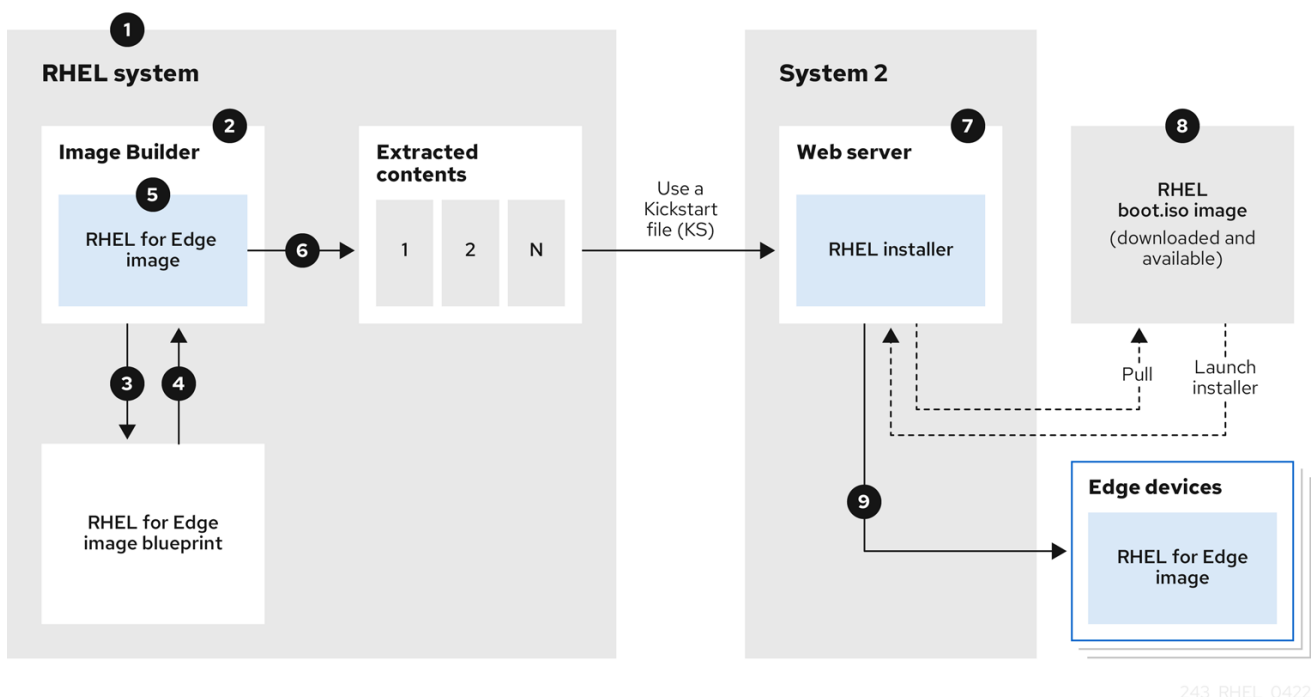
1. Install and register a RHEL system
2. Install RHEL image builder
3. Using RHEL image builder, create a blueprint for RHEL for Edge image
4. Import the RHEL for Edge blueprint in RHEL image builder
5. Create a RHEL for Edge Commit (**.tar**) image
6. Download the RHEL for Edge image file
7. On the same system where you have installed RHEL image builder, install a web server that you want to serve the the RHEL for Edge Commit content. For instructions, see [Setting up and configuring NGINX](#)

8. Extract the RHEL for Edge Commit (**.tar**) content to the running web server
9. Create a Kickstart file that pulls the OSTree content from the running web server. For details on how to modify the Kickstart to pull the OSTree content, see [Extracting the RHEL for Edge image commit](#)
10. Boot the RHEL installer ISO on the edge device and provide the Kickstart to it.

For an unattended installation, you can customize the RHEL installation ISO and embed the Kickstart file to it.

The following diagram represents the RHEL for Edge network image deployment workflow:

Figure 1.2. Deploying RHEL for Edge in network-base environment



1.5. DIFFERENCE BETWEEN RHEL RPM IMAGES AND RHEL FOR EDGE IMAGES

You can create RHEL system images in traditional package-based RPM format and also as RHEL for Edge (**rpm-ostree**) images.

You can use the traditional package-based RPMs to deploy RHEL on traditional data centers. However, with RHEL for Edge images you can deploy RHEL on servers other than traditional data centers. These servers include systems where processing of large amounts of data is done closest to the source where data is generated—Edge servers.

The RHEL for Edge (**rpm-ostree**) images are not a package manager. They only support complete bootable file system trees, not individual files. These images do not have information regarding the individual files such as how these files were generated or anything related to their origin.

The **rpm-ostree** images need a separate mechanism, the package manager, to install additional applications in the **/var** directory. With that, the **rpm-ostree** image keeps the operating system unchanged, while maintaining the state of the **/var** and **/etc** directories. The atomic updates enable rollbacks and background staging of updates.

Refer to the following table to know how RHEL for Edge images differ from the package-based RHEL RPM images.

Table 1.2. Difference between RHEL RPM images and RHEL for Edge images

Key attributes	RHEL RPM image	RHEL for Edge image
OS assembly	You can assemble the packages locally to form an image.	The packages are assembled in an ostree which you can install on a system.
OS updates	You can use yum update to apply the available updates from the enabled repositories.	You can use rpm-ostree upgrade to stage an update if any new commit is available in the ostree remote at /etc/ostree/remotes.d/ . The update takes effect on system reboot.
Repository	The package contains YUM repositories	The package contains Ostree remote repository
User access permissions	Read write	Read-only (/usr)
Data persistence	You can mount the image to any non tmpfs mount point	/etc & /var are read-write enabled and include persisting data.

CHAPTER 2. SETTING UP RHEL IMAGE BUILDER

Use RHEL image builder to create your customized RHEL for Edge images. After you install RHEL image builder on a RHEL system, RHEL image builder is available as an application in RHEL web console. You can also access RHEL image builder through a command line interface in the **composer-cli** tool.



NOTE

It is recommended to install RHEL image builder on a virtual machine.

2.1. IMAGE BUILDER SYSTEM REQUIREMENTS

The environment where RHEL image builder runs, for example a virtual machine, must meet the requirements that are listed in the following table.



NOTE

Running RHEL image builder inside a container is not supported.

Table 2.1. Image builder system requirements

Parameter	Minimal Required Value
System type	A dedicated virtual machine
Processor	2 cores
Memory	4 GiB
Disk space	20 GiB
Access privileges	Administrator level (root)
Network	Connectivity to Internet



NOTE

The 20 GiB disk space requirement is enough to install and run RHEL image builder in the host. To build and deploy image builds, you must allocate additional dedicated disk space.

2.2. INSTALLING RHEL IMAGE BUILDER

To install RHEL image builder on a dedicated virtual machine, follow these steps:

Prerequisites

- The virtual machine is created and is powered on.
- You have installed RHEL and you have subscribed to RHSM or Red Hat Satellite.

- You have enabled the **BaseOS** and **AppStream** repositories to be able to install the RHEL image builder packages.

Procedure

1. Install the following packages on the virtual machine.

- `osbuild-composer`
- `composer-cli`
- `cockpit-composer`
- `bash-completion`
- `firewalld`

```
# yum install osbuild-composer composer-cli cockpit-composer bash-completion firewalld
```

RHEL image builder is installed as an application in RHEL web console.

2. Reboot the virtual machine
3. Configure the system firewall to allow access to the web console:

```
# firewall-cmd --add-service=cockpit && firewall-cmd --add-service=cockpit --permanent
```

4. Enable RHEL image builder.

```
# systemctl enable osbuild-composer.socket cockpit.socket --now
```

The `osbuild-composer` and `cockpit` services start automatically on first access.

5. Load the shell configuration script so that the autocomplete feature for the **composer-cli** command starts working immediately without reboot:

```
$ source /etc/bash_completion.d/composer-cli
```

Additional resources

- [Managing repositories](#)

CHAPTER 3. CONFIGURING RHEL IMAGE BUILDER REPOSITORIES

To use RHEL image builder, you must ensure that the repositories are configured. You can use the following types of repositories in RHEL image builder:

Official repository overrides

Use these if you want to download base system RPMs from elsewhere than the Red Hat Content Delivery Network (CDN) official repositories, for example, a custom mirror in your network. Using official repository overrides disables the default repositories, and your custom mirror must contain all the necessary packages.

Custom third-party repositories

Use these to include packages that are not available in the official RHEL repositories.

3.1. ADDING CUSTOM THIRD-PARTY REPOSITORIES TO RHEL IMAGE BUILDER

You can add custom third-party sources to your repositories and manage these repositories by using the **composer-cli**.

Prerequisites

- You have the URL of the custom third-party repository.

Procedure

1. Create a repository source file, such as **/root/repo.toml**. For example:

```
id = "k8s"
name = "Kubernetes"
type = "yum-baseurl"
url = "https://server.example.com/repos/company_internal_packages/"
check_gpg = false
check_ssl = false
system = false
```

The **type** field accepts the following valid values: **yum-baseurl**, **yum-mirrorlist**, and **yum-metalink**.

2. Save the file in the TOML format.
3. Add the new third-party source to RHEL image builder:

```
$ composer-cli sources add <file-name>.toml
```

Verification

1. Check if the new source was successfully added:

```
$ composer-cli sources list
```

2. Check the new source content:

```
$ composer-cli sources info <source_id>
```

3.2. ADDING THIRD-PARTY REPOSITORIES WITH SPECIFIC DISTRIBUTIONS TO RHEL IMAGE BUILDER

You can specify a list of distributions in the custom third-party source file by using the optional field **distro**. The repository file uses the distribution string list while resolving dependencies during the image building.

Any request that specifies **rhel-8** uses this source. For example, if you list packages and specify **rhel-8**, it includes this source. However, listing packages for the host distribution do not include this source.

Prerequisites

- You have the URL of the custom third-party repository.
- You have the list of distributions that you want to specify.

Procedure

1. Create a repository source file, such as **/root/repo.toml**. For example, to specify the distribution:

```
check_gpg = true
check_ssl = true
distros = ["rhel-8"]
id = "rh9-local"
name = "packages for RHEL"
system = false
type = "yum-baseurl"
url = "https://local/repos/rhel8/projectrepo/"
```

2. Save the file in the TOML format.
3. Add the new third-party source to RHEL image builder:

```
$ composer-cli sources add <file-name>.toml
```

Verification

1. Check if the new source was successfully added:

```
$ composer-cli sources list
```

2. Check the new source content:

```
$ composer-cli sources info <source_id>
```

3.3. CHECKING REPOSITORIES METADATA WITH GPG

To detect and avoid corrupted packages, you can use the DNF package manager to check the GNU Privacy Guard (GPG) signature on RPM packages, and also to check if the repository metadata has been signed with a GPG key.

You can either enter the **gpgkey** that you want to do the check over **https** by setting the **gpgkeys** field with the key URL. Alternatively, to improve security, you can also embed the whole key into the **gpgkeys** field, to import it directly instead of fetching the key from the URL.

Prerequisites

- The directory that you want to use as a repository exists and contains packages.

Procedure

1. Access the folder where you want to create a repository:

```
$ cd repo/
```

2. Run the **createrepo_c** to create a repository from RPM packages:

```
$ createrepo_c .
```

3. Access the directory where the repodata is:

```
$ cd repodata/
```

4. Sign your **repomd.xml** file:

```
$ gpg -u <_gpg-key-email_> --yes --detach-sign --armor /srv/repo/example/repomd.xml
```

5. To enable GPG signature checks in the repository:

- a. Set **check_repogpg = true** in the repository source.
- b. Enter the **gpgkey** that you want to do the check. If your key is available over **https**, set the **gpgkeys** field with the key URL for the key. You can add as many URL keys as you need. The following is an example:

```
check_gpg = true
check_ssl = true
id = "signed local packages"
name = "repository_name"
type = "yum-baseurl"
url = "https://local/repos/projectrepo/"
check_repogpg = true
gpgkeys=["https://local/keys/repokey.pub"]
```

As an alternative, add the GPG key directly in the **gpgkeys** field, for example:

```
check_gpg = true
check_ssl = true
check_repogpg
id = "custom-local"
name = "signed local packages"
type = "yum-baseurl"
url = "https://local/repos/projectrepo/"
gpgkeys=["https://remote/keys/other-repokey.pub",
```

```
"-----BEGIN PGP PUBLIC KEY BLOCK-----
...
-----END PGP PUBLIC KEY BLOCK-----"]
```

- If the test does not find the signature, the GPG tool shows an error similar to the following one:

```
$ GPG verification is enabled, but GPG signature is not available.
This may be an error or the repository does not support GPG verification:
Status code: 404 for http://repo-server/rhel/repodata/repomd.xml.asc (IP:
192.168.1.3)
```

- If the signature is invalid, the GPG tool shows an error similar to the following one:

```
repomd.xml GPG signature verification error: Bad GPG signature
```

Verification

- Test the signature of the repository manually:

```
$ gpg --verify /srv/repo/example/repomd.xml.asc
```

3.4. RHEL IMAGE BUILDER OFFICIAL REPOSITORY OVERRIDES

RHEL image builder **osbuild-composer** back end does not inherit the system repositories located in the **/etc/yum.repos.d/** directory. Instead, it has its own set of official repositories defined in the **/usr/share/osbuild-composer/repositories** directory. This includes the Red Hat official repository, which contains the base system RPMs to install additional software or update already installed programs to newer versions. If you want to override the official repositories, you must define overrides in **/etc/osbuild-composer/repositories/**. This directory is for user defined overrides and the files located there take precedence over those in the **/usr/share/osbuild-composer/repositories/** directory.

The configuration files are not in the usual YUM repository format known from the files in **/etc/yum.repos.d/**. Instead, they are JSON files.

3.5. OVERRIDING A SYSTEM REPOSITORY

You can configure your own repository override for RHEL image builder in the **/etc/osbuild-composer/repositories** directory.



NOTE

Prior to RHEL 8.5 release, the name of the repository overrides is **rhel-8.json**. Starting from RHEL 8.5, the names also respect the minor version: **rhel-84.json**, **rhel-85.json**, and so on.

Prerequisites

- You have a custom repository that is accessible from your host system.

Procedure

1. Create the **/etc/osbuild-composer/repositories/** directory to store your repository overrides:

```
$ sudo mkdir -p /etc/osbuild-composer/repositories
```

2. Create a JSON file, using a name corresponding to your RHEL version. Alternatively, you can copy the file for your distribution from `/usr/share/osbuild-composer/` and modify its content. For RHEL 8.9, use `/etc/osbuild-composer/repositories/rhel-89.json`.
3. Add the following structure to your JSON file. Specify only one of the following attributes, in the string format:

- **baseurl** - The base URL of the repository.
- **metalink** - The URL of a metalink file that contains a list of valid mirror repositories.
- **mirrorlist** - The URL of a mirrorlist file that contains a list of valid mirror repositories. The remaining fields, such as **gpgkey**, and **metadata_expire**, are optional. For example:

```
{
  "x86_64": [
    {
      "name": "baseos",
      "baseurl": "http://mirror.example.com/composes/released/RHEL-
8/8.0/BaseOS/x86_64/os/",
      "gpgkey": "-----BEGIN PGP PUBLIC KEY BLOCK-----\n\n (...)",
      "check_gpg": true
    }
  ]
}
```

Alternatively, you can copy the JSON file for your distribution, by replacing **rhel-version.json** with your RHEL version, for example: `rhel-8.json`.

```
$ cp /usr/share/osbuild-composer/repositories/rhel-version.json /etc/osbuild-
composer/repositories/
```

4. Optional: Verify the JSON file:

```
$ json_verify /etc/osbuild-composer/repositories/<file>.json
```

5. Edit the **baseurl** paths in the **rhel-8.json** file and save it. For example:

```
$ /etc/osbuild-composer/repositories/rhel-version.json
```

6. Restart the **osbuild-composer.service**:

```
$ sudo systemctl restart osbuild-composer.service
```

Verification

- Check if the repository points to the correct URLs:

```
$ cat /etc/yum.repos.d/redhat.repo
```


You can see that the repository points to the correct URLs which are copied from the `/etc/yum.repos.d/redhat.repo` file.

Additional resources

- [The latest RPMs version available in repository not visible for `osbuild-composer`](#).

3.6. OVERRIDING A SYSTEM REPOSITORY THAT REQUIRES SUBSCRIPTIONS

You can set up the `osbuild-composer` service to use system subscriptions that are defined in the `/etc/yum.repos.d/redhat.repo` file. To use a system subscription in `osbuild-composer`, define a repository override that has the following details:

- The same **baseurl** as the repository defined in `/etc/yum.repos.d/redhat.repo`.
- The value of **"rhsm": true** defined in the JSON object.



NOTE

`osbuild-composer` does not automatically use repositories defined in `/etc/yum.repos.d/`. You need to manually specify them either as a system repository override or as an additional **source** by using `composer-cli`. The "BaseOS" and "AppStream" repositories usually use system repository overrides, whereas all the other repositories use `composer-cli` sources.

Prerequisites

- Your system has a subscription defined in `/etc/yum.repos.d/redhat.repo`
- You have created a repository override. See [Overriding a system repository](#).

Procedure

1. Get the **baseurl** from the `/etc/yum.repos.d/redhat.repo` file:

```
# cat /etc/yum.repos.d/redhat.repo
[AppStream]
name = AppStream mirror example
baseurl = https://mirror.example.com/RHEL-8/8.0/AppStream/x86_64/os/
enabled = 1
gpgcheck = 0
sslvcrify = 1
sslcacert = /etc/pki/ca1/ca.crt
sslclientkey = /etc/pki/ca1/client.key
sslclientcert = /etc/pki/ca1/client.crt
metadata_expire = 86400
enabled_metadata = 0
```

2. Configure the repository override to use the same **baseurl** and set **rhsm** to true:

```
{
  "x86_64": [
    {
```

```
{
  "name": "AppStream mirror example",
  "baseurl": "https://mirror.example.com/RHEL-8/8.0/AppStream/x86_64/os/",
  "gpgkey": "-----BEGIN PGP PUBLIC KEY BLOCK-----\n\n (...)",
  "check_gpg": true,
  "rhsm": true
}
```

3. Restart the **osbuild-composer.service**:

```
$ sudo systemctl restart osbuild-composer.service
```

Additional resources

- [RHEL image builder uses CDN repositories when host is registered to Satellite 6](#)

CHAPTER 4. COMPOSING A RHEL FOR EDGE IMAGE USING IMAGE BUILDER IN RHEL WEB CONSOLE

Use RHEL image builder to create a custom RHEL for Edge image (OSTree commit).

To access RHEL image builder and to create your custom RHEL for Edge image, you can either use the RHEL web console interface or the command-line interface.

You can compose RHEL for Edge images using RHEL image builder in RHEL web console by performing the following high-level steps:

1. Access RHEL image builder in RHEL web console
2. Create a blueprint for RHEL for Edge image.
3. Create a RHEL for Edge image. You can create the following images:
 - RHEL for Edge Commit image.
 - RHEL for Edge Container image.
 - RHEL for Edge Installer image.
4. Download the RHEL for Edge image

4.1. ACCESSING RHEL IMAGE BUILDER IN THE RHEL WEB CONSOLE

To access RHEL image builder in RHEL web console, ensure that you have met the following prerequisites and then follow the procedure.

Prerequisites

- You have installed a RHEL system.
- You have administrative rights on the system.
- You have subscribed the RHEL system to Red Hat Subscription Manager (RHSM) or to Red Hat Satellite Server.
- The system is powered on and accessible over network.
- You have installed RHEL image builder on the system.

Procedure

1. On your RHEL system, access `https://localhost:9090/` in a web browser.
2. For more information about how to remotely access RHEL image builder, see [Managing systems using the RHEL 8 web console](#) document.
3. Log in to the web console using an administrative user account.
4. On the web console, in the left hand menu, click **Apps**.
5. Click **Image Builder**.

The RHEL image builder dashboard opens in the right pane. You can now proceed to create a blueprint for the RHEL for Edge images.

4.2. CREATING A BLUEPRINT FOR A RHEL FOR EDGE IMAGE USING IMAGE BUILDER IN THE WEB CONSOLE

To create a blueprint for a RHEL for Edge image by using RHEL image builder in RHEL web console, ensure that you have met the following prerequisites and then follow the procedure.

Prerequisites

- On a RHEL system, you have opened the RHEL image builder dashboard.

Procedure

1. On the RHEL image builder dashboard, click **Create Blueprint**.
The **Create Blueprint** dialogue box opens.
2. On the **Details** page:
 - a. Enter the name of the blueprint and, optionally, its description. Click **Next**.
3. Optional: In the **Packages** page:
 - a. On the **Available packages** search, enter the package name and click the **>** button to move it to the Chosen packages field. Search and include as many packages as you want. Click **Next**.



NOTE

These customizations are all optional unless otherwise specified.

4. On the **Kernel** page, enter a kernel name and the command-line arguments.
5. On the **File system** page, select **Use automatic partitioning**. OSTree systems do not support filesystem customization, because OSTree images have their own mount rule, such as read-only. Click **Next**.
6. On the **Services** page, you can enable or disable services:
 - a. Enter the service names you want to enable or disable, separating them by a comma, by space, or by pressing the **Enter** key. Click **Next**.
7. On the **Firewall** page, set up your firewall setting:
 - a. Enter the **Ports**, and the firewall services you want to enable or disable.
 - b. Click the **Add zone** button to manage your firewall rules for each zone independently. Click **Next**.
8. On the **Users** page, add a users by following the steps:
 - a. Click **Add user**.
 - b. Enter a **Username**, a **password**, and a **SSH key**. You can also mark the user as a privileged user, by clicking the **Server administrator** checkbox. Click **Next**.

9. On the **Groups** page, add groups by completing the following steps:
 - a. Click the **Add groups** button:
 - i. Enter a **Group name** and a **Group ID**. You can add more groups. Click **Next**.
10. On the **SSH keys** page, add a key:
 - a. Click the **Add key** button.
 - i. Enter the SSH key.
 - ii. Enter a **User**. Click **Next**.
11. On the **Timezone** page, set your timezone settings:
 - a. On the **Timezone** field, enter the timezone you want to add to your system image. For example, add the following timezone format: "US/Eastern".
If you do not set a timezone, the system uses Universal Time, Coordinated (UTC) as default.
 - b. Enter the **NTP** servers. Click **Next**.
12. On the **Locale** page, complete the following steps:
 - a. On the **Keyboard** search field, enter the package name you want to add to your system image. For example: ["en_US.UTF-8"].
 - b. On the **Languages** search field, enter the package name you want to add to your system image. For example: "us". Click **Next**.
13. On the **Others** page, complete the following steps:
 - a. On the **Hostname** field, enter the hostname you want to add to your system image. If you do not add a hostname, the operating system determines the hostname.
 - b. Mandatory only for the Simplifier Installer image: On the **Installation Devices** field, enter a valid node for your system image. For example: **dev/sda**. Click **Next**.
14. Mandatory only when building FIDO images: On the **FIDO device onboarding** page, complete the following steps:
 - a. On the **Manufacturing server URL** field, enter the following information:
 - i. On the **DIUN public key insecure** field, enter the insecure public key.
 - ii. On the **DIUN public key hash** field, enter the public key hash.
 - iii. On the **DIUN public key root certs** field, enter the public key root certs. Click **Next**.
15. On the **OpenSCAP** page, complete the following steps:
 - a. On the **Datastream** field, enter the **datastream** remediation instructions you want to add to your system image.
 - b. On the **Profile ID** field, enter the **profile_id** security profile you want to add to your system image. Click **Next**.

16. Mandatory only when building Ignition images: On the **Ignition** page, complete the following steps:
 - a. On the **Firstboot URL** field, enter the package name you want to add to your system image.
 - b. On the **Embedded Data** field, drag or upload your file. Click **Next**.
17. . On the **Review** page, review the details about the blueprint. Click **Create**.

The RHEL image builder view opens, listing existing blueprints.

4.3. CREATING A RHEL FOR EDGE COMMIT IMAGE BY USING IMAGE BUILDER IN WEB CONSOLE

You can create a “RHEL for Edge Commit” image by using RHEL image builder in RHEL web console. The “RHEL for Edge Commit (.tar)” image type contains a full operating system, but it is not directly bootable. To boot the Commit image type, you must deploy it in a running container.

Prerequisites

- On a RHEL system, you have accessed the RHEL image builder dashboard.

Procedure

1. On the RHEL image builder dashboard click **Create Image**.
2. On the **Image output** page, perform the following steps:
 - a. From the **Select a blueprint** dropdown menu, select the blueprint you want to use.
 - b. From the **Image output type** dropdown list, select “RHEL for Edge Commit (.tar)” for network-based deployment.
 - c. Click **Next**.
 - d. On the **OSTree settings** page, enter:
 - i. **Repository URL**: specify the URL to the OSTree repository of the commit to embed in the image. For example, `http://10.0.2.2:8080/repo/`.
 - ii. **Parent commit**: specify a previous commit, or leave it empty if you do not have a commit at this time.
 - iii. In the **Ref** textbox, specify a reference path for where your commit is going to be created. By default, the web console specifies `rhel/8/$ARCH/edge`. The “\$ARCH” value is determined by the host machine. Click **Next**.
 - e. On the **Review** page, check the customizations and click **Create**.
RHEL image builder starts to create a RHEL for Edge Commit image for the blueprint that you created.



NOTE

The image creation process takes up to 20 minutes to complete.

Verification

1. To check the RHEL for Edge Commit image creation progress:

- a. Click the **Images** tab.

After the image creation process is complete, you can download the resulting “RHEL for Edge Commit (.tar)” image.

Additional resources

- [Downloading a RHEL for Edge image](#)

4.4. CREATING A RHEL FOR EDGE CONTAINER IMAGE BY USING RHEL IMAGE BUILDER IN RHEL WEB CONSOLE

You can create RHEL for Edge images by selecting “RHEL for Edge Container (.tar)”. The **RHEL for Edge Container (.tar)** image type creates an OSTree commit and embeds it into an OCI container with a web server. When the container is started, the web server serves the commit as an OSTree repository.

Follow the steps in this procedure to create a RHEL for Edge Container image using image builder in RHEL web console.

Prerequisites

- On a RHEL system, you have accessed the RHEL image builder dashboard.
- You have created a blueprint.

Procedure

1. On the RHEL image builder dashboard click **Create Image**.
2. On the **Image output** page, perform the following steps:
3. From the **Select a blueprint** dropdown menu, select the blueprint you want to use.
 - a. From the **Image output type** dropdown list, select “RHEL for Edge Container (.tar)” for network-based deployment.
 - b. Click **Next**.
 - c. On the **OSTree** page, enter:
 - i. **Repository URL:** specify the URL to the OSTree repository of the commit to embed in the image. For example, `http://10.0.2.2:8080/repo/`. By default, the repository folder for a RHEL for Edge Container image is `/repo`.
To find the correct URL to use, access the running container and check the **nginx.conf** file. To find which URL to use, access the running container and check the **nginx.conf** file. Inside the **nginx.conf** file, find the **root** directory entry to search for the `/repo/` folder information. Note that, if you do not specify a repository URL when creating a RHEL for Edge Container image **(.tar)** by using RHEL image builder, the default `/repo/` entry is created in the **nginx.conf** file.
 - ii. **Parent commit:** specify a previous commit, or leave it empty if you do not have a commit at this time.

- iii. In the **Ref** textbox, specify a reference path for where your commit is going to be created. By default, the web console specifies **rhel/8/\$ARCH/edge**. The "\$ARCH" value is determined by the host machine. Click **Next**.
 - d. On the **Review** page, check the customizations. Click **Save blueprint**.
4. Click **Create**.
RHEL image builder starts to create a RHEL for Edge Container image for the blueprint that you created.

**NOTE**

The image creation process takes up to 20 minutes to complete.

Verification

1. To check the RHEL for Edge Container image creation progress:
 - a. Click the **Images** tab.

After the image creation process is complete, you can download the resulting “**RHEL for Edge Container (.tar)**” image.

Additional resources

- [Downloading a RHEL for Edge image](#)

4.5. CREATING A RHEL FOR EDGE INSTALLER IMAGE BY USING IMAGE BUILDER IN RHEL WEB CONSOLE

You can create RHEL for Edge Installer images for non-network-based deployment by selecting **RHEL for Edge Installer (.iso)**. The **RHEL for Edge Installer (.iso)** image type pulls the OSTree commit repository from the running container served by the **RHEL for Edge Container (.tar)** and creates an installable boot ISO image with a Kickstart file that is configured to use the embedded OSTree commit.

Follow the steps in this procedure to create a RHEL for Edge image using image builder in RHEL web console.

Prerequisites

- On a RHEL system, you have accessed the image builder dashboard.
- You created a blueprint.
- You created a RHEL for Edge Container image and loaded it into a running container. See [Creating a RHEL for Edge Container image for non-network-based deployments](#) .

Procedure

1. On the RHEL image builder dashboard click **Create Image**.
2. On the **Image output** page, perform the following steps:
 - a. From the **Select a blueprint** dropdown menu, select the blueprint you want to use.

- b. From the **Image output type** dropdown list, select **RHEL for Edge Installer (.iso)** image.
 - c. Click **Next**.
 - d. On the **OSTree settings** page, enter:
 - i. **Repository URL**: specify the URL to the OSTree repository of the commit to embed in the image. For example, `http://10.0.2.2:8080/repo/`.
 - ii. In the **Ref** textbox, specify a reference path for where your commit is going to be created. By default, the web console specifies **rhel/8/\$ARCH/edge**. The "\$ARCH" value is determined by the host machine. Click **Next**.
 - e. On the **Review** page, check the customizations. Click **Save blueprint**.
3. Click **Create**.
RHEL image builder starts to create a RHEL for Edge Installer image for the blueprint that you created.

**NOTE**

The image creation process takes up to 20 minutes to complete.

Verification

After the image creation process is complete, you can download the resulting **RHEL for Edge Installer (.iso)** image.

1. To check the RHEL for Edge Installer image creation progress:
 - a. Click the **Images** tab.

After the image creation process is complete, you can download the resulting **RHEL for Edge Installer (.iso)** image and boot the ISO image into a device.

Additional resources

- [Downloading a RHEL for Edge image](#)

4.6. DOWNLOADING A RHEL FOR EDGE IMAGE

After you successfully create the RHEL for Edge image by using RHEL image builder, download the image on the local host.

Procedure

To download an image:

1. From the **More Options** menu, click **Download**.
The RHEL image builder tool downloads the file at your default download location.

The downloaded file consists of a **.tar** file with an OSTree repository for RHEL for Edge Commit and RHEL for Edge Container images, or a **.iso** file for RHEL for Edge Installer images, with an OSTree repository. This repository contains the commit and a **json** file which contains information metadata about the repository content.

4.7. ADDITIONAL RESOURCES

- [Composing a RHEL for Edge image using image builder command-line](#) .

CHAPTER 5. COMPOSING A RHEL FOR EDGE IMAGE USING IMAGE BUILDER COMMAND-LINE

You can use image builder to create a customized RHEL for Edge image (OSTree commit).

To access image builder and to create your custom RHEL for Edge image, you can either use the RHEL web console interface or the command-line interface.


For Network-based deployments, the workflow to compose RHEL for Edge images using the CLI, involves the following high-level steps:

1. Create a blueprint for RHEL for Edge image
2. Create a RHEL for Edge Commit image
3. Download the RHEL for Edge Commit image

For Non-Network-based deployments, the workflow to compose RHEL for Edge images using the CLI, involves the following high-level steps:

1. Create a blueprint for RHEL for Edge image
2. Create a blueprint for the RHEL for Edge Installer image
3. Create a RHEL for Edge Container image
4. Create a RHEL for Edge Installer image
5. Download the RHEL for Edge image

To perform the steps, use the **composer-cli** package.



NOTE

To run the **composer-cli** commands as non-root, you must be part of the **weldr** group or you must have administrator access to the system.

5.1. NETWORK-BASED DEPLOYMENTS WORKFLOW

This provides steps on how to build **OSTree** commits. These **OSTree** commits contain a full operating system, but are not directly bootable. To boot them, you need to deploy them using a Kickstart file.

5.1.1. Creating a RHEL for Edge Commit image blueprint using image builder command-line interface

Create a blueprint for RHEL for Edge Commit image using the CLI.

Prerequisite

- You do not have an existing blueprint. To verify that, list the existing blueprints:

```
$ sudo composer-cli blueprints list
```

Procedure

1. Create a plain text file in the TOML format, with the following content:

```
name = "blueprint-name"
description = "blueprint-text-description"
version = "0.0.1"
modules = [ ]
groups = [ ]
```

Where,

- *blueprint-name* is the name and *blueprint-text-description* is the description for your blueprint.
- *0.0.1* is the version number according to the Semantic Versioning scheme.
- *Modules* describe the package name and matching version glob to be installed into the image, for example, the package name = "tmux" and the matching version glob is version = "2.9a".
Notice that currently there are no differences between packages and modules.
- *Groups* are packages groups to be installed into the image, for example the group package *anaconda-tools*.
At this time, if you do not know the modules and groups, leave them empty.

2. Include the required packages and customize the other details in the blueprint to suit your requirements.

For every package that you want to include in the blueprint, add the following lines to the file:

```
[[packages]]
name = "package-name"
version = "package-version"
```

Where,

- *package-name* is the name of the package, such as *httpd*, *gdb-doc*, or *coreutils*.
- *package-version* is the version number of the package that you want to use.
The *package-version* supports the following *dnf* version specifications:
- For a specific version, use the exact version number such as *8.0*.
- For the latest available version, use the asterisk ***.
- For the latest minor version, use formats such as *8.**.

3. Push (import) the blueprint to the RHEL image builder server:

```
# composer-cli blueprints push blueprint-name.toml
```

4. List the existing blueprints to check whether the created blueprint is successfully pushed and exists.

```
# composer-cli blueprints show BLUEPRINT-NAME
```

5. Check whether the components and versions listed in the blueprint and their dependencies are valid:

```
# composer-cli blueprints depsolve blueprint-name
```

Additional resources

- [Supported Image Customizations](#)

5.1.2. Creating a RHEL for Edge Commit image using image builder command-line interface

To create a RHEL for Edge Commit image by using RHEL image builder command-line interface, ensure that you have met the following prerequisites and follow the procedure.

Prerequisites

- You have created a blueprint for RHEL for Edge Commit image.

Procedure

1. Create the RHEL for Edge Commit image.

```
# composer-cli compose start blueprint-name image-type
```

Where,


- *blueprint-name* is the RHEL for Edge blueprint name.
- *image-type* is **edge-commit** for **network-based deployment**.
A confirmation that the composer process has been added to the queue appears. It also shows a Universally Unique Identifier (UUID) number for the image created. Use the UUID number to track your build. Also keep the UUID number handy for further tasks.

2. Check the image compose status.

```
# composer-cli compose status
```

The output displays the status in the following format:

```
<UUID> RUNNING date blueprint-name blueprint-version image-type
```



NOTE

The image creation process takes up to 20 minutes to complete.

To interrupt the image creation process, run:

```
# composer-cli compose cancel <UUID>
```

To delete an existing image, run:

```
# composer-cli compose delete <UUID>
```

After the image is ready, you can download it and use the image on your **network deployments**.

Additional resources

- [Composing a RHEL for Edge image using RHEL image builder command-line](#)

5.1.3. Creating a RHEL for Edge image update with a ref commit by using RHEL image builder CLI

If you performed a change in an existing blueprint, for example, you added a new package, and you want to update an existing RHEL for Edge image with this new package, you can use the **--parent** argument to generate an updated **RHEL for Edge Commit (.tar)** image. The **--parent** argument can be either a **ref** that exists in the repository specified by the **URL** argument, or you can use the **Commit ID**, which you can find in the extracted **.tar** image file. Both **ref** and **Commit ID** arguments retrieve a parent for the new commit that you are building. RHEL image builder can read information from the parent commit that will affect parts of the new commit that you are building. As a result, RHEL image builder reads the parent commit's user database and preserves UIDs and GIDs for the package-created system users and groups.

Prerequisites

- You have updated an existing blueprint for RHEL for Edge image.
- You have an existing RHEL for Edge image (OSTree commit). See [Extracting RHEL for Edge image commit](#).
- The **ref** being built is available at the **OSTree** repository specified by the URL.

Procedure

1. Create the RHEL for Edge commit image:

```
# composer-cli compose start-ostree --ref rhel/8/x86_64/edge --parent parent-OSTree-REF -url URL blueprint-name image-type
```

For example:

- To create a new RHEL for Edge commit based on a **parent** and with a new **ref**, run the following command:

```
# composer-cli compose start-ostree --ref rhel/8/x86_64/edge --parent rhel/8/x86_64/edge --url http://10.0.2.2:8080/repo rhel_update edge-commit
```

- To create a new RHEL for Edge commit based on the same **ref**, run the following command:

```
# composer-cli compose start-ostree --ref rhel/8/x86_64/edge --url http://10.0.2.2:8080/repo rhel_update edge-commit
```

Where:

- The **--ref** argument specifies the same path value that you used to build an OSTree repository.

- The `--parent` argument specifies the parent commit. It can be ref to be resolved and pulled, for example **rhel/8/x86_64/edge**, or the **Commit ID** that you can find in the extracted **.tar** file.
- *blueprint-name* is the RHEL for Edge blueprint name.
- The `--url` argument specifies the URL to the OSTree repository of the commit to embed in the image, for example, `http://10.0.2.2:8080/repo`.
- *image-type* is **edge-commit** for **network-based deployment**.



NOTE

- The `--parent` argument can only be used for the **RHEL for Edge Commit (.tar)** image type. Using the `--url` and `--parent` arguments together results in errors with the **RHEL for Edge Container (.tar)** image type.
- If you omit the **parent ref** argument, the system falls back to the **ref** specified by the `--ref` argument.

A confirmation that the composer process has been added to the queue appears. It also shows a Universally Unique Identifier (UUID) number for the image created. Use the UUID number to track your build. Also keep the UUID number handy for further tasks.

2. Check the image compose status.

```
# composer-cli compose status
```

The output displays the status in the following format:

```
<UUID> RUNNING date blueprint-name blueprint-version image-type
```



NOTE

The image creation process takes a few minutes to complete.

(Optional) To interrupt the image creation process, run:

```
# composer-cli compose cancel <UUID>
```

(Optional) To delete an existing image, run:

```
# composer-cli compose delete <UUID>
```

After the image creation is complete, to upgrade an existing OSTree deployment, you need:

- Set up a repository. See [Deploying a RHEL for Edge image](#) .
- Add this repository as a remote, that is, the http or https endpoint that hosts the OSTree content.

- Pull the new OSTree commit onto their existing running instance. See [Deploying RHEL for Edge image updates manually](#).

Additional resources

- [Creating a system image with RHEL image builder in the command-line interface](#)
- [Downloading a RHEL for Edge image by using the RHEL image builder command-line interface](#)

5.1.4. Downloading a RHEL for Edge image using the image builder command-line interface

To download a RHEL for Edge image by using RHEL image builder command line interface, ensure that you have met the following prerequisites and then follow the procedure.

Prerequisites

- You have created a RHEL for Edge image.

Procedure

1. Review the RHEL for Edge image status.

```
# composer-cli compose status
```

The output must display the following:

```
$ <UUID> FINISHED date blueprint-name blueprint-version image-type
```

2. Download the image.

```
# composer-cli compose image <UUID>
```

RHEL image builder downloads the image as a **tar** file to the current directory.

The UUID number and the image size is displayed alongside.

```
$ <UUID>-commit.tar: size MB
```

The image contains a commit and a **json** file with information metadata about the repository content.

Additional resources

- [Deploying a RHEL for Edge image in a network-based environment](#)

5.2. NON-NETWORK-BASED DEPLOYMENTS WORKFLOW

To build a boot ISO image that installs an OSTree-based system using the "RHEL for Edge Container" and the "RHEL for Edge Installer" images and that can be later deployed to a device in disconnected environments, follow the steps.

5.2.1. Creating a RHEL for Edge Container blueprint by using image builder CLI

To create a blueprint for RHEL for Edge Container image, perform the following steps:

Procedure

1. Create a plain text file in the TOML format, with the following content:

```
name = "blueprint-name"
description = "blueprint-text-description"
version = "0.0.1"
modules = [ ]
groups = [ ]
```

Where,

- *blueprint-name* is the name and *blueprint-text-description* is the description for your blueprint.
 - *0.0.1* is the version number according to the Semantic Versioning scheme.
 - *Modules* describe the package name and matching version glob to be installed into the image, for example, the package name = "tmux" and the matching version glob is version = "2.9a".
Notice that currently there are no differences between packages and modules.
 - *Groups* are packages groups to be installed into the image, for example the group package *anaconda-tools*.
At this time, if you do not know the modules and groups, leave them empty.
2. Include the required packages and customize the other details in the blueprint to suit your requirements.

For every package that you want to include in the blueprint, add the following lines to the file:

```
[[packages]]
name = "package-name"
version = "package-version"
```

Where,

- *package-name* is the name of the package, such as *httpd*, *gdb-doc*, or *coreutils*.
 - *package-version* is the version number of the package that you want to use.
The *package-version* supports the following dnf version specifications:
 - For a specific version, use the exact version number such as *8.0*.
 - For the latest available version, use the asterisk ***.
 - For the latest minor version, use formats such as *8.**.
3. Push (import) the blueprint to the RHEL image builder server:

```
# composer-cli blueprints push blueprint-name.toml
```

4. List the existing blueprints to check whether the created blueprint is successfully pushed and exists.

```
# composer-cli blueprints show BLUEPRINT-NAME
```

5. Check whether the components and versions listed in the blueprint and their dependencies are valid:

```
# composer-cli blueprints depsolve blueprint-name
```

Additional resources

- [Supported Image Customizations](#)

5.2.2. Creating a RHEL for Edge Installer blueprint using image builder CLI

You can create a blueprint to build a **RHEL for Edge Installer (.iso)** image, and specify user accounts to automatically create one or more users on the system at installation time.



WARNING

When you create a user in the blueprint with the **customizations.user** customization, the blueprint creates the user under the **/usr/lib/passwd** directory and the password, under the **/usr/etc/shadow** directory. Note that you cannot change the password in further versions of the image in a running system using **OSTree** updates. The users you create with blueprints must be used only to gain access to the created system. After you access the system, you need to create users, for example, using the **useradd** command.

To create a blueprint for RHEL for Edge Installer image, perform the following steps:

Procedure

1. Create a plain text file in the TOML format, with the following content:

```
name = "blueprint-installer"
description = "blueprint-for-installer-image"
version = "0.0.1"

[[customizations.user]]
name = "user"
description = "account"
password = "user-password"
key = "user-ssh-key"
home = "path"
groups = ["user-groups"]
```

Where,

- *blueprint-name* is the name and *blueprint-text-description* is the description for your blueprint.

- 0.0.1 is the version number according to the Semantic Versioning scheme.
2. Push (import) the blueprint to the RHEL image builder server:

```
# composer-cli blueprints push blueprint-name.toml
```

3. List the existing blueprints to check whether the created blueprint is successfully pushed and exists.

```
# composer-cli blueprints show blueprint-name
```

4. Check whether the components and versions listed in the blueprint and their dependencies are valid:

```
# composer-cli blueprints depsolve blueprint-name
```

Additional resources

- [Supported Image Customizations](#)

5.2.3. Creating a RHEL for Edge Container image by using image builder CLI

To create a RHEL for Edge Container image by using RHEL image builder command-line interface, ensure that you have met the following prerequisites and follow the procedure.

Prerequisites

- You have created a blueprint for RHEL for Edge Container image.

Procedure

1. Create the RHEL for Edge Container image.

```
# composer-cli compose start-ostree --ref rhel/8/x86_64/edge --url URL-OSTree-repository blueprint-name image-type
```

Where,

- **--ref** is the same value that customer used to build ostree repository
- **--url** is the URL to the OSTree repository of the commit to embed in the image. For example, `http://10.0.2.2:8080/repo/`. By default, the repository folder for a RHEL for Edge Container image is `/repo`. See [Setting up a web server to install RHEL for Edge image](#) . To find the correct URL to use, access the running container and check the **nginx.conf** file. To find which URL to use, access the running container and check the **nginx.conf** file. Inside the **nginx.conf** file, find the **root** directory entry to search for the **/repo/** folder information. Note that, if you do not specify a repository URL when creating a RHEL for Edge Container image (**.tar**) by using RHEL image builder, the default **/repo/** entry is created in the **nginx.conf** file.
- *blueprint-name* is the RHEL for Edge blueprint name.
- *image-type* is **edge-container** for **non-network-based deployment**


A confirmation that the composer process has been added to the queue appears. It also shows a Universally Unique Identifier (UUID) number for the image created. Use the UUID number to track your build. Also keep the UUID number handy for further tasks.

2. Check the image compose status.

```
# composer-cli compose status
```

The output displays the status in the following format:

```
<UUID> RUNNING date blueprint-name blueprint-version image-type
```



NOTE

The image creation process takes up to 20 minutes to complete.

To interrupt the image creation process, run:

```
# composer-cli compose cancel <UUID>
```

To delete an existing image, run:

```
# composer-cli compose delete <UUID>
```

After the image is ready, it can be used for **non-network deployments**. See [Creating a RHEL for Edge Container image for non-network-based deployments](#).

Additional resources

- [Composing a RHEL for Edge image using RHEL image builder command-line](#)

5.2.4. Creating a RHEL for Edge Installer image using command-line interface for non-network-based deployments

To create a RHEL for Edge Installer image that embeds the **OSTree** commit, use the RHELImage builder command-line interface, and ensure that you have met the following prerequisites and then follow the procedure.

Prerequisites

- You have created a blueprint for RHEL for Edge Installer image.
- You have created a RHEL for Edge Edge Container image and deployed it using a web server.

Procedure

1. Begin to create the RHEL for Edge Installer image.

```
# composer-cli compose start-ostree --ref rhel/8/x86_64/edge --url URL-OSTree-repository  
blueprint-name image-type
```

Where,

- *ref* is the same value that customer used to build ostree repository
- *URL-OSTree-repository* is the URL to the OSTree repository of the commit to embed in the image. For example, <http://10.0.2.2:8080/repo>. See [Creating a RHEL for Edge Container image for non-network-based deployments](#).
- *blueprint-name* is the RHEL for Edge Installer blueprint name.
- *image-type* is **edge-installer**.
A confirmation that the composer process has been added to the queue appears. It also shows a Universally Unique Identifier (UUID) number for the image created. Use the UUID number to track your build. Also keep the UUID number handy for further tasks.

2. Check the image compose status.

```
# composer-cli compose status
```

The command output displays the status in the following format:

```
<UUID> RUNNING date blueprint-name blueprint-version image-type
```



NOTE

The image creation process takes a few minutes to complete.

To interrupt the image creation process, run:

```
# composer-cli compose cancel <UUID>
```

To delete an existing image, run:

```
# composer-cli compose delete <UUID>
```

After the image is ready, you can use it for **non-network deployments**. See [Installing the RHEL for Edge image for non-network-based deployments](#).

5.2.5. Downloading a RHEL for Edge Installer image using the image builder CLI

To download a RHEL for Edge Installer image by using RHEL image builder command line interface, ensure that you have met the following prerequisites and then follow the procedure.

Prerequisites

- You have created a RHEL for Edge Installer image.

Procedure

1. Review the RHEL for Edge image status.

```
# composer-cli compose status
```

The output must display the following:

```
■
```

```
$ <UUID> FINISHED date blueprint-name blueprint-version image-type
```

2. Download the image.

```
# composer-cli compose image <UUID>
```

RHEL image builder downloads the image as an **.iso** file to the current directory.

The UUID number and the image size is displayed alongside.

```
$ <UUID>-boot.iso: size MB
```

The resulting image is a bootable ISO image.

Additional resources

- [Deploying a RHEL for Edge image in a non-network-based environment](#) .

5.3. SUPPORTED IMAGE CUSTOMIZATIONS

You can customize your image by adding customizations to your blueprint, such as:

- Adding an additional RPM package
- Enabling a service
- Customizing a kernel command line parameter.

Between others. You can use several image customizations within blueprints. By using the customizations, you can add packages and groups to the image that are not available in the default packages. To use these options, configure the customizations in the blueprint and import (push) it to RHEL image builder.

Additional resources

- [Blueprint import fails after adding filesystem customization "size"](#) .

5.3.1. Selecting a distribution

You can use the **distro** field to select the distribution to use when composing your images, or depsolving the blueprint. If **distro** is left blank it will use the host distribution. If you do not specify a distribution, the blueprint uses the host distribution. In case you upgrade the host operating system, the blueprints with no distribution set build images using the new operating system version. You cannot build an operating system image that differs from the RHEL image builder host.

Procedure

- Customize the blueprint with the **distro** to always build the specified RHEL image:

```
name = "blueprint_name"
description = "blueprint_version"
version = "0.1"
distro = "different_minor_version"
```

Replace "*different_minor_version*" to build a different minor version, for example, if you want to build a RHEL 8.8 image, use **distro** = "rhel-88". On RHEL 8.9 image, you can build minor versions such as RHEL 8.8 and earlier releases.

5.3.2. Selecting a package group

Customize the blueprint with packages and modules. The **name** attribute is a required string. The **version** attribute is an optional string that, if not provided, uses the latest version in the repositories.



NOTE

Currently, there are no differences between packages and modules in **osbuild-composer**. Both are treated as an RPM package dependency.

Procedure

- Customize your blueprint with a package:

```
[[packages]]
name = "package_group_name"
```

Replace "*package_group_name*" with the name of the group. For example, "tmux".

```
[[packages]]
name = "tmux"
version = "2.9a"
```

5.3.3. Setting the image hostname

The **customizations.hostname** is an optional string that you can use to configure the final image hostname. This customization is optional, and if you do not set it, the blueprint uses the default hostname.

Procedure

- Customize the blueprint to configure the hostname:

```
[customizations]
hostname = "baseimage"
```

5.3.4. Specifying additional users

Add a user to the image, and optionally, set their SSH key. All fields for this section are optional except for the **name**.

Procedure

- Customize the blueprint to add a user to the image:

```
[[customizations.user]]
name = "USER-NAME"
description = "USER-DESCRIPTION"
```

```
password = "PASSWORD-HASH"  
key = "PUBLIC-SSH-KEY"  
home = "/home/USER-NAME/"  
shell = "/usr/bin/bash"  
groups = ["users", "wheel"]  
uid = NUMBER  
gid = NUMBER
```

The GID is optional and must already exist in the image. Optionally, a package creates it, or the blueprint creates the GID by using the `[[customizations.group]]` entry.

Replace `PASSWORD-HASH` with the actual **password hash**. To generate the **password hash**, use a command such as:

```
$ python3 -c 'import crypt;pw=getpass.getpass();print(crypt.crypt(pw) if  
(pw==getpass.getpass("Confirm: "))) else exit()'
```

Replace the other placeholders with suitable values.

Enter the **name** value and omit any lines you do not need.

Repeat this block for every user to include.

5.3.5. Specifying additional groups

Specify a group for the resulting system image. Both the **name** and the **gid** attributes are mandatory.

Procedure

- Customize the blueprint with a group:

```
[[customizations.group]]  
name = "GROUP-NAME"  
gid = NUMBER
```

Repeat this block for every group to include.

5.3.6. Setting SSH key for existing users

You can use **customizations.sshkey** to set an SSH key for the existing users in the final image. Both **user** and **key** attributes are mandatory.

Procedure

- Customize the blueprint by setting an SSH key for existing users:

```
[[customizations.sshkey]]  
user = "root"  
key = "PUBLIC-SSH-KEY"
```




NOTE

You can only configure the **customizations.sshkey** customization for existing users. To create a user and set an SSH key, see the **User specifications for the resulting system image** customization.

5.3.7. Appending a kernel argument

You can append arguments to the boot loader kernel command line. By default, RHEL image builder builds a default kernel into the image. However, you can customize the kernel by configuring it in the blueprint.

Procedure

- Append a kernel boot parameter option to the defaults:

```
[customizations.kernel]
append = "KERNEL-OPTION"
```

- Define a kernel name to use in an image

```
[customizations.kernel]
name = "KERNEL-rt"
```

5.3.8. Setting time zone and NTP

You can customize your blueprint to configure the time zone and the *Network Time Protocol* (NTP). Both **timezone** and **ntpserver**s attributes are optional strings. If you do not customize the time zone, the system uses *Universal Time, Coordinated* (UTC). If you do not set NTP servers, the system uses the default distribution.

Procedure

- Customize the blueprint with the **timezone** and the **ntpserver**s you want:

```
[customizations.timezone]
timezone = "TIMEZONE"
ntpserver = "NTP_SERVER"
```

For example:

```
[customizations.timezone]
timezone = "US/Eastern"
ntpserver = ["0.north-america.pool.ntp.org", "1.north-america.pool.ntp.org"]
```



NOTE

Some image types, such as Google Cloud, already have NTP servers set up. You cannot override it because the image requires the NTP servers to boot in the selected environment. However, you can customize the time zone in the blueprint.

5.3.9. Customizing the locale settings

You can customize the locale settings for your resulting system image. Both **language** and the **keyboard** attributes are mandatory. You can add many other languages. The first language you add is the primary language and the other languages are secondary.

Procedure

- Set the locale settings:

```
[customizations.locale]
languages = ["LANGUAGE"]
keyboard = "KEYBOARD"
```

For example:

```
[customizations.locale]
languages = ["en_US.UTF-8"]
keyboard = "us"
```

- To list the values supported by the languages, run the following command:

```
$ localectl list-locales
```

- To list the values supported by the keyboard, run the following command:

```
$ localectl list-keymaps
```

5.3.10. Customizing firewall

Set the firewall for the resulting system image. By default, the firewall blocks incoming connections, except for services that enable their ports explicitly, such as **sshd**.

If you do not want to use the **[customizations.firewall]** or the **[customizations.firewall.services]**, either remove the attributes, or set them to an empty list []. If you only want to use the default firewall setup, you can omit the customization from the blueprint.



NOTE

The Google and OpenStack templates explicitly disable the firewall for their environment. You cannot override this behavior by setting the blueprint.

Procedure

- Customize the blueprint with the following settings to open other ports and services:

```
[customizations.firewall]
ports = ["PORTS"]
```

Where **ports** is an optional list of strings that contain ports or a range of ports and protocols to open. You can configure the ports by using the following format: **port:protocol** format. You can configure the port ranges by using the **portA-portB:protocol** format. For example:

```
[customizations.firewall]
ports = ["22:tcp", "80:tcp", "imap:tcp", "53:tcp", "53:udp", "30000-32767:tcp", "30000-32767:udp"]
```

You can use numeric ports, or their names from the **/etc/services** to enable or disable port lists.

- Specify which firewall services to enable or disable in the **customizations.firewall.service** section:

```
[customizations.firewall.services]
enabled = ["SERVICES"]
disabled = ["SERVICES"]
```

- You can check the available firewall services:

```
$ firewall-cmd --get-services
```

For example:

```
[customizations.firewall.services]
enabled = ["ftp", "ntp", "dhcp"]
disabled = ["telnet"]
```



NOTE

The services listed in **firewall.services** are different from the **service-names** available in the **/etc/services** file.

5.3.11. Enabling or disabling services

You can control which services to enable during the boot time. Some image types already have services enabled or disabled to ensure that the image works correctly and you cannot override this setup. The **[customizations.services]** settings in the blueprint do not replace these services, but add the services to the list of services already present in the image templates.

Procedure

- Customize which services to enable during the boot time:

```
[customizations.services]
enabled = ["SERVICES"]
disabled = ["SERVICES"]
```

For example:

```
[customizations.services]
enabled = ["sshd", "cockpit.socket", "httpd"]
disabled = ["postfix", "telnetd"]
```

5.3.12. Specifying a custom filesystem configuration

You can specify a custom filesystem configuration in your blueprints and therefore create images with a specific disk layout, instead of the default layout configuration. By using the non-default layout configuration in your blueprints, you can benefit from:

- Security benchmark compliance
- Protection against out-of-disk errors
- Improved performance
- Consistency with existing setups



NOTE

The filesystem customization is not supported for OSTree systems, because OSTree images have their own mount rule, such as read-only.

The blueprint supports the following **mountpoints** and their sub-directories:

- `/` - the root mount point
- `/var`
- `/home`
- `/opt`
- `/srv`
- `/usr`
- `/app`
- `/data`
- `/boot` - The blueprint supports the `/boot` customization from RHEL 8.7 and RHEL 9.1 onward.



NOTE

Customizing mount points is only supported from RHEL 8.5 distributions onward, by using the CLI. In earlier distributions, you can only specify the **root** partition as a mount point and specify the **size** argument as an alias for the image size. Beginning with RHEL 8.6, for the **osbuild-composer-46.1-1.el8** RPM and later version, the physical partitions are no longer available and filesystem customizations create logical volumes.

If you have more than one partition in the customized image, you can create images with a customized file system partition on LVM and resize those partitions at runtime. To do this, you can specify a customized filesystem configuration in your blueprint and therefore create images with the required disk layout. The default filesystem layout remains unchanged - if you use plain images without file system customization, and **cloud-init** resizes the root partition.

The blueprint automatically converts the file system customization to an LVM partition.

You can use the custom file blueprint customization to create new files or to replace existing files. The parent directory of the file you specify must exist, otherwise, the image build fails. Ensure that the parent directory exists by specifying it in the **[[customizations.directories]]** customization.

**WARNING**

If you combine the files customizations with other blueprint customizations, it might affect the functioning of the other customizations, or it might override the current files customizations.

With the **[[customizations.files]]** blueprint customization you can:

- Create new text files.
- Modifying existing files. **WARNING:** this can override the existing content.
- Set user and group ownership for the file you are creating.
- Set the mode permission in the octal format.

You cannot create or replace the following files:

- **/etc/fstab**
- **/etc/shadow**
- **/etc/passwd**
- **/etc/group**

You can create customized files and directories in your image, by using the **[[customizations.files]]** and the **[[customizations.directories]]** blueprint customizations. You can use these customizations only in the **/etc** directory.

**NOTE**

These blueprint customizations are supported by all image types, except the image types that deploy OSTree commits, such as **edge-raw-image**, **edge-installer**, and **edge-simplified-installer**.

**WARNING**

If you use the **customizations.directories** with a directory path which already exists in the image with **mode**, **user** or **group** already set, the image build fails to prevent changing the ownership or permissions of the existing directory.

With the **[[customizations.directories]]** blueprint customization you can:

- Create new directories.
- Set user and group ownership for the directory you are creating.

- Set the directory mode permission in the octal format.
- Ensure that parent directories are created as needed.

With the **[[customizations.files]]** blueprint customization you can:

- Create new text files.
- Modifying existing files. WARNING: this can override the existing content.
- Set user and group ownership for the file you are creating.
- Set the mode permission in the octal format.



NOTE

You cannot create or replace the following files:

- **/etc/fstab**
- **/etc/shadow**
- **/etc/passwd**
- **/etc/group**

Procedure

- Customize the filesystem configuration in your blueprint:

```
[[customizations.filesystem]]
mountpoint = "MOUNTPOINT"
size = MINIMUM-PARTITION-SIZE
```

The **MINIMUM-PARTITION-SIZE** value has no default size format. The blueprint customization supports the following values and units: kB to TB and KiB to TiB. For example, you can define the mount point size in bytes:

```
[[customizations.filesystem]]
mountpoint = "/var"
size = 1073741824
```

- Define the mount point size by using units. For example:

```
[[customizations.filesystem]]
mountpoint = "/opt"
size = "20 GiB"
```

```
[[customizations.filesystem]]
mountpoint = "/boot"
size = "1 GiB"
```

- Create customized directories under the **/etc** directory for your image by using **[[customizations.directories]]**:

```
[[customizations.directories]]
path = "/etc/directory_name"
mode = "octal_access_permission"
user = "user_string_or_integer"
group = "group_string_or_integer"
ensure_parents = boolean
```

The blueprint entries are described as following:

- **path** - Mandatory - enter the path to the directory that you want to create. It must be an absolute path under the **/etc** directory.
- **mode** - Optional - set the access permission on the directory, in the octal format. If you do not specify a permission, it defaults to 0755. The leading zero is optional.
- **user** - Optional - set a user as the owner of the directory. If you do not specify a user, it defaults to **root**. You can specify the user as a string or as an integer.
- **group** - Optional - set a group as the owner of the directory. If you do not specify a group, it defaults to **root**. You can specify the group as a string or as an integer.
- **ensure_parents** - Optional - Specify whether you want to create parent directories as needed. If you do not specify a value, it defaults to **false**.
- Create customized file under the **/etc** directory for your image by using **[[customizations.files]]**:

```
[[customizations.files]]
path = "/etc/directory_name"
mode = "octal_access_permission"
user = "user_string_or_integer"
group = "group_string_or_integer"
data = "Hello world!"
```

The blueprint entries are described as following:

- **path** - Mandatory - enter the path to the file that you want to create. It must be an absolute path under the **/etc** directory.
- **mode** - Optional - set the access permission on the file, in the octal format. If you do not specify a permission, it defaults to 0644. The leading zero is optional.
- **user** - Optional - set a user as the owner of the file. If you do not specify a user, it defaults to **root**. You can specify the user as a string or as an integer.
- **group** - Optional - set a group as the owner of the file. If you do not specify a group, it defaults to **root**. You can specify the group as a string or as an integer.
- **data** - Optional - Specify the content of a plain text file. If you do not specify a content, it creates an empty file.

5.4. PACKAGES INSTALLED BY RHEL IMAGE BUILDER

When you create a system image using RHEL image builder, the system installs a set of base package groups.

**NOTE**

When you add additional components to your blueprint, ensure that the packages in the components you added do not conflict with any other package components. Otherwise, the system fails to solve dependencies and creating your customized image fails. You can check if there is no conflict between the packages by running the command:

```
# composer-cli blueprints depsolve BLUEPRINT-NAME
```

By default, RHEL image builder uses the **Core** group as the base list of packages.

Table 5.1. Default packages to support image type creation

Image type	Default Packages
ami	checkpolicy, chrony, cloud-init, cloud-utils-growpart, @Core, dhcp-client, gdisk, insights-client, kernel, langpacks-en, net-tools, NetworkManager, redhat-release, redhat-release-eula, rng-tools, rsync, selinux-policy-targeted, tar, yum-utils
openstack	@core, langpacks-en
qcow2	@core, chrony, dnf, kernel, yum, nfs-utils, dnf-utils, cloud-init, python3-jsonschema, qemu-guest-agent, cloud-utils-growpart, dracut-norescue, tar, tcpdump, rsync, dnf-plugin-spacewalk, rhn-client-tools, rhnlib, rhnsd, rhn-setup, NetworkManager, dhcp-client, cockpit-ws, cockpit-system, subscription-manager-cockpit, redhat-release, redhat-release-eula, rng-tools, insights-client
tar	policycoreutils, selinux-policy-targeted
vhd	@core, langpacks-en
vmdk	@core, chrony, cloud-init, firewallld, langpacks-en, open-vm-tools, selinux-policy-targeted

Image type	Default Packages
edge-commit	attr, audit, basesystem, bash, bash-completion, chrony, clevis, clevis-dracut, clevis-luks, container-selinux, coreutils, criu, cryptsetup, curl, dnsmasq, dosfstools, dracut-config-generic, dracut-network, e2fsprogs, firewalld, fuse-overlayfs, fwupd, glibc, glibc-minimal-langpack, gnupg2, greenboot, gzip, hostname, ima-evm-utils, iproute, iptables, iputils, keyutils, less, lvm2, NetworkManager, NetworkManager-wifi, NetworkManager-wwan, nss-altfiles, openssh-clients, openssh-server, passwd, pinentry, platform-python, podman, policycoreutils, policycoreutils-python-utils, polkit, procps-ng, redhat-release, rootfiles, rpm, rpm-ostree, rsync, selinux-policy-targeted, setools-console, setup, shadow-utils, shadow-utils, skopeo, slirp4netns, sudo, systemd, tar, tmux, traceroute, usbguard, util-linux, vim-minimal, wpa_supplicant, xz
edge-container	dnf, dosfstools, e2fsprogs, glibc, lorax-templates-generic, lorax-templates-rhel, lvm2, policycoreutils, python36, python3-iniparse, qemu-img, selinux-policy-targeted, systemd, tar, xfsprogs, xz

Image type	Default Packages
edge-installer	aajohan-comfortaa-fonts, abattis-cantarell-fonts, alsa-firmware, alsa-tools-firmware, anaconda, anaconda-install-env-deps, anaconda-widgets, audit, bind-utils, bitmap-fangsongti-fonts, bzip2, cryptsetup, dbus-x11, dejavu-sans-fonts, dejavu-sans-mono-fonts, device-mapper-persistent-data, dnf, dump, ethtool, fcoe-utils, ftp, gdb-gdbserver, gdisk, gfs2-utils, glibc-all-langpacks, google-noto-sans-cjk-ttc-fonts, gsettings-desktop-schemas, hdparm, hexedit, initscripts, ipmitool, iwl3945-firmware, iwl4965-firmware, iwl6000g2a-firmware, iwl6000g2b-firmware, jomolhari-fonts, kacst-farsi-fonts, kacst-qurn-fonts, kbd, kbd-misc, kdump-anaconda-addon, khmeros-base-fonts, libblockdev-lvm-dbus, libertas-sd8686-firmware, libertas-sd8787-firmware, libertas-usb8388-firmware, libertas-usb8388-olpc-firmware, libibverbs, libreport-plugin-bugzilla, libreport-plugin-reportuploader, libreport-rhel-anaconda-bugzilla, librsvg2, linux-firmware, lklug-fonts, lldpad, lohit-assamese-fonts, lohit-bengali-fonts, lohit-devanagari-fonts, lohit-gujarati-fonts, lohit-gurmukhi-fonts, lohit-kannada-fonts, lohit-odia-fonts, lohit-tamil-fonts, lohit-telugu-fonts, lsof, madan-fonts, metacity, mtr, mt-st, net-tools, nmap-ncat, nm-connection-editor, nss-tools, openssh-server, oscap-anaconda-addon, pciutils, perl-interpreter, pigz, python3-pyatspi, rdma-core, redhat-release-eula, rpm-ostree, rsync, rsyslog, sg3_utils, sil-abyssinica-fonts, sil-padauk-fonts, sil-scheherazade-fonts, smartmontools, smc-meera-fonts, spice-vdagent, strace, system-storage-manager, thai-scalable-waree-fonts, tigervnc-server-minimal, tigervnc-server-module, udisks2, udisks2-iscsi, usbutils, vim-minimal, volume_key, wget, xfsdump, xorg-x11-drivers,xorg-x11-fonts-misc,xorg-x11-server-utils,xorg-x11-server-Xorg, xorg-x11-xauth

Image type	Default Packages
edge-simplified-installer	attr, basesystem, binutils, bsdtar, clevis-dracut, clevis-luks, cloud-utils-growpart, coreos-installer, coreos-installer-dracut, coreutils, device-mapper-multipath, dnsmasq, dosfstools, dracut-live, e2fsprogs, fcoe-utils, fdo-init, gzip, ima-evm-utils, iproute, iptables, iputils, iscsi-initiator-utils, keyutils, lldpad, lvm2, passwd, policycoreutils, policycoreutils-python-utils, procps-ng, rootfiles, setools-console, sudo, traceroute, util-linux
image-installer	anaconda-dracut, curl, dracut-config-generic, dracut-network, hostname, iwl100-firmware, iwl1000-firmware, iwl105-firmware, iwl135-firmware, iwl2000-firmware, iwl2030-firmware, iwl3160-firmware, iwl5000-firmware, iwl5150-firmware, iwl6000-firmware, iwl6050-firmware, iwl7260-firmware, kernel, less, nfs-utils, openssh-clients, ostree, plymouth, prefixdevname, rng-tools, rpcbind, selinux-policy-targeted, systemd, tar, xfsprogs, xz
edge-raw-image	dnf, dosfstools, e2fsprogs, glibc, lorax-templates-generic, lorax-templates-rhel, lvm2, policycoreutils, python36, python3-iniparse, qemu-img, selinux-policy-targeted, systemd, tar, xfsprogs, xz
gce	@core, langpacks-en, acpid, dhcp-client, dnf-automatic, net-tools, python3, rng-tools, tar, vim

Additional resources

- [RHEL image builder description](#)

5.5. ADDITIONAL RESOURCES

- [Composing a RHEL for Edge image using image builder in RHEL web console](#) .

CHAPTER 6. BUILDING SIMPLIFIED INSTALLER IMAGES TO PROVISION A RHEL FOR EDGE IMAGE

You can build a RHEL for Edge Simplified Installer image, which is optimized for unattended installation to a device, and provision the image to a RHEL for Edge image.

6.1. SIMPLIFIED INSTALLER IMAGE BUILD AND DEPLOYMENT

Build a RHEL for Edge Simplified Installer image by using the **edge-simplified-installer** image type,.

To build a RHEL for Edge Simplified Installer image, provide an existing **OSTree** commit. The resulting simplified image contains a raw image that has the OSTree commit deployed. After you boot the Simplified installer ISO image, it provisions a RHEL for Edge system that you can use on a hard disk or as a boot image in a virtual machine. You can log in to the deployed system with the user name and password that you specified in the blueprint that you used to create the Simplified Installer image.

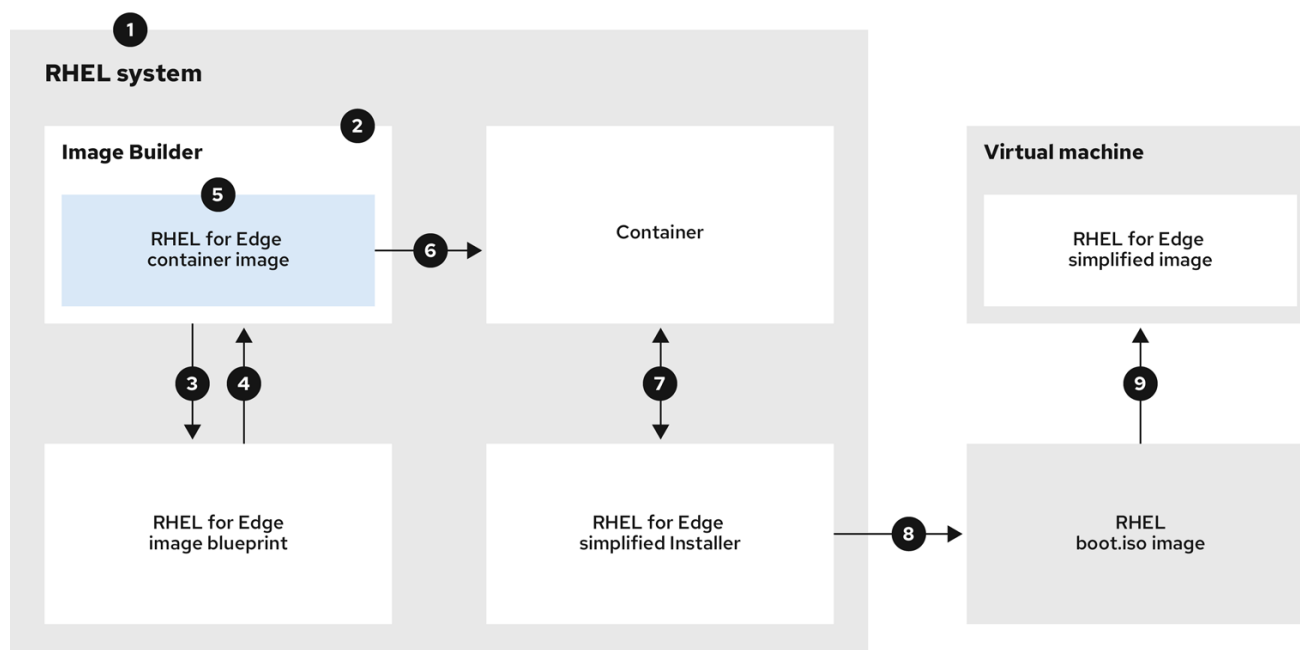
The RHEL for Edge Simplified Installer image is optimized for unattended installation to a device and supports both network-based deployment and non-network-based deployments. However, for network-based deployment, it supports only UEFI HTTP boot.

Composing and deploying a simplified RHEL for Edge image involves the following high-level steps:

1. Install and register a RHEL system
2. Install RHEL image builder
3. Using RHEL image builder, create a blueprint with customizations for RHEL for Edge Container image
4. Import the RHEL for Edge blueprint in RHEL image builder
5. Create a RHEL for Edge image embed in an OCI container with a web server ready to deploy the commit as an OSTree repository
6. Create a blueprint for the **edge-simplified-installer** image
7. Build a simplified RHEL for Edge image
8. Download the RHEL for Edge simplified image
9. Install the raw image with the **edge-simplified-installer** virt-install

The following diagram represents the RHEL for Edge Simplified building and provisioning workflow:

Figure 6.1. Building and provisioning RHEL for Edge in network-based environment



243_RHEL_0422

6.2. CREATING A BLUEPRINT FOR A SIMPLIFIED IMAGE USING RHEL IMAGE BUILDER CLI

To create a blueprint for a simplified RHEL for Edge image, you must customize it with a **device file** location to enable an unattended installation to a device and a **URL** to perform the initial device credential exchange. You also must specify users and user groups in the blueprint. For that, follow the steps:

Procedure

1. Create a plain text file in the Tom's Obvious, Minimal Language (TOML) format, with the following content:

```
name = "simplified-installer-blueprint"
description = "blueprint for the simplified installer image"
version = "0.0.1"
packages = []
modules = []
groups = []
distro = ""

[customizations]
installation_device = "/dev/vda"

[[customizations.user]]
name = "admin"
password = "admin"
groups = ["users", "wheel"]
```

```
[customizations.fdo]
manufacturing_server_url = "http://10.0.0.2:8080"
diun_pub_key_insecure = "true"
```



NOTE

The FDO customization in the blueprints is optional, and you can build your RHEL for Edge Simplified Installer image with no errors.

- *name* is the name and *description* is the description for your blueprint.
 - *0.0.1* is the version number according to the Semantic Versioning scheme.
 - *Modules* describe the package name and matching version glob to be installed into the image, for example, the package name = "tmux" and the matching version glob is version = "2.9a". Notice that currently there are no differences between packages and modules.
 - *Groups* are packages groups to be installed into the image, for example the **anaconda-tools** group package. If you do not know the modules and groups, leave them empty.
 - *installation-device* is the customization to enable an unattended installation to your device.
 - *manufacturing_server_url* is the URL to perform the initial device credential exchange.
 - *name* is the user name to login to the image.
 - *password* is a password of your choice.
 - *groups* are any user groups, such as "widget".
2. Push (import) the blueprint to the RHEL image builder server:

```
# composer-cli blueprints push blueprint-name.toml
```

3. List the existing blueprints to check whether the created blueprint is successfully pushed and exists.

```
# composer-cli blueprints show blueprint-name
```

4. Check whether the components and versions listed in the blueprint and their dependencies are valid:

```
# composer-cli blueprints depsolve blueprint-name
```

Additional resources

- [Composing a RHEL for Edge image using RHEL image builder command-line](#) .

6.3. CREATING A RHEL FOR EDGE SIMPLIFIED INSTALLER IMAGE USING IMAGE BUILDER CLI

To create a RHEL for Edge Simplified image by using RHEL image builder command-line interface, ensure that you have met the following prerequisites and then follow the procedure.

Prerequisites

- You created a blueprint for the RHEL for Edge Simplified image.
- You served an OSTree repository of the commit to embed in the image. For example, <http://10.0.2.2:8080/repo>. See [Setting up a web server to install RHEL for Edge image](#) .

Procedure

1. Create the bootable ISO image.

```
# composer-cli compose start-ostree \
  blueprint-name \
  edge-simplified-installer \
  --ref rhel/8/x86_64/edge \
  --url URL-OSTree-repository \
```

Where,

- **blueprint-name** is the RHEL for Edge blueprint name.
- **edge-simplified-installer** is the image type .
- **--ref** is the reference for where your commit is going to be created.
- **--url** is the URL to the OSTree repository of the commit to embed in the image. For example, <http://10.0.2.2:8080/repo/>. You can either start a RHEL for Edge Container or set up a web server. See [Creating a RHEL for Edge Container image for non-network-based deployments](#) and [Setting up a web server to install RHEL for Edge image](#) .
A confirmation that the composer process has been added to the queue appears. It also shows a Universally Unique Identifier (UUID) number for the image created. Use the UUID number to track your build. Also keep the UUID number handy for further tasks.

2. Check the image compose status.

```
# composer-cli compose status
```

The output displays the status in the following format:

```
<UUID> RUNNING date blueprint-name blueprint-version image-type
```



NOTE

The image creation processes can take up to ten minutes to complete.

To interrupt the image creation process, run:

```
# composer-cli compose cancel <UUID>
```

To delete an existing image, run:

```
# composer-cli compose delete <UUID>
```

Additional resources

- [Composing a RHEL for Edge image using RHEL image builder command-line](#)

6.4. DOWNLOADING A SIMPLIFIED RHEL FOR EDGE IMAGE USING THE IMAGE BUILDER COMMAND-LINE INTERFACE

To download a RHEL for Edge image by using RHEL image builder command line interface, ensure that you have met the following prerequisites and then follow the procedure.

Prerequisites

- You have created a RHEL for Edge image.

Procedure

1. Review the RHEL for Edge image status.

```
# composer-cli compose status
```

The output must display the following:

```
$ <UUID> FINISHED date blueprint-name blueprint-version image-type
```

2. Download the image.

```
# composer-cli compose image <UUID>
```

RHEL image builder downloads the image as an **.iso** file at the current directory path where you run the command.

The UUID number and the image size is displayed alongside.


```
$ <UUID>-simplified-installer.iso: size MB
```

As a result, you downloaded a RHEL for Edge Simplified Installer ISO image. You can use it directly as a boot ISO to install a RHEL for Edge system.

6.5. CREATING A BLUEPRINT FOR A SIMPLIFIED IMAGE RHEL USING IMAGE BUILDER GUI

To create a RHEL for Edge Simplified Installer image, you must create a blueprint and ensure that you customize it with:

- A device node location to enable an unattended installation to your device.
- A URL to perform the initial device credential exchange.
- A user or user group.

**NOTE**

You can also add any other customizations that your image requires.

To create a blueprint for a simplified RHEL for Edge image in the RHEL image builder GUI, complete the following steps:

Prerequisites

- You have opened the image builder app from the web console in a browser. See [Accessing the RHEL image builder GUI in the RHEL web console](#).

Procedure

1. Click **Create Blueprint** in the upper-right corner of the RHEL image builder app. A dialog wizard with fields for the blueprint name and description opens.
2. On the **Details** page:
 - a. Enter the name of the blueprint and, optionally, its description. Click **Next**.
3. Optional: On the **Packages** page, complete the following steps:
 - a. In the **Available packages** search, enter the package name and click the **>** button to move it to the **Chosen packages** field. Search and include as many packages as you want. Click **Next**.

**NOTE**

The customizations are all optional unless otherwise specified.

4. Optional: On the **Kernel** page, enter a kernel name and the command-line arguments.
5. Optional: On the **File system** page, select **Use automatic partitioning**. The filesystem customization is not supported for OSTree systems, because OSTree images have their own mount rule, such as read-only. Click **Next**.
6. Optional: On the **Services** page, you can enable or disable services:
 - a. Enter the service names you want to enable or disable, separating them by a comma, by space, or by pressing the **Enter** key. Click **Next**.
7. Optional: On the **Firewall** page, set up your firewall setting:
 - a. Enter the **Ports**, and the firewall services you want to enable or disable.
 - b. Click the **Add zone** button to manage your firewall rules for each zone independently. Click **Next**.
8. On the **Users** page, add a users by following the steps:
 - a. Click **Add user**.
 - b. Enter a **Username**, a **password**, and a **SSH key**. You can also mark the user as a privileged user, by clicking the **Server administrator** checkbox.



NOTE

When you specify the user in the blueprint customization and then create an image from that blueprint, the blueprint creates the user under the **/usr/lib/passwd** directory and the password under the **/usr/etc/shadow** during installation time. You can log in to the device with the username and password you created for the blueprint. After you access the system, you must create users, for example, using the **useradd** command.

Click **Next**.

9. Optional: On the **Groups** page, add groups by completing the following steps:
 - a. Click the **Add groups** button:
 - i. Enter a **Group name** and a **Group ID**. You can add more groups. Click **Next**.
10. Optional: On the **SSH keys** page, add a key:
 - a. Click the **Add key** button.
 - i. Enter the SSH key.
 - ii. Enter a **User**. Click **Next**.
11. Optional: On the **Timezone** page, set your timezone settings:
 - a. On the **Timezone** field, enter the timezone you want to add to your system image. For example, add the following timezone format: "US/Eastern".
If you do not set a timezone, the system uses Universal Time, Coordinated (UTC) as default.
 - b. Enter the **NTP** servers. Click **Next**.
12. Optional: On the **Locale** page, complete the following steps:
 - a. On the **Keyboard** search field, enter the package name you want to add to your system image. For example: ["en_US.UTF-8"].
 - b. On the **Languages** search field, enter the package name you want to add to your system image. For example: "us". Click **Next**.
13. Mandatory: On the **Others** page, complete the following steps:
 - a. In the **Hostname** field, enter the hostname you want to add to your system image. If you do not add a hostname, the operating system determines the hostname.
 - b. Mandatory: In the **Installation Devices** field, enter a valid node for your system image to enable an unattended installation to your device. For example: **dev/sda1**. Click **Next**.
14. Optional: On the **FIDO device onboarding** page, complete the following steps:
 - a. On the **Manufacturing server URL** field, enter the **manufacturing server URL** to perform the initial device credential exchange, for example: "http://10.0.0.2:8080". The FDO customization in the blueprints is optional, and you can build your RHEL for Edge Simplified Installer image with no errors.
 - b. On the **DIUN public key insecure** field, enter the certification public key hash to perform

the initial device credential exchange. This field accepts "true" as value, which means this is an insecure connection to the manufacturing server. For example:

manufacturing_server_url="http://{FDO_SERVER}:8080"

diun_pub_key_insecure="true". You must use only one of these three options: "key insecure", "key hash" and "key root certs".

- c. On the **DIUN public key hash** field, enter the hashed version of your public key. For example:

17BD05952222C421D6F1BB1256E0C925310CED4CE1C4FFD6E5CB968F4B73BF73.

You can get the key hash by generating it based on the certificate of the manufacturing server. To generate the key hash, run the command:

```
# openssl x509 -fingerprint -sha256 -noout -in /etc/fdo/aio/keys/diun_cert.pem | cut -d"=" -f2 | sed 's://g'
```

The **/etc/fdo/aio/keys/diun_cert.pem** is the certificate that is stored in the manufacturing server.

- d. On the **DIUN public key root certs** field, enter the public key root certs. This field accepts the content of the certification file that is stored in the manufacturing server. To get the content of certificate file, run the command:

```
$ cat /etc/fdo/aio/keys/diun_cert.pem.
```

15. Click **Next**.

16. On the **Review** page, review the details about the blueprint. Click **Create**.

The RHEL image builder view opens, listing existing blueprints.

6.6. CREATING A RHEL FOR EDGE SIMPLIFIED INSTALLER IMAGE USING IMAGE BUILDER GUI

To create a RHEL for Edge Simplified image by using RHEL image builder GUI, ensure that you have met the following prerequisites and then follow the procedure.

Prerequisites

- You opened the RHEL image builder app from the web console in a browser.
- You created a blueprint for the RHEL for Edge Simplified image.
- You served an OSTree repository of the commit to embed in the image, for example, **http://10.0.2.2:8080/repo**. See [Setting up a web server to install RHEL for Edge image](#) .
- The FDO manufacturing server is up and running.

Procedure

1. Access mage builder dashboard.
2. On the blueprint table, find the blueprint you want to build an image for.
3. Navigate to the **Images** tab and click **Create Image**. The **Create image** wizard opens.

4. On the **Image output** page, complete the following steps:
 - a. From the **Select a blueprint** list, select the blueprint you created for the RHEL for Edge Simplified image.
 - b. From the **Image output type** list, select **RHEL for Edge Simplified Installer (.iso)**.
 - c. In the **Image Size** field, enter the image size. Minimum image size required for Simplified Installer image is:
5. Click **Next**.
6. In the **OSTree settings** page, complete the following steps:
 - a. In the **Repository URL** field, enter the repository URL to where the parent OSTree commit will be pulled from.
 - b. In the **Ref** field, enter the **ref** branch name path. If you do not enter a **ref**, the default **ref** for the distro is used.
7. On the **Review** page, review the image customization and click **Create**.

The image build starts and takes up to 20 minutes to complete. To stop the building, click **Stop build**.

6.7. DOWNLOADING A SIMPLIFIED RHEL FOR EDGE IMAGE USING THE IMAGE BUILDER GUI


To download a RHEL for Edge image by using RHEL image builder GUI, ensure that you have met the following prerequisites and then follow the procedure.

Prerequisites

- You have successfully created a RHEL for Edge image. See link.

Procedure

1. Access the RHEL image builder dashboard. The blueprint list dashboard opens.
2. In the blueprint table, find the blueprint you built your RHEL for Edge Simplified Installer image for.
3. Navigate to the **Images** tab.
4. Choose one of the options:
 - Download the image.
 - Download the logs of the image to inspect the elements and verify if any issue is found.



NOTE

You can use the RHEL for Edge Simplified Installer ISO image that you downloaded directly as a boot ISO to install a RHEL for Edge system.

6.8. SETTING UP AN UEFI HTTP BOOT SERVER

To set up an **UEFI HTTP Boot** server, so that you can start to provision a RHEL for Edge Virtual Machine over network by connecting to this UEFI HTTP Boot server, follow the steps:

Prerequisites

- You have created the ISO simplified installer image.
- An http server that serves the ISO content.

Procedure

1. Mount the ISO image to the directory of your choice:

```
# mkdir /mnt/rhel8-install/
# mount -o loop,ro -t iso9660 /path_directory/installer.iso /mnt/rhel8-install/
```

Replace **/path_directory/installer.iso** with the path to the RHEL for Edge bootable ISO image.

2. Copy the files from the mounted image to the HTTP server root. This command creates the **/var/www/html/rhel8-install/** directory with the contents of the image.

```
# mkdir /var/www/html/httpboot/
# cp -R /mnt/rhel8-install/* /var/www/html/httpboot/
# chmod -R +r /var/www/html/httpboot/*
```



NOTE

Some copying methods can skip the **.treeinfo** file which is required for a valid installation source. Running the **cp** command for whole directories as shown in this procedure will copy **.treeinfo** correctly.

3. Update the **/var/www/html/EFI/BOOT/grub.cfg** file, by replacing:
 - a. **coreos.inst.install_dev=/dev/sda** with **coreos.inst.install_dev=/dev/vda**
 - b. **linux /images/pxeboot/vmlinuz** with **linuxefi /images/pxeboot/vmlinuz**
 - c. **initrd /images/pxeboot/initrd.img** with **initrdefi /images/pxeboot/initrd.img**
 - d. **coreos.inst.image_file=/run/media/iso/disk.img.xz** with **coreos.inst.image_url=http://{IP-ADDRESS}/disk.img.xz**
The *IP-ADDRESS* is the ip address of this machine, which will serve as a http boot server.
4. Start the httpd service:

```
# systemctl start httpd.service
```

As a result, after you set up an **UEFI HTTP Boot** server, you can install your RHEL for Edge devices by using **UEFI HTTP** boot.

6.9. DEPLOYING THE SIMPLIFIED ISO IMAGE IN A VIRTUAL MACHINE

Deploy the RHEL for Edge ISO image you generated by creating a RHEL for Edge Simplified image by using any the following installation sources:

- **UEFI HTTP Boot**
- **virt-install**

This example shows how to create a **virt-install** installation source from your ISO image for a **network-based** installation .

Prerequisites

- You have created an ISO image.
- You set up a network configuration to support UEFI HTTP boot.

Procedure

1. Set up a network configuration to support **UEFI HTTP** boot. See [Setting up UEFI HTTP boot with libvirt](#).
2. Use the **virt-install** command to create a RHEL for Edge Virtual Machine from the UEFI HTTP Boot.

```
# virt-install \
  --name edge-install-image \
  --disk path=" ",format=qcow2 \
  --ram 3072 \
  --memory 4096 \
  --vcpus 2 \
  --network network=integration,mac=mac_address \
  --os-type linux \
  --os-variant rhel8 \
  --cdrom "/var/lib/libvirt/images/"ISO_FILENAME" \
  --boot
  uefi,loader_ro=yes,loader_type=pflash,nvram_template=/usr/share/edk2/ovmf/OVMF_VARS.fd,
  loader_secure=no \
  --virt-type kvm \
  --graphics none \
  --wait=-1 \
  --noreboot
```

After you run the command, the Virtual Machine installation starts.

Verification

- Log in to the created Virtual Machine.

6.10. DEPLOYING THE SIMPLIFIED ISO IMAGE FROM A USB FLASH DRIVE

Deploy the RHEL for Edge ISO image you generated by creating a RHEL for Edge Simplified image by using an **USB installation**.


This example shows how to create a **USB installation** source from your ISO image.

Prerequisites

- You have created a simplified installer image, which is an ISO image.
- You have a 8 GB USB flash drive.

Procedure

1. Copy the ISO image file to a USB flash drive.
2. Connect the USB flash drive to the port of the computer you want to boot.
3. Boot the ISO image from the USB flash drive. The boot menu shows you the following options:

A vertical bar representing a cursor is positioned to the left of two menu items. The first item is "Install Red Hat Enterprise Linux 8" and the second item is "Test this media & install Red Hat Enterprise Linux 8".

```
Install Red Hat Enterprise Linux 8
Test this media & install Red Hat Enterprise Linux 8
```

4. Choose Install Red Hat Enterprise Linux 8. This starts the system installation.

Additional resources

- [Booting the installation.](#)

CHAPTER 7. AUTOMATICALLY PROVISIONING AND ONBOARDING RHEL FOR EDGE DEVICES WITH FDO

You can build a RHEL for Edge Simplified Installer image, and provision it to a RHEL for Edge image. The FIDO Device Onboarding (FDO) process automatically provisions and onboards your Edge devices, and exchanges data with other devices and systems connected on the networks.



IMPORTANT

Red Hat provides the **FDO** process as a Technology Preview feature and should run on secure networks. Technology Preview features are not supported with Red Hat production service level agreements (SLAs) and might not be functionally complete. These features provide early access to upcoming product features, enabling customers to test functionality and provide feedback during the development process. See [Technology Preview Features Support Scope](#) on the Red Hat Customer Portal for information about the support scope for Technology Preview features.

7.1. THE FIDO DEVICE ONBOARDING (FDO) PROCESS

The FIDO Device Onboarding (FDO) is the process that:

- Provisions and onboards a device.
- Automatically configures credentials for this device. The FDO process is an automatic onboarding mechanism that is triggered by the installation of a new device.
- Enables this device to securely connect and interact on the network.

With FIDO Device Onboarding (FDO), you can perform a secure device onboarding by adding new devices into your IoT architecture. This includes the specified device configuration that needs to be trusted and integrated with the rest of the running systems. The FDO process is an automatic onboarding mechanism that is triggered by the installation of a new device.

The FDO protocol performs the following tasks:

- Solves the trust and chain of ownership along with the automation needed to securely onboard a device at scale.
- Performs device initialization at the manufacturing stage and late device binding for its actual use. This means that actual binding of the device to a management system happens on the first boot of the device without requiring manual configuration on the device.
- Supports automated secure devices onboarding, that is, zero touch installation and onboarding that does not need any specialized person at the edge location. After the device is onboarded, the management platform can connect to it and apply patches, updates, and rollbacks.

With FDO, you can benefit from the following:

- FDO is a secure and simple way to enroll a device to a management platform. Instead of embedding a Kickstart configuration to the image, FDO applies the device credentials during the device first boot directly to the ISO image.
- FDO solves the issue of late binding to a device, enabling any sensitive data to be shared over a secure FDO channel.

- FDO cryptographically identifies the system identity and ownership before enrolling and passing the configuration and other secrets to the system. That enables non-technical users to power-on the system.

To build a RHEL for Edge Simplified Installer image and automatically onboard it, provide an existing OSTree commit. The resulting simplified image contains a raw image that has the OSTree commit deployed. After you boot the Simplified installer ISO image, it provisions a RHEL for Edge system that you can use on a hard disk or as a boot image in a virtual machine.

The RHEL for Edge Simplified Installer image is optimized for unattended installation to a device and supports both network-based deployment and non-network-based deployments. However, for network-based deployment, it supports only UEFI HTTP boot.

The FDO protocol is based on the following servers:

Manufacturing server

1. Generates the device credentials.
2. Creates an Ownership voucher that is used to set the ownership of the device, later in the process.
3. Binds the device to a specific management platform.

Owner management system

1. Receives the Ownership voucher from the Manufacturing server and becomes the owner of the associated device.
2. Later in the process, it creates a secure channel between the device and the Owner onboarding server after the device authentication.
3. Uses the secure channel to send the required information, such as files and scripts for the onboarding automation to the device.

Service-info API server

Based on Service-info API server's configuration and modules available on the client, it performs the final steps of onboarding on target client devices, such as copying SSH keys and files, executing commands, creating users, encrypting disks and so on

Rendezvous server

1. Gets the Ownership voucher from the Owner management system and makes a mapping of the device UUID to the Owner server IP. Then, the Rendezvous server matches the device UUID with a target platform and informs the device about which Owner onboarding server endpoint this device must use.
2. During the first boot, the Rendezvous server will be the contact point for the device and it will direct the device to the owner, so that the device and the owner can establish a secure channel.


Device client

This is installed on the device. The Device client performs the following actions:

1. Starts the queries to the multiple servers where the onboarding automation will be executed.
2. Uses TCP/IP protocols to communicate with the servers.

The following diagram represents the FIDO device onboarding workflow:

Figure 7.1. Deploying RHEL for Edge in non-network environment



356_RHEL_0823

At the **Device Initialization**, the device contacts the Manufacturing server to get the FDO credentials, a set of certificates and keys to be installed on the operating system with the Rendezvous server endpoint (URL). It also gets the Ownership Voucher, that is maintained separately in case you need to change the owner assignment.

1. The Device contacts the Manufacturing server
2. The Manufacturing server generates an Ownership Voucher and the Device Credentials for the Device.
3. The Ownership Voucher is transferred to the Owner onboarding server.

At the **On-site onboarding**, the Device gets the Rendezvous server endpoint (URL) from its device credentials and contacts Rendezvous server endpoint to start the onboarding process, which will redirect it to the Owner management system, that is formed by the Owner onboarding server and the Service Info API server.

4. The Owner onboarding server transfers the Ownership Voucher to the Rendezvous server, which makes a mapping of the Ownership Voucher to the Owner.
5. The device client reads device credentials.
6. The device client connects to the network.
7. After connecting to the network, the Device client contacts the Rendezvous server.
8. The Rendezvous server sends the owner endpoint URL to the Device Client, and registers the device.
9. The Device client connects to the Owner onboarding server shared by the Rendezvous server.
10. The Device proves that it is the correct device by signing a statement with a device key.
11. The Owner onboarding server proves itself correct by signing a statement with the last key of the Owner Voucher.
12. The Owner onboarding server transfers the information of the Device to the Service Info API server.
13. The Service info API server sends the configuration for the Device.
14. The Device is onboarded.

7.2. AUTOMATICALLY PROVISIONING AND ONBOARDING RHEL FOR EDGE DEVICES

To build a RHEL for Edge Simplified Installer image and automatically onboard it, provide an existing OSTree commit. The resulting simplified image contains a raw image that has the OSTree commit deployed. After you boot the Simplified installer ISO image, it provisions a RHEL for Edge system that you can use on a hard disk or as a boot image in a virtual machine.

The RHEL for Edge Simplified Installer image is optimized for unattended installation to a device and supports both network-based deployment and non-network-based deployments. However, for network-based deployment, it supports only UEFI HTTP boot.

Automatically provisioning and onboarding a RHEL for Edge device involves the following high-level steps:

1. Install and register a RHEL system
2. Install RHEL image builder
3. By using RHEL image builder, create a blueprint with customizations for RHEL for a **rhel-edge-container** image type.

```
name = "rhel-edge-container"
description = "Minimal RHEL for Edge Container blueprint"
version = "0.0.1"
```

4. Import the RHEL for Edge Container blueprint in RHEL image builder
5. Create a RHEL for Edge Container image
6. Use the RHEL for Edge Container image to serve the OSTree commit, which will be later used when building the RHEL for Edge Simplified Installer image type
7. Create a blueprint for and **edge-simplified-installer** image type with customizations for storage device path and FDO customizations

```
name = "rhel-edge-simplified-installer-with-fdo"
description = "Minimal RHEL for Edge Simplified Installer with FDO blueprint"
version = "0.0.1"
packages = []
modules = []
groups = []
distro = ""

[customizations]
installation_device = "/dev/vda"

[customizations.fdo]
manufacturing_server_url = "http://10.0.0.2:8080"
diun_pub_key_insecure = "true"
```


8. Build a simplified installer RHEL for Edge image
9. Download the RHEL for Edge simplified installer image
10. At this point, the FDO server infrastructure should be up and running, and the specific onboarding details handled by the **service-info API** server, that is part of the owner's infrastructure, are configured
11. Install the simplified installer ISO image to a device. The FDO client runs on the Simplified Installer ISO and the UEFI directory structure makes the image bootable.
12. The network configuration enables the device to reach out to the manufacturing server to perform the initial device credential exchange.
13. After the system reaches the endpoint, the device credentials are created for the device.
14. The device uses the device credentials to reach the Rendezvous server, where it checks the cryptographic credentials based on the vouchers that the Rendezvous server has, and then the Rendezvous server redirects the device to the Owner server.
15. The device contacts the Owner server. They establish a mutual trust and the final steps of onboarding happen based on the configuration of the Service-info API server. For example, it installs the SSH keys in the device, transfer the files, create the users, run the commands, encrypt the filesystem, and so on.

Additional resources

- [FDO automatic onboarding technologies](#)

7.3. GENERATING KEY AND CERTIFICATES

To run the FIDO Device Onboarding (FDO) infrastructure, you need to generate keys and certificates. FDO generates these keys and certificates to configure the manufacturing server. FDO automatically generates the certificates and **.yaml** configuration files when you install the services, and re-creating them is optional. After you install and start the services, it runs with the default settings.



IMPORTANT

Red Hat provides the **fdo-admin-tool** tool as a Technology Preview feature and should run on secure networks. Technology Preview features are not supported with Red Hat production service level agreements (SLAs) and might not be functionally complete. These features provide early access to upcoming product features, enabling customers to test functionality and provide feedback during the development process. See [Technology Preview Features Support Scope](#) on the Red Hat Customer Portal for information about the support scope for Technology Preview features.

Prerequisites

- You installed the **fdo-admin-cli** RPM package

Procedure

1. Generate the keys and certificates in the **/etc/fdo** directory:

```
$ for i in "diun" "manufacturer" "device-ca" "owner"; do fdo-admin-tool generate-key-and-cert
$i; done
$ ls keys
device_ca_cert.pem device_ca_key.der diun_cert.pem diun_key.der manufacturer_cert.pem
manufacturer_key.der owner_cert.pem owner_key.der
```

2. Check the key and certificates that were created in the **/etc/fdo/keys** directory:

```
$ tree keys
```

You can see the following output:

```
– device_ca_cert.pem
– device_ca_key.der
– diun_cert.pem
– diun_key.dre
– manufacturer_cert.pem
– manufacturer_key.der
– owner_cert.pem
– owner_key.pem
```


Additional resources

- See the **fdo-admin-tool generate-key-and-cert --help** manual page

7.4. INSTALLING AND RUNNING THE MANUFACTURING SERVER

The **fdo-manufacturing-server** RPM package enables you to run the Manufacturing Server component of the FDO protocol. It also stores other components, such as the owner vouchers, the manufacturer keys, and information about the manufacturing sessions. During the device installation, the

Manufacturing server generates the device credentials for the specific device, including its GUID, rendezvous information and other metadata. Later on in the process, the device uses this rendezvous information to contact the Rendezvous server.



IMPORTANT

Red Hat provides the **fdo-manufacturing-server** tool as a Technology Preview feature and should run on secure networks because Technology Preview features are not supported with Red Hat production service level agreements (SLAs) and might not be functionally complete. These features provide early access to upcoming product features, enabling customers to test functionality and provide feedback during the development process. See [Technology Preview Features Support Scope](#) on the Red Hat Customer Portal for information about the support scope for Technology Preview features.

To install the **manufacturing server** RPM package, complete the following steps:

Procedure

1. Install the **fdo-admin-cli** package:

```
# yum install -y fdo-admin-cli
```

2. Check if the **fdo-manufacturing-server** RPM package is installed:

```
$ rpm -qa | grep fdo-manufacturing-server --refresh
```

3. Check if the files were correctly installed:

```
$ ls /usr/share/doc/fdo
```

You can see the following output:

```
Output:
manufacturing-server.yml
owner-onboarding-server.yml
rendezvous-info.yml
rendezvous-server.yml
serviceinfo-api-server.yml
```

4. Optional: Check the content of each file, for example:

```
$ cat /usr/share/doc/fdo/manufacturing-server.yml
```

5. Configure the Manufacturing server. You must provide the following information:

- The Manufacturing server URL
- The IP address or DNS name for the Rendezvous server
- The path to the keys and certificates you generated. See [Generating key and certificates](#). You can find an example of a Manufacturing server configuration file in the **/usr/share/doc/fdo/manufacturing-server.yml** directory. The following is a **manufacturing**

server.yml example that is created and saved in the **/etc/fdo** directory. It contains paths to the directories, certificates, keys that you created, the Rendezvous server IP address and the default port.

```
session_store_driver:
  Directory:
    path: /etc/fdo/stores/manufacturing_sessions/
ownership_voucher_store_driver:
  Directory:
    path: /etc/fdo/stores/owner_vouchers
public_key_store_driver:
  Directory:
    path: /etc/fdo/stores/manufacturer_keys
bind: "0.0.0.0:8080"
protocols:
  plain_di: false
diun:
  mfg_string_type: SerialNumber
  key_type: SECP384R1
  allowed_key_storage_types:
    - Tpm
    - FileSystem
  key_path: /etc/fdo/keys/diun_key.der
  cert_path: /etc/fdo/keys/diun_cert.pem
rendezvous_info:
  - deviceport: 8082
    ip_address: 192.168.122.99
    ownerport: 8082
    protocol: http
manufacturing:
  manufacturer_cert_path: /etc/fdo/keys/manufacturer_cert.pem
  device_cert_ca_private_key: /etc/fdo/keys/device_ca_key.der
  device_cert_ca_chain: /etc/fdo/keys/device_ca_cert.pem
  owner_cert_path: /etc/fdo/keys/owner_cert.pem
  manufacturer_private_key: /etc/fdo/keys/manufacturer_key.der
```

6. Start the Manufacturing server.

- a. Check if the systemd unit file are in the server:

```
# systemctl list-unit-files | grep fdo | grep manufacturing
fdo-manufacturing-server.service disabled disabled
```

- b. Enable and start the manufacturing server.

```
# systemctl enable --now fdo-manufacturing-server.service
```

- c. Open the default ports in your firewall:

```
# firewall-cmd --add-port=8080/tcp --permanent
# systemctl restart firewalld
```

- d. Ensure that the service is listening on the port 8080:

```
# ss -ltn
```

7. Install RHEL for Edge onto your system using the simplified installer. See [Building simplified installer images to provision a RHEL for Edge image](#).

Additional resources

- The [manufacturing-server.yml](#) example
- [FDO automatic onboarding terminology](#)

7.5. INSTALLING, CONFIGURING, AND RUNNING THE RENDEZVOUS SERVER

Install the **fdo-rendezvous-server** RPM package to enable the systems to receive the voucher generated by the Manufacturing server during the first device boot. The Rendezvous server then matches the device UUID with the target platform or cloud and informs the device about which Owner server endpoint the device must use.

Prerequisites

- You created a **manufacturer_cert.pem** certificate. See [Generating key and certificates](#).
- You copied the **manufacturer_cert.pem** certificate to the **/etc/fdo/keys** directory in the Rendezvous server.

Procedure

1. Install the **fdo-rendezvous-server** RPM packages:

```
# yum install -y fdo-rendezvous-server
```

2. Create the **rendezvous-server.yml** configuration file, including the path to the manufacturer certificate. You can find an example in **/usr/share/doc/fdo/rendezvous-server.yml**. The following example shows a configuration file that is saved in **/etc/fdo/rendezvous-server.yml**.

```
storage_driver:
  Directory:
    path: /etc/fdo/stores/rendezvous_registered
session_store_driver:
  Directory:
    path: /etc/fdo/stores/rendezvous_sessions
trusted_manufacturer_keys_path: /etc/fdo/keys/manufacturer_cert.pem
max_wait_seconds: ~
bind: "0.0.0.0:8082"
```

3. Check the Rendezvous server service status:

```
# systemctl list-unit-files | grep fdo | grep rende
fdo-rendezvous-server.service disabled disabled
```

- a. If the service is stopped and disabled, enable and start it:


```
# systemctl enable --now fdo-rendezvous-server.service
```

4. Check that the server is listening on the default configured port 8082:

```
# ss -ltn
```

5. Open the port if you have a firewall configured on this server:

```
# firewall-cmd --add-port=8082/tcp --permanent
# systemctl restart firewalld
```

7.6. INSTALLING, CONFIGURING, AND RUNNING THE OWNER SERVER

Install the **fdo-owner-cli** and **fdo-owner-onboarding-server** RPM package to enable the systems to receive the voucher generated by the Manufacturing server during the first device boot. The Rendezvous server then matches the device UUID with the target platform or cloud and informs the device about which Owner server endpoint the device must use.

Prerequisites

- The device where the server will be deployed has a Trusted Platform Module (TPM) device to encrypt the disk. If not, you will get an error when booting the RHEL for Edge device.
- You created the **device_ca_cert.pem**, **owner_key.der**, and **owner_cert.pem** with keys and certificates and copied them into the **/etc/fdo/keys** directory.

Procedure

1. Install the required RPMs in this server:

```
# dnf install -y fdo-owner-cli fdo-owner-onboarding-server
```

2. Prepare the **owner-onboarding-server.yml** configuration file and save it to the **/etc/fdo/** directory. Include the path to the certificates you already copied and information about where to publish the Owner server service in this file.

The following is an example available in **/usr/share/doc/fdo/owner-onboarding-server.yml**. You can find references to the Service Info API, such as the URL or the authentication token.

```
---
ownership_voucher_store_driver:
  Directory:
    path: /etc/fdo/stores/owner_vouchers
session_store_driver:
  Directory:
    path: /etc/fdo/stores/owner_onboarding_sessions
trusted_device_keys_path: /etc/fdo/keys/device_ca_cert.pem
owner_private_key_path: /etc/fdo/keys/owner_key.der
owner_public_key_path: /etc/fdo/keys/owner_cert.pem
bind: "0.0.0.0:8081"
service_info_api_url: "http://localhost:8083/device_info"
service_info_api_authentication:
  BearerToken:
    token: Kpt5P/5fIBkaiNSvDYS3cEdBQXJn2Zv9n1D50431/lo=
```

```
owner_addresses:
- transport: http
  addresses:
    - ip_address: 192.168.122.149
```

3. Create and configure the Service Info API.

- a. Add the automated information for onboarding, such as user creation, files to be copied or created, commands to be executed, disk to be encrypted, and so on. Use the Service Info API configuration file example in **/usr/share/doc/fdo/serviceinfo-api-server.yml** as a template to create the configuration file under **/etc/fdo/**.

```
---
service_info:
  initial_user:
    username: admin
    sshkeys:
      - "ssh-rsa AAAA...."
  files:
    - path: /root/resolv.conf
      source_path: /etc/resolv.conf
  commands:
    - command: touch
      args:
        - /root/test
      return_stdout: true
      return_stderr: true
  diskencryption_clevis:
    - disk_label: /dev/vda4
      binding:
        pin: tpm2
        config: "{}"
      reencrypt: true
  additional_serviceinfo: ~
  bind: "0.0.0.0:8083"
  device_specific_store_driver:
    Directory:
      path: /etc/fdo/stores/serviceinfo_api_devices
  service_info_auth_token: Kpt5P/5fIBkaiNSvDYS3cEdBQXJn2Zv9n1D50431/lo=
  admin_auth_token: zJNoErq7aa0RusJ1w0tkTjdITdMCWYkndzVv7F0V42Q=
```

4. Check the status of the systemd units:

```
# systemctl list-unit-files | grep fdo
fdo-owner-onboarding-server.service    disabled    disabled
fdo-serviceinfo-api-server.service      disabled    disabled
```

- a. If the service is stopped and disabled, enable and start it:

```
# systemctl enable --now fdo-owner-onboarding-server.service
# systemctl enable --now fdo-serviceinfo-api-server.service
```

**NOTE**

You must restart the **systemd** services every time you change the configuration files.

5. Check that the server is listening on the default configured port 8083:

```
# ss -ltn
```

6. Open the port if you have a firewall configured on this server:

```
# firewall-cmd --add-port=8081/tcp --permanent
# firewall-cmd --add-port=8083/tcp --permanent
# systemctl restart firewalld
```

7.7. AUTOMATICALLY ONBOARDING A RHEL FOR EDGE DEVICE BY USING FDO AUTHENTICATION

To prepare your device to automatically onboard a RHEL for Edge device and provision it as part of the installation process, complete the following steps:

Prerequisites

- You built an OSTree commit for RHEL for Edge and used that to generate an **edge-simplified-installer** artifact.
- Your device is assembled.
- You installed the **fdo-manufacturing-server** RPM package. See [Installing the manufacturing server package](#).

Procedure

1. Start the installation process by booting the RHEL for Edge simplified installer image on your device. You can install it from a CD-ROM or from a USB flash drive, for example.
2. Verify through the terminal that the device has reached the manufacturing service to perform the initial device credential exchange and has produced an ownership voucher.
You can find the ownership voucher at the storage location configured by the **ownership_voucher_store_driver** parameter at the **manufacturing-sever.yml** file.

The directory should have an **ownership_voucher** file with a name in the GUID format which indicates that the correct device credentials were added to the device.

The onboarding server uses the device credential to authenticate against the onboarding server. It then passes the configuration to the device. After the device receives the configuration from the onboarding server, it receives an SSH key and installs the operating system on the device. Finally, the system automatically reboots, encrypts it with a strong key stored at TPM.

Verification

After the device automatically reboots, you can log in to the device with the credentials that you created as part of the FDO process.

1. Log in to the device by providing the username and password you created in the Service Info API.

Additional resources

- [Deploying the Simplified ISO image from a USB flash drive](#)

CHAPTER 8. DEPLOYING A RHEL FOR EDGE IMAGE IN A NETWORK-BASED ENVIRONMENT

You can deploy a RHEL for Edge image using the RHEL installer graphical user interface or a Kickstart file. The overall process for deploying a RHEL for Edge image depends on whether your deployment environment is network-based or non-network-based.



NOTE

To deploy the images on bare metal, use a Kickstart file.

Network-based deployments

Deploying a RHEL for Edge image in a network-based environment involves the following high-level steps:

- a. Extract the image file contents.
- b. Set up a web server
- c. Install the image

8.1. EXTRACTING THE RHEL FOR EDGE IMAGE COMMIT

After you download the commit, extract the **.tar** file and note the ref name and the commit ID.

The downloaded commit file consists of a **.tar** file with an **OSTree** repository. The **OSTree** repository has a commit and a **compose.json** file.

The **compose.json** file has information metadata about the commit with information such as the "Ref", the reference ID and the commit ID. The commit ID has the RPM packages.

To extract the package contents, perform the following the steps:

Prerequisites

- Create a Kickstart file or use an existing one.

Procedure

1. Extract the downloaded image **.tar** file:

```
# tar xvf <UUID>-commit.tar
```

2. Go to the directory where you have extracted the **.tar** file.
It has a **compose.json** file and an **OSTree** directory. The **compose.json** file has the commit number and the **OSTree** directory has the RPM packages.
3. Open the **compose.json** file and note the commit ID number. You need this number handy when you proceed to set up a web server.
If you have the **jq** JSON processor installed, you can also retrieve the commit ID by using the **jq** tool:

```
# jq '["ostree-commit"]' < compose.json
```

4. List the RPM packages in the commit.

```
# rpm-ostree db list rhel/8/x86_64/edge --repo=repo
```

5. Use a Kickstart file to run the RHEL installer. Optionally, you can use any existing file or can create one by using the Kickstart Generator tool.
In the Kickstart file, ensure that you include the details about how to provision the file system, create a user, and how to fetch and deploy the RHEL for Edge image. The RHEL installer uses this information during the installation process.

The following is a Kickstart file example:

```
lang en_US.UTF-8
keyboard us
timezone Etc/UTC --isUtc
text
zerombr
clearpart --all --initlabel
autopart
reboot
user --name=core --group=wheel
sshkey --username=core "ssh-rsa AAAA3Nza...."
rootpw --lock
network --bootproto=dhcp

ostreesetup --nogpg --osname=rhel --remote=edge --url=https://mirror.example.com/repo/ --
ref=rhel/8/x86_64/edge
```

The OSTree-based installation uses the **ostreesetup** command to set up the configuration. It fetches the OSTree commit, by using the following flags:

- **--nogpg** - Disable GNU Privacy Guard (GPG) key verification.
- **--osname** - Management root for the operating system installation.
- **--remote** - Management root for the operating system installation
- **--url** - URL of the repository to install from.
- **--ref** - Name of the branch from the repository that the installation uses.
- **--url=http://mirror.example.com/repo/** - is the address of the host system where you extracted the edge commit and served it over **nginx**. You can use the address to reach the host system from the guest computer.
For example, if you extract the commit image in the **/var/www/html** directory and serve the commit over **nginx** on a computer whose hostname is **www.example.com**, the value of the **--url** parameter is **http://www.example.com/repo**.



NOTE

Use the http protocol to start a service to serve the commit, because https is not enabled on the Apache HTTP Server.

Additional resources

- [Downloading a RHEL for Edge image](#)
- [Creating Kickstart files](#)

8.2. SETTING UP A WEB SERVER TO INSTALL RHEL FOR EDGE IMAGES

After you have extracted the RHEL for Edge image contents, set up a web server to provide the image commit details to the RHEL installer by using HTTP.

The following example provides the steps to set up a web server by using a container.

Prerequisites

- You have installed Podman on your system. See [How do I install Podman in RHEL](#)

Procedure

1. Create the **nginx** configuration file with the following instructions:

```
events {
}

http {
    server{
        listen 8080;
        root /usr/share/nginx/html;
    }
}

pid /run/nginx.pid;
daemon off;
```

2. Create a Dockerfile with the following instructions:

```
FROM registry.access.redhat.com/ubi8/ubi
RUN yum -y install nginx && yum clean all
COPY kickstart.ks /usr/share/nginx/html/
COPY repo /usr/share/nginx/html/
COPY nginx /etc/nginx.conf
EXPOSE 8080
CMD ["/usr/sbin/nginx", "-c", "/etc/nginx.conf"]
ARG commit
ADD ${commit} /usr/share/nginx/html/
```

Where,

- **kickstart.ks** is the name of the Kickstart file from the RHEL for Edge image. The Kickstart file includes directive information. To help you manage the images later, it is advisable to include the checks and settings for greenboot checks. For that, you can update the Kickstart file to include the following settings:

■

```

lang en_US.UTF-8
keyboard us
timezone Etc/UTC --isUtc
text
zerombr
clearpart --all --initlabel
autopart
reboot
user --name=core --group=wheel
sshkey --username=core "ssh-rsa AAAA3Nza...."

ostreesetup --nogpg --osname=rhel --remote=edge
--url=https://mirror.example.com/repo/
--ref=rhel/8/x86_64/edge

%post
cat << EOF > /etc/greenboot/check/required.d/check-dns.sh
#!/bin/bash

DNS_SERVER=$(grep nameserver /etc/resolv.conf | cut -f2 -d" ")
COUNT=0

# check DNS server is available
ping -c1 $DNS_SERVER
while [ $? != '0' ] && [ $COUNT -lt 10 ]; do

    COUNT++
    echo "Checking for DNS: Attempt $COUNT ."
    sleep 10
    ping -c 1 $DNS_SERVER
done
EOF
%end

```

Any HTTP service can host the OSTree repository, and the example, which uses a container, is just an option for how to do this. The Dockerfile performs the following tasks:

- a. Uses the latest Universal Base Image (UBI)
 - b. Installs the web server (nginx)
 - c. Adds the Kickstart file to the server
 - d. Adds the RHEL for Edge image commit to the server
3. Build a Docker container

```
# podman build -t name-of-container-image --build-arg commit=uuid-commit.tar .
```

4. Run the container

```
# podman run --rm -d -p port:8080 localhost/name-of-container-image
```


As a result, the server is set up and ready to start the RHEL Installer by using the **commit.tar** repository and the Kickstart file.

8.3. PERFORMING AN ATTENDED INSTALLATION TO AN EDGE DEVICE BY USING KICKSTART

For an attended installation in a network-based environment, you can install the RHEL for Edge image to a device by using the RHEL Installer ISO, a Kickstart file, and a web server. The web server serves the RHEL for Edge Commit and the Kickstart file to boot the [RHEL Installer ISO](#) image.

Prerequisites

- You have made the RHEL for Edge Commit available by running a web server. See [Setting up a web server to install RHEL for Edge images](#).
- You have created a **.qcow2** disk image to be used as the target of the attended installation. See [Creating a virtual disk image by using qemu-img](#).


Procedure

1. Run the RHEL Anaconda Installer by using the **libvirt virt-install** utility to create a virtual machine (VM) with a RHEL operating system. Use the **.qcow2** disk image as the target disk in the attended installation:

```
virt-install \
--name rhel-edge-test-1 \
--memory 2048 \
--vcpus 2 \
--disk path=prepared_disk_image.qcow2,format=qcow2,size=8 \
--os-variant rhel8 \
--cdrom /home/username/Downloads/rhel-8-x86_64-boot.iso
```

2. On the installation screen:

Figure 8.1. Red Hat Enterprise Linux boot menu



- a. Press the **e** key to add an additional kernel parameter:

```
inst.ks=http://edge_device_ip:port/kickstart.ks
```

The kernel parameter specifies that you want to install RHEL by using the Kickstart file and not the RHEL image contained in the RHEL Installer.

- b. After adding the kernel parameters, press **Ctrl+X** to boot the RHEL installation by using the Kickstart file.

The RHEL Installer starts, fetches the Kickstart file from the server (HTTP) endpoint and executes the commands, including the command to install the RHEL for Edge image commit from the HTTP endpoint. After the installation completes, the RHEL Installer prompts you for login details.

Verification

1. On the Login screen, enter your user account credentials and click **Enter**.
2. Verify whether the RHEL for Edge image is successfully installed.

```
$ rpm-ostree status
```

The command output provides the image commit ID and shows that the installation is successful.

Following is a sample output:

```

State: idle
Deployments:
* ostree://edge:rhel/8/x86_64/edge
  Timestamp: 2020-09-18T20:06:54Z
  Commit: 836e637095554e0b634a0a48ea05c75280519dd6576a392635e6fa7d4d5e96

```

Additional resources

- [How to embed a Kickstart file into an ISO image](#) .
- [Booting the installation.](#)

8.4. PERFORMING AN UNATTENDED INSTALLATION TO AN EDGE DEVICE BY USING KICKSTART

For an unattended installation in a network-based environment, you can install the RHEL for Edge image to an Edge device by using a Kickstart file and a web server. The web server serves the RHEL for Edge Commit and the Kickstart file, and both artifacts are used to start the [RHEL Installer ISO](#) image.

Prerequisites

- You have the **qemu-img** utility installed on your host system.
- You have created a **.qcow2** disk image to install the commit you created. See [Creating a system image with RHEL image builder in the CLI](#).
- You have a running web server. See [Creating a RHEL for Edge Container image for non-network-based deployments](#).

Procedure

1. Run a RHEL for Edge Container image to start a web server. The server fetches the commit in the RHEL for Edge Container image and becomes available and running.
2. Run the RHEL Anaconda Installer, passing the customized **.qcow2** disk image, by using **libvirt virt-install**:


```

virt-install \
  --name rhel-edge-test-1 \
  --memory 2048 \
  --vcpus 2 \
  --disk path=prepared_disk_image.qcow2,format=qcow2,size=8 \
  --os-variant rhel8 \
  --cdrom /home/username/Downloads/rhel-8-x86_64-boot.iso

```

3. On the installation screen:

Figure 8.2. Red Hat Enterprise Linux boot menu



- a. Press the **TAB** key and add the Kickstart kernel argument:

```
inst.ks=http://web-server_device_ip:port/kickstart.ks
```

The kernel parameter specifies that you want to install RHEL by using the Kickstart file and not the RHEL image contained in the RHEL Installer.

- b. After adding the kernel parameters, press **Ctrl+X** to boot the RHEL installation by using the Kickstart file.
The RHEL Installer starts, fetches the Kickstart file from the server (HTTP) endpoint, and executes the commands, including the command to install the RHEL for Edge image commit from the HTTP endpoint. After the installation completes, the RHEL Installer prompts you for login details.

Verification

1. On the Login screen, enter your user account credentials and click **Enter**.
2. Verify whether the RHEL for Edge image is successfully installed.

```
$ rpm-ostree status
```

The command output provides the image commit ID and shows that the installation is successful.

The following is a sample output:

State: idle

Deployments:

* ostree://edge:rhel/8/x86_64/edge

Timestamp: 2020-09-18T20:06:54Z

Commit: 836e637095554e0b634a0a48ea05c75280519dd6576a392635e6fa7d4d5e96

Additional resources

- [How to embed a Kickstart file into an ISO image](#) .
- [Booting the installation.](#)

CHAPTER 9. DEPLOYING A RHEL FOR EDGE IMAGE IN A NON-NETWORK-BASED ENVIRONMENT

The RHEL for Edge Container (**.tar**) in combination with the RHEL for Edge Installer (**.iso**) image type result in a ISO image. The ISO image can be used in disconnected environments during the image deployment to a device. However, network access might require network access to build the different artifacts.

Deploying a RHEL for Edge image in a non-network-based environment involves the following high-level steps:

1. Download the RHEL for Edge Container. See [Downloading a RHEL for Edge image](#) for information about how to download the RHEL for Edge image.
2. Load the RHEL for Edge Container image into Podman
3. Run the RHEL for Edge Container image into Podman
4. Load the RHEL for Edge Installer blueprint
5. Build the RHEL for Edge Installer image
6. Prepare a **.qcow2** disk
7. Boot the Virtual Machine (VM)
8. Install the image

9.1. CREATING A RHEL FOR EDGE CONTAINER IMAGE FOR NON-NETWORK-BASED DEPLOYMENTS

You can build a running container by loading the downloaded RHEL for Edge Container OSTree commit into Podman. For that, follow the steps:

Prerequisites

- You created and downloaded a RHEL for Edge Container OSTree commit.
- You have installed **Podman** on your system. See [How do I install Podman in RHEL](#) .

Procedure

1. Navigate to the directory where you have downloaded the RHEL for Edge Container OSTree commit.
2. Load the RHEL for Edge Container OSTree commit into **Podman**.

```
$ sudo podman load -i UUID-container.tar
```

The command output provides the image ID, for example:

@8e0d51f061ff1a51d157804362bc875b649b27f2ae1e66566a15e7e6530cec63

3. Tag the new RHEL for Edge Container image, using the image ID generated by the previous step.

—

```
$ sudo podman tag image-ID localhost/edge-container
```

The **podman tag** command assigns an additional name to the local image.

4. Run the container named **edge-container**.

```
$ sudo podman run -d --name=edge-container -p 8080:8080 localhost/edge-container
```

The **podman run -d --name=edge-container** command assigns a name to your container-based on the **localhost/edge-container** image.

5. List containers:

```
$ sudo podman ps -a
```

CONTAINER ID	IMAGE	COMMAND	CREATED	STATUS
2988198c4c4b/localhost/edge-container	/bin/bash	3 seconds ago	Up 2 seconds ago
	edge-container			

As a result, **Podman** runs a container that serves an OSTree repository with the RHEL for Edge Container commit.

9.2. CREATING A RHEL FOR EDGE INSTALLER IMAGE FOR NON-NETWORK-BASED DEPLOYMENTS

After you have built a running container to serve a repository with the **RHEL for Edge Container** commit, create an **RHEL for Edge Installer (.iso)** image. The **RHEL for Edge Installer (.iso)** pulls the commit served by **RHEL for Edge Container (.tar)**. After the **RHEL for Edge Container** commit is loaded in Podman, it exposes the **OSTree** in the URL format.

To create the RHEL for Edge Installer image in the CLI, follow the steps:

Prerequisites

- You created a blueprint for RHEL for Edge image.
- You created a RHEL for Edge Edge Container image and deployed it using a web server.

Procedure

1. Begin to create the RHEL for Edge Installer image.

```
# composer-cli compose start-ostree --ref rhel/8/x86_64/edge --url URL-OSTree-repository
blueprint-name image-type
```

Where,

- *ref* is the same value that customer used to build ostree repository
- *URL-OSTree-repository* is the URL to the OSTree repository of the commit to embed in the image. For example, <http://10.0.2.2:8080/repo/>. See [Creating a RHEL for Edge Container image for non-network-based deployments](#).
- *blueprint-name* is the RHEL for Edge Installer blueprint name.

- *image-type* is **edge-installer**.

A confirmation that the composer process has been added to the queue appears. It also shows a Universally Unique Identifier (UUID) number for the image created. Use the UUID number to track your build. Also keep the UUID number handy for further tasks.

2. Check the image compose status.

```
# composer-cli compose status
```

The command output displays the status in the following format:

```
<UUID> RUNNING date blueprint-name blueprint-version image-type
```



NOTE

The image creation process takes a few minutes to complete.

To interrupt the image creation process, run:

```
# composer-cli compose cancel <UUID>
```

To delete an existing image, run:

```
# composer-cli compose delete <UUID>
```

RHEL image builder pulls the commit that is being served by the running container during the image build.

After the image build is complete, you can download the resulting ISO image.

3. Download the image. See [Downloading a RHEL for Edge image](#).

After the image is ready, you can use it for **non-network deployments**. See [Installing the RHEL for Edge image for non-network-based deployments](#).

Additional resources

- [Creating a RHEL for Edge Installer image using command-line interface for non-network-based deployments](#).

9.3. INSTALLING THE RHEL FOR EDGE IMAGE FOR NON-NETWORK-BASED DEPLOYMENTS

To install the RHEL for Edge image, follow the steps:

Prerequisites

- You created a RHEL for Edge Installer ISO image.
- You stopped the running container.
- A disk image to install the commit you created.

- You installed the **edk2-ovmf** package.
- You installed the **virt-viewer** package.
- You customized your blueprint with a user account. See [Creating a blueprint for a RHEL for Edge image using RHEL image builder in RHEL web console](#).



WARNING

If you do not define a user account customization in your blueprint, you will not be able to login to the ISO image.

Procedure

1. Create a **qcow** VM disk file to install the (**.iso**) image. That is an image of a hard disk drive for the virtual machine (VM). For example:

```
$ qemu-img create -f qcow2 diskfile.qcow2 20G
```

2. Use the **virt-install** command to boot the VM using the disk as a drive and the installer ISO image as a CD-ROM. For example:

```
$ virt-install \
--boot uefi \
--name VM_NAME \
--memory 2048 \
--vcpus 2 \
--disk path=diskfile.qcow2 \
--cdrom /var/lib/libvirt/images/UUID-installer.iso \
--os-variant rhel9.0
```

This command instructs **virt-install** to:

- Instructs the VM to use UEFI to boot, instead of the BIOS.
- Mount the installation ISO.
- Use the hard disk drive image created in the first step.
It gives you an Anaconda Installer. The RHEL Installer starts, loads the Kickstart file from the ISO and executes the commands, including the command to install the RHEL for Edge image commit. Once the installation is complete, the RHEL installer prompts for login details.



NOTE

Anaconda is preconfigured to use the Container commit during the installation. However, you need to set up system configurations, such as disk partition, timezone, between others.

3. Connect to Anaconda GUI with **virt-viewer** to setup the system configuration:

```
$ virt-viewer --connect qemu:///system --wait VM_NAME
```

4. Reboot the system to finish the installation.
5. On the login screen, specify your user account credentials and click **Enter**.

Verification steps

1. Verify whether the RHEL for Edge image is successfully installed.

```
$ rpm-ostree status
```

The command output provides the image commit ID and shows that the installation is successful.

CHAPTER 10. MANAGING RHEL FOR EDGE IMAGES

To manage the RHEL for Edge images, you can perform any of the following administrative tasks:

- Edit the RHEL for Edge image blueprint by using image builder in RHEL web console or in the command-line
- Build a commit update by using image builder command-line
- Update the RHEL for Edge images
- Configure **rpm-ostree** remotes on nodes, to update node policy
- Restore RHEL for Edge images manually or automatically by using greenboot

10.1. EDITING A RHEL FOR EDGE IMAGE BLUEPRINT BY USING IMAGE BUILDER

You can edit the RHEL for Edge image blueprint to:

- Add additional components that you might require
- Modify the version of any existing component
- Remove any existing component

10.1.1. Adding a component to RHEL for Edge blueprint using image builder in RHEL web console

To add a component to a RHEL for Edge image blueprint, ensure that you have met the following prerequisites and then follow the procedure to edit the corresponding blueprint.

Prerequisites

- On a RHEL system, you have accessed the RHEL image builder dashboard.
- You have created a blueprint for RHEL for Edge image.

Procedure

1. On the RHEL image builder dashboard, click the blueprint that you want to edit.
To search for a specific blueprint, enter the blueprint name in the filter text box, and click **Enter**.
2. On the upper right side of the blueprint, click **Edit Packages**.
The **Edit blueprints** wizard opens.
3. On the **Details** page, update the blueprint name and click **Next**.
4. On the **Packages** page, follow the steps:
 - a. In the **Available Packages**, enter the package name that you want to add in the filter text box, and click **Enter**.
A list with the component name appears.
 - b. Click **>** to add the component to the blueprint.

5. On the **Review** page, click **Save**.
The blueprint is now updated with the new package.

10.1.2. Removing a component from a blueprint using RHEL image builder in the web console

To remove one or more unwanted components from a blueprint that you created by using RHEL image builder, ensure that you have met the following prerequisites and then follow the procedure.

Prerequisites

- On a RHEL system, you have accessed the RHEL image builder dashboard.
- You have created a blueprint for RHEL for Edge image.
- You have added at least one component to the RHEL for Edge blueprint.

Procedure

1. On the RHEL image builder dashboard, click the blueprint that you want to edit.
To search for a specific blueprint, enter the blueprint name in the filter text box, and click **Enter**.
2. On the upper right side of the blueprint, click **Edit Packages**.
The **Edit blueprints** wizard opens.
3. On the **Details** page, update the blueprint name and click **Next**.
4. On the **Packages** page, follow the steps:
 - a. From the **Chosen packages**, click **<** to remove the chosen component. You can also click **<<** to remove all the packages at once.
5. On the **Review** page, click **Save**.
The blueprint is now updated.

10.1.3. Editing a RHEL for Edge image blueprint using command-line interface

You can change the specifications for your RHEL for Edge image blueprint by using RHEL image builder command-line. To do so, ensure that you have met the following prerequisites and then follow the procedure to edit the corresponding blueprint.

Prerequisites

- You have access to the RHEL image builder command-line.
- You have created a RHEL for Edge image blueprint.

Procedure


1. Save (export) the blueprint to a local text file:

```
# composer-cli blueprints save BLUEPRINT-NAME
```

2. Edit the **BLUEPRINT-NAME.toml** file with a text editor of your choice and make your changes.

Before finishing with the edits, verify that the file is a valid blueprint:

3. Increase the version number.
Ensure that you use a Semantic Versioning scheme.



NOTE

if you do not change the version, the patch component of the version is increased automatically.

4. Check if the contents are valid TOML specifications. See the TOML documentation for more information.




NOTE

TOML documentation is a community product and is not supported by Red Hat. You can report any issues with the tool at <https://github.com/toml-lang/toml/issues>.

5. Save the file and close the editor.
6. Push (import) the blueprint back into RHEL image builder server:

```
# composer-cli blueprints push BLUEPRINT-NAME.toml
```



NOTE

When pushing the blueprint back into the RHEL image builder server, provide the file name including the **.toml** extension.

7. Verify that the contents uploaded to RHEL image builder match your edits:

```
# composer-cli blueprints show BLUEPRINT-NAME
```

8. Check whether the components and versions listed in the blueprint and their dependencies are valid:

```
# composer-cli blueprints depsolve BLUEPRINT-NAME
```

10.2. UPDATING RHEL FOR EDGE IMAGES

10.2.1. How RHEL for Edge image updates are deployed

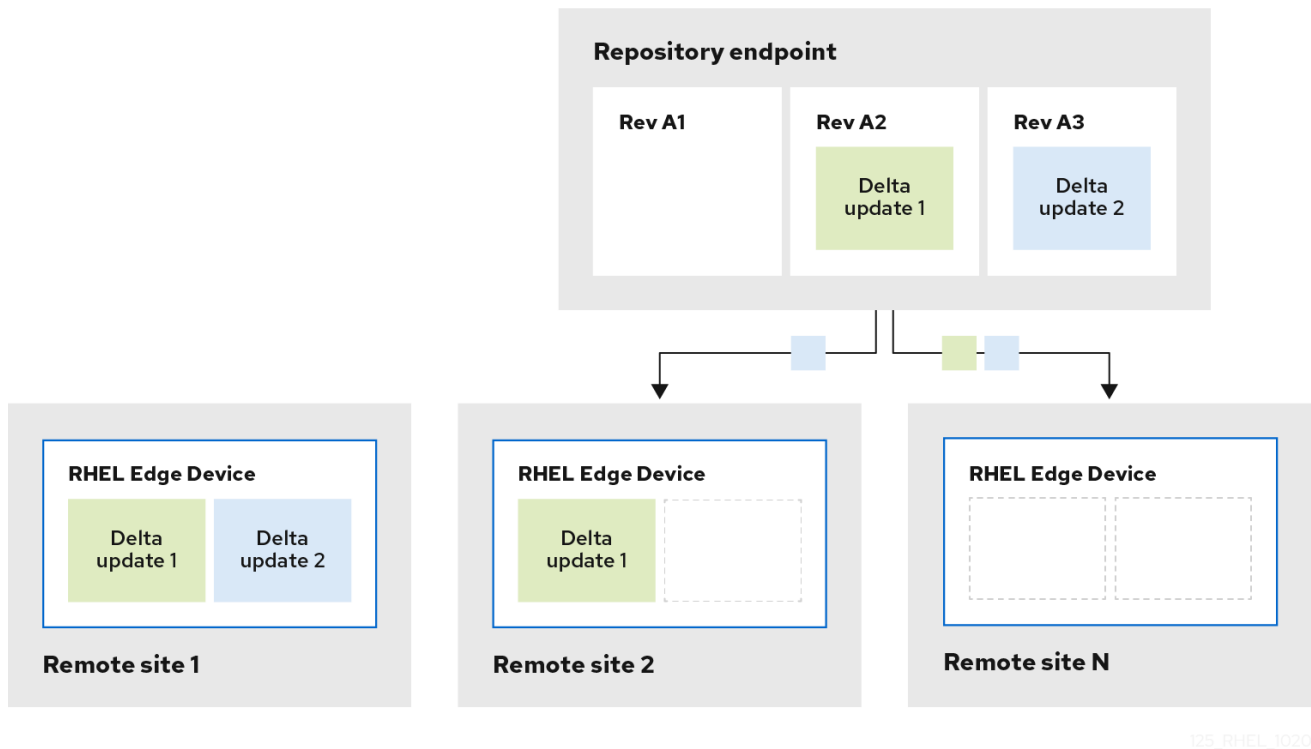
With RHEL for Edge images, you can either deploy the updates manually or can automate the deployment process. The updates are applied in an atomic manner, where the state of each update is known, and the updates are staged and applied only upon reboot. Because no changes are seen until you reboot the device, you can schedule a reboot to ensure the highest possible uptime.

During the image update, only the updated operating system content is transferred over the network. This makes the deployment process more efficient compared to transferring the entire image. The operating system binaries and libraries in **/usr** are **read-only**, and the **read and write** state is maintained

in **/var** and **/etc** directories.

When moving to a new deployment, the **/etc** and the **/var** directories are copied to the new deployment with **read and write** permissions. The **/usr** directory is copied as a soft link to the new deployment directory, with **read-only** permissions.

The following diagram illustrates the RHEL for Edge image update deployment process:



By default, the new system is booted using a procedure similar to a **chroot** operation, that is, the system enables control access to a filesystem while controlling the exposure to the underlying server environment. The new **/sysroot** directory mainly has the following parts:

- Repository database at the **/sysroot/ostree/repo** directory.
- File system revisions at the **/sysroot/ostree/deploy/rhel/deploy** directory, which are created by each operation in the system update.
- The **/sysroot/ostree/boot** directory, which links to deployments on the previous point. Note that **/ostree** is a soft link to **/sysroot/ostree**. The files from the **/sysroot/ostree/boot** directory are not duplicated. The same file is used if it is not changed during the deployment. The files are hard-links to another file stored in the **/sysroot/ostree/repo/objects** directory.

The operating system selects the deployment in the following way:

1. The **dracut** tool parses the **ostree** kernel argument in the **initramfs root** file system and sets up the **/usr** directory as a **read-only** bind mount.
2. Bind the deployment directory in **/sysroot** to **/** directory.
3. Re-mount the operating system already mounted **dirs** using the **MS_MOVE** mount flag

If anything goes wrong, you can perform a deployment rollback by removing the old deployments with the **rpm-ostree** cleanup command. Each client machine contains an **OSTree** repository stored in **/ostree/repo**, and a set of deployments stored in **/ostree/deploy/\$STATEROOT/\$CHECKSUM**.

With the deployment updates in RHEL for Edge image, you can benefit from a better system consistency across multiple devices, easier reproducibility, and better isolation between the pre and post system states change.

10.2.2. Building a commit update

You can build a commit update after making a change in the blueprint, such as:

- Adding an additional package that your system requires
- Modifying the package version of any existing component
- Removing any existing package.

Prerequisites

- You have updated a system which is running RHEL image builder.
- You created a blueprint update.
- You have previously created an OSTree repository and served it through HTTP. See [Setting up a web server to install RHEL for Edge images](#).

Procedure

1. Start the compose of the new commit image, with the following arguments: **--url, --ref, blueprint-name, edge-commit**.

```
# composer-cli compose start-ostree --ref rhel/8/x86_64/edge --url http://localhost:8080/repo
<blueprint-name> edge-commit
```

The command instructs the compose process to fetch the metadata from the OSTree repo before starting the compose. The resulting new OSTree commit contains a reference of the original OSTree commit as a parent image.

2. After the compose process finishes, fetch the **.tar** file.

```
# composer-cli compose image <UUID>
```

3. Extract the commit to a temporary directory, so that you can store the commit history in the OSTree repo.

```
$ tar -xf UUID.tar -C /var/tmp
```

4. Inspect the resulting OSTree repo commit, by using the **tar -xf** command. It extracts the tar file to disk so you can inspect the resulting OSTree repo:

```
$ ostree --repo=/var/tmp/repo log rhel/8/x86_64/edge
commit d523ef801e8b1df69ddbf73ce810521b5c44e9127a379a4e3bba5889829546fa
Parent: f47842de7e6859cee07d743d3c67949420874727883fa9dbbaeb5824ad949d91
ContentChecksum:
```

```
f0f6703696331b661fa22d97358db48ba5f8b62711d9db83a00a79b3ae0dfe16
Date: 2023-06-04 20:22:28 /+0000
Version: 8
```

In the output example, there is a single OSTree commit in the repo that references a parent commit. The parent commit is the same checksum from the original OSTree commit that you previously made.

5. Merge the two commits by using the **ostree pull-local** command:

```
$ sudo ostree --repo=/var/srv/httpd/repo pull-local /var/tmp/repo
20 metadata, 22 content objects imported; 0 bytes content written
```

This command copies any new metadata and content from the location on the disk, for example, **/var/tmp**, to a destination OSTree repo in **/var/srv/httpd**.

Verification

1. Inspect the target OSTree repo:

```
$ ostree --repo=/var/srv/httpd/repo log rhel/8/x86_64/edge
commit d523ef801e8b1df69ddb73ce810521b5c44e9127a379a4e3bba5889829546fa
Parent: f47842de7e6859cee07d743d3c67949420874727883fa9dbbaeb5824ad949d91
ContentChecksum:
f0f6703696331b661fa22d97358db48ba5f8b62711d9db83a00a79b3ae0dfe16
Date: 2023-06-04 20:22:28 /+0000
Version: 8
(no subject)

commit f47842de7e6859cee07d743d3c67949420874727883fa9dbbaeb5824ad949d91
ContentChecksum:
9054de3fe5f1210e3e52b38955bea0510915f89971e3b1ba121e15559d5f3a63
Date: 2023-06-04 20:01:08 /+0000
Version: 8
(no subject)
```

You can see that the target OSTree repo now contains two commits in the repository, in a logical order. After successful verification, you can update your RHEL for Edge system.

10.2.3. Deploying RHEL for Edge image updates manually

After you have edited a RHEL for Edge blueprint, you can update the image commit. RHEL image builder generates a new commit for the updated RHEL for Edge image. Use this new commit to deploy the image with latest package versions or with additional packages.

To deploy RHEL for Edge images updates, ensure that you meet the prerequisites and then follow the procedure.

Prerequisites

- On a RHEL system, you have accessed the RHEL image builder dashboard.
- You have created a RHEL for Edge image blueprint.
- You have edited the RHEL for Edge image blueprint.

Procedure

1. On the RHEL image builder dashboard click **Create Image**.
2. On the **Create Image** window, perform the following steps:
 - a. In the Image output page:
 - i. From the **Select a blueprint** dropdown list, select the blueprint that you edited.
 - ii. From the **Image output type** dropdown list, select **RHEL for Edge Commit (.tar)**. Click **Next**.
 - b. In the **OSTree settings** page, enter:
 - i. In the **Repository URL** field, enter the URL to the OSTree repository of the commit to embed in the image. For example, `http://10.0.2.2:8080/repo/`. See [Setting up a web server to install RHEL for Edge image](#).
 - ii. In the **Parent commit** field, specify the parent commit ID that was previously generated. See [Extracting RHEL for Edge image commit](#).
 - iii. In the **Ref** field, you can either specify a name for your commit or leave it empty. By default, the web console specifies the **Ref** as **rhel/8/arch_name/edge**. Click **Next**.
 - c. In the **Review** page, check the customizations and click **Create image**. RHEL image builder starts to create a RHEL for Edge image for the updated blueprint. The image creation process takes a few minutes to complete.
To view the RHEL for Edge image creation progress, click the blueprint name from the breadcrumbs, and then click the **Images** tab.

The resulting image includes the latest packages that you have added, if any, and have the original **commit ID** as a parent.
3. Download the resulting RHEL for Edge Commit (.tar) image.
 - a. From the **Images** tab, click **Download** to save the RHEL for Edge Commit (.tar) image to your system.

4. Extract the OSTree commit (.tar) file.

```
# tar -xf UUID-commit.tar -C UPGRADE_FOLDER
```

5. Upgrade the OSTree repo:

```
# ostree pull-local --repo http://10.0.2.2:8080/repo UPGRADE_FOLDER/repo OSTREE_REF
# ostree summary --update --repo http://10.0.2.2:8080/repo
```

6. Build a docker container, serving the child commit ID this time.

```
# podman build -t name-of-server --build-arg commit=UUID-child_commit.tar .
```

7. Run the container.

```
# podman run --rm -p 8000:80 name-of-server
```

8. On the RHEL system provisioned, from the original edge image, verify the current status.

```
$ rpm-ostree status
```

If there is no new commit ID, run the following command to verify if there is any upgrade available:

```
$ rpm-ostree upgrade --check
```

The command output provides the current active OSTree commit ID.

9. Update OSTree to make the new OSTree commit ID available.

```
$ rpm-ostree upgrade
```

OSTree verifies if there is an update on the repository. If yes, it fetches the update and requests you to reboot your system so that you can activate the deployment of this new commit update.

10. Check the current status again:

```
$ rpm-ostree status
```

You can now see that there are 2 commits available:

- The active parent commit.
- A new commit that is not active and contains 1 added difference.

11. To activate the new deployment and to make the new commit active, reboot your system.

```
# systemctl reboot
```

The Anaconda Installer reboots into the new deployment. On the login screen, you can see a new deployment available for you to boot.

12. If you want to boot into the newest deployment (commit), the **rpm-ostree** upgrade command automatically orders the boot entries so that the new deployment is first in the list. Optionally, you can use the arrow key on your keyboard to select the GRUB menu entry and press **Enter**.

13. Provide your login user account credentials.

14. Verify the OSTree status:

```
$ rpm-ostree status
```

The command output provides the active commit ID.

15. To view the changed packages, if any, run a diff between the parent commit and the new commit:

```
$ rpm-ostree db diff parent_commit new_commit
```

The update shows that the package you have installed is available and ready for use.

10.2.4. Deploying RHEL for Edge image updates manually using the command-line

After you have edited a RHEL for Edge blueprint, you can update the image commit. RHEL image builder generates a new commit for the updated RHEL for Edge image. Use the new commit to deploy the image with latest package versions or with additional packages using the CLI.

To deploy RHEL for Edge image updates using the CLI, ensure that you meet the prerequisites, and then follow the procedure.

Prerequisites

- You created the RHEL for Edge image blueprint.
- You edited the RHEL for Edge image blueprint. See [Editing a RHEL for Edge image blueprint using command-line interface](#).

Procedure

1. Create the RHEL for Edge Commit (**.tar**) image with the following arguments:

```
# composer-cli compose start-ostree --ref ostree_ref --url URL-OSTree-repository -  
blueprint_name_ image-type
```

where

- **ref** is the reference you provided during the creation of the RHEL for Edge Container commit. For example, **rhel/8/x86_64/edge**.
- **URL-OSTree-repository** is the URL to the OSTree repository of the commit to embed in the image. For example, <http://10.0.2.2:8080/repo/>. See [Setting up a web server to install RHEL for Edge image](#).
- **image-type** is **edge-commit**.
RHEL image builder creates a RHEL for Edge image for the updated blueprint.

2. Check the RHEL for Edge image creation progress:

```
# composer-cli compose status
```



NOTE

The image creation processes can take up to ten to thirty minutes to complete.

The resulting image includes the latest packages that you have added, if any, and has the original **commit ID** as a parent.

3. Download the resulting RHEL for Edge image. For more information, see [Downloading a RHEL for Edge image using the RHEL image builder command-line interface](#).
4. Extract the OSTree commit.

```
# tar -xf UUID-commit.tar -C UPGRADE_FOLDER
```

5. Serve the OSTree commit by using httpd. See [Setting up a web server to install RHEL for Edge image](#).

6. Upgrade the OSTree repo:

```
# ostree pull-local --repo http://10.0.2.2:8080/repo UPGRADE_FOLDER/repo OSTREE_REF
# ostree summary --update --repo http://10.0.2.2:8080/repo
```

7. On the RHEL system provisioned from the original edge image, verify the current status:

```
$ rpm-ostree status
```

If there is no new commit ID, run the following command to verify if there is any upgrade available:

```
$ rpm-ostree upgrade --check
```

The command output provides the current active OSTree commit ID.

8. Update OSTree to make the new OSTree commit ID available:

```
$ rpm-ostree upgrade
```

OSTree verifies if there is an update on the repository. If yes, it fetches the update and requests you to reboot your system so that you can activate the deployment of the new commit update.

9. Check the current status again:

```
$ rpm-ostree status
```

You should now see that there are 2 commits available:

- The active parent commit
- A new commit that is not active and contains 1 added difference

10. To activate the new deployment and make the new commit active, reboot your system:

```
# systemctl reboot
```

The Anaconda Installer reboots into the new deployment. On the login screen, you can see a new deployment available for you to boot.

11. If you want to boot into the newest deployment, the **rpm-ostree upgrade** command automatically orders the boot entries so that the new deployment is first in the list. Optionally, you can use the arrow key on your keyboard to select the GRUB menu entry and press **Enter**.
12. Log in using your account credentials.
13. Verify the OSTree status:

```
$ rpm-ostree status
```

The command output provides the active commit ID.

14. To view the changed packages, if any, run a diff between the parent commit and the new commit:

```
$ rpm-ostree db diff parent_commit new_commit
```

The update shows that the package you have installed is available and ready for use.

10.2.5. Deploying RHEL for Edge image updates manually for non-network-base deployments

After editing a RHEL for Edge blueprint, you can update your RHEL for Edge Commit image with those updates. Use RHEL image builder to generate a new commit to update your RHEL for Edge image that is already deployed in a VM, for example. Use this new commit to deploy the image with latest package versions or with additional packages.

To deploy RHEL for Edge images updates, ensure that you meet the prerequisites and then follow the procedure.

Prerequisites

- On your host, you have opened the RHEL image builder app from the web console in a browser.
- You have a RHEL for Edge system provisioned that is up and running.
- You have an OSTree repository that is being served over HTTP.
- You have edited a previously created RHEL for Edge image blueprint.

Procedure

1. On your system host, on the RHEL image builder dashboard, click **Create Image**.
2. On the **Create Image** window, perform the following steps:
 - a. In the **Image output** page:
 - i. From the **Select a blueprint** dropdown list, select the blueprint that you edited.
 - ii. From the **Image output type** dropdown list, select **RHEL for Edge Container (.tar)**.
 - iii. Click **Next**.
 - b. In the **OSTree settings** page, enter:
 - i. In the **Repository URL** field, enter the URL to the OSTree repository of the commit to embed in the image. For example, <http://10.0.2.2:8080/repo/>. See [Setting up a web server to install RHEL for Edge image](#).
 - ii. In the **Parent commit** field, specify the parent commit ID that was previously generated. See [Extracting RHEL for Edge image commit](#).
 - iii. In the **Ref** field, you can either specify a name for your commit or leave it empty. By default, the web console specifies the **Ref** as **rhel/8/arch_name/edge**.
 - iv. Click **Next**.

- c. In the **Review** page, check the customizations and click **Create**.
RHEL image builder creates a RHEL for Edge image for the updated blueprint.
- d. Click the **Images** tab to view the progress of RHEL for Edge image creation.

**NOTE**

The image creation process takes a few minutes to complete.

The resulting image includes the latest packages that you have added, if any, and has the original **commit ID** as a parent.

3. Download the resulting RHEL for Edge image on your host:
 - a. From the **Images** tab, click **Download** to save the RHEL for Edge Container (**.tar**) image to your host system.
4. On the RHEL system provisioned from the original edge image, perform the following steps:
 - a. Load the RHEL for Edge Container image into Podman, serving the child commit ID this time.

```
$ cat ./child-commit_ID-container.tar | sudo podman load
```

- b. Run **Podman**.

```
# sudo podman run -p 8080:8080 localhost/edge-test
```

- c. Upgrade the OSTree repo:

```
# ostree pull-local --repo http://10.0.2.2:8080/repo UPGRADE_FOLDER/repo
OSTREE_REF
# ostree summary --update --repo http://10.0.2.2:8080/repo
```

- d. On the RHEL system provisioned, from the original edge image, verify the current status.

```
$ rpm-ostree status
```

If there is no new commit ID, run the following command to verify if there is any upgrade available:

```
$ rpm-ostree upgrade --check
```

If there are updates available, the command output provides information about the available updates in the OSTree repository, such as the current active OSTree commit ID. Else, it prompts a message informing that there are no updates available.

- e. Update OSTree to make the new OSTree commit ID available.

```
$ rpm-ostree upgrade
```

OSTree verifies if there is an update on the repository. If yes, it fetches the update and requests you to reboot your system so that you can activate the deployment of this new commit update.

- f. Check the current system status:

```
$ rpm-ostree status
```

You can now see that there are 2 commits available:

- The active parent commit.
- A new commit that is not active and contains 1 added difference.

- g. To activate the new deployment and to make the new commit active, reboot your system.

```
# systemctl reboot
```

The Anaconda Installer reboots into the new deployment. On the login screen, you can see a new deployment available for you to boot.

- h. To boot into the newest commit, run the following command to automatically order the boot entries so that the new deployment is first in the list:

```
$ rpm-ostree upgrade
```

Optionally, you can use the arrow key on your keyboard to select the GRUB menu entry and press **Enter**.

5. Provide your login user account credentials.

6. Verify the OSTree status:

```
$ rpm-ostree status
```

The command output provides the active commit ID.

7. To view the changed packages, if any, run a diff between the parent commit and the new commit:

```
$ rpm-ostree db diff parent_commit new_commit
```

The update shows that the package you have installed is available and ready for use.

10.3. DEPLOYING RHEL FOR EDGE AUTOMATIC IMAGE UPDATES

After you install a RHEL for Edge image on an Edge device, you can check for image updates available, if any, and can auto-apply them.

The **rpm-ostreed-automatic.service** (systemd service) and **rpm-ostreed-automatic.timer** (systemd timer) control the frequency of checks and upgrades. The available updates, if any, appear as staged deployments.

Deploying automatic image updates involves the following high-level steps:

- Update the image update policy
- Enable automatic download and staging of updates

10.3.1. Updating the RHEL for Edge image update policy

To update the image update policy, use the **AutomaticUpdatePolicy** and an **IdleExitTimeout** setting from the **rpm-ostreed.conf** file at **/etc/rpm-ostreed.conf** location on an Edge device.

The **AutomaticUpdatePolicy** settings controls the automatic update policy and has the following update checks options:

- **none**: Disables automatic updates. By default, the **AutomaticUpdatePolicy** setting is set to **none**.
- **check**: Downloads enough metadata to display available updates with **rpm-ostree** status.
- **stage**: Downloads and unpacks the updates that are applied on a reboot.

The **IdleExitTimeout** setting controls the time in seconds of inactivity before the daemon exit and has the following options:

- 0: Disables auto-exit.
- 60: By default, the **IdleExitTimeout** setting is set to **60**.

To enable automatic updates, perform the following steps:

Procedure

1. In the **/etc/rpm-ostreed.conf** file, update the following:
 - Change the value of **AutomaticUpdatePolicy** to **check**.
 - To run the update checks, specify a value in seconds for **IdleExitTimeout**.
2. Reload the **rpm-ostreed** service and enable the **systemd** timer.

```
# systemctl reload rpm-ostreed
# systemctl enable rpm-ostreed-automatic.timer --now
```

3. Verify the **rpm-ostree** status to ensure the automatic update policy is configured and time is active.

```
# rpm-ostree status
```

The command output shows the following:

```
State: idle; auto updates enabled (check; last run <minutes> ago)
```

Additionally, the output also displays information about the available updates.

10.3.2. Enabling RHEL for Edge automatic download and staging of updates

After you update the image update policy to check for image updates, the updates if any are displayed along with the update details. If you decide to apply the updates, enable the policy to automatically download and stage the updates. The available image updates are then downloaded and staged for deployment. The updates are applied and take effect when you reboot the Edge device.

To enable the policy for automatic download and staging of updates, perform the following updates:

Procedure

1. In the `/etc/rpm-ostreed.conf` file, update 'AutomaticUpdatePolicy' to **stage**.
2. Reload the rpm-ostreed service.

```
# systemctl enable rpm-ostreed-automatic.timer --now
```

3. Verify the rpm-ostree status

```
# rpm-ostree status
```

The command output shows the following:

```
State: idle
AutomaticUpdates: stage; rpm-ostreed-automatic.timer: last run <time> ago
```

4. To initiate the updates, you can either wait for the timer to initiate the updates, or can manually start the service.

```
# systemctl start rpm-ostreed-automatic.service
```

After the updates are initiated, the rpm-ostree status shows the following:

```
# rpm-ostree status
State: busy
AutomaticUpdates: stage; rpm-ostreed-automatic.service: running
Transaction: automatic (stage)
```

When the update is complete, a new deployment is staged in the list of deployments, and the original booted deployment is left untouched. You can decide if you want to boot the system using the new deployment or can wait for the next update.

To view the list of deployments, run the **rpm-ostree status** command.

Following is a sample output.

```
# rpm-ostree status
State: idle
AutomaticUpdates: stage; rpm-ostreed-automatic.timer: last run <time> ago
Deployments:
```

To view the list of deployments with the updated package details, run the **rpm-ostree status -v** command.

10.4. ROLLING BACK RHEL FOR EDGE IMAGES

Because RHEL for Edge applies transactional updates to the operating system, you can either manually or automatically roll back the unsuccessful updates to the last known good state, which prevents system failure during updates. You can automate the verification and rollback process by using the **greenboot** framework.

The **greenboot** health check framework leverages **rpm-ostree** to run custom health checks on system startup. In case of an issue, the system rolls back to the last working state. When you deploy a **rpm-ostree** update, it runs scripts to check that critical services can still work after the update. If the system does not work, for example, due to some failed package, you can roll back the system to a previous stable version of the system. This process ensures that your RHEL for Edge device is in an operational state.

After you update an image, it creates a new image deployment while preserving the previous image deployment. You can verify whether the update was successful. If the update is unsuccessful, for example, due to a failed package, you can roll back the system to a previous stable version.

10.4.1. Introducing the greenboot checks

Greenboot is a Generic Health Check Framework for **systemd** available on **rpm-ostree** based systems. It contains the following RPM packages that you can install on your system:

- **greenboot** - a package that contains the following functionalities:
 - Checking provided scripts
 - Reboot the system if the check fails
 - Rollback to a previous deployment the reboot did not solve the issue.
- **greenboot-default-health-checks** - a set of optional and selected health checks provided by your **greenboot** system maintainers.

Greenboot works in a RHEL for Edge system by using health check scripts that run on the system to assess the system health and automate a rollback to the last healthy state in case of some software fails. These health checks scripts are available in the **/etc/greenboot/check/required.d** directory.

Greenboot supports shell scripts for the health checks. Having a health check framework is especially useful when you need to check for software problems and perform system rollbacks on edge devices where direct serviceability is either limited or non-existent. When you install and configure health check scripts, it triggers the health checks to run every time the system starts.

You can create your own 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.



NOTE

You cannot use rollback in case of any health check failure on a system that is not using OSTree.

10.4.2. RHEL for Edge images roll back with greenboot

With RHEL for Edge images, only transactional updates are applied to the operating system. The transactional updates are atomic, which means that the updates are applied only if all the updates are successful, and there is support for rollbacks. With the transactional updates, you can easily rollback the unsuccessful updates to the last known good state, preventing system failure during updates.

Performing health checks is especially useful when you need to check for software problems and perform system rollbacks on edge devices where direct serviceability is limited or non-existent.



NOTE

You cannot use rollback in case of an update failure on a system that is not using OSTree, even if health checks might run.

You can use intelligent rollbacks with the **greenboot** health check framework to automatically assess system health every time the system starts. You can obtain pre-configured health from the **greenboot-default-health-checks** subpackage. These checks are located in the `/usr/lib/greenboot/check` read-only directory in **rpm-ostree** systems.

Greenboot leverages **rpm-ostree** and runs custom health checks that run on system startup. In case of an issue, the system rolls back the changes and preserves the last working state. When you deploy an **rpm-ostree** update, it runs scripts to check that critical services can still work after the update. If the system does not work, the update rolls back to the last known working version of the system. This process ensures that your RHEL for Edge device is in an operational state.

You can obtain pre-configured health from the **greenboot-default-health-checks** subpackage. These checks are located in the `/usr/lib/greenboot/check` read-only directory in **rpm-ostree** systems. You can also configure shell scripts as the following types of checks:

Example 10.1. The greenboot directory structure

```

etc
├── greenboot
│   ├── check
│   │   ├── required.d
│   │   ├── init.py
│   ├── green.d
│   └── red.d

```

Required

Contains the health checks that must not fail. Place required shell scripts in the `/etc/greenboot/check/required.d` directory. If the scripts fail, greenboot retries them three times by default. You can configure the number of retries in the `/etc/greenboot/greenboot.conf` file by setting the **GREENBOOT_MAX_BOOTS** parameter to the number of retries you want.

After all retries fail, **greenboot** automatically initiates a rollback if one is available. If a rollback is not available, the system log output shows that you need to perform a manual intervention.

Wanted

Contains the health checks that might fail without causing the system to roll back. Place wanted shell scripts in the `/etc/greenboot/check/wanted.d` directory. **Greenboot** informs that the script fails, the system health status remains unaffected and it does not perform a rollback neither a reboot.

You can also specify shell scripts that will run after a check:

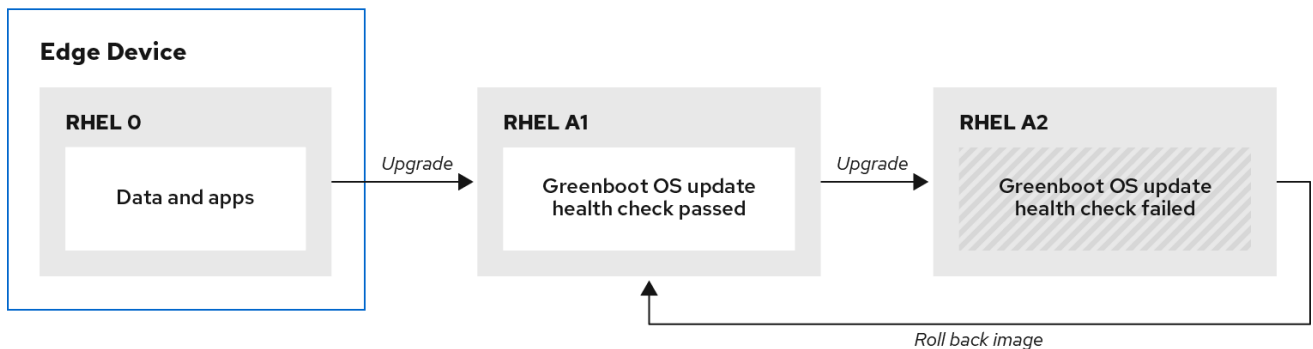
Green

Contains the scripts to run after a successful boot. Place these scripts into the `/etc/greenboot/green.d` directory. **Greenboot** informs that the boot was successful.

Red

Contains the scripts to run after a failed boot. Place these scripts into the `/etc/greenboot/red.d` directory. The system attempts to boot three times and in case of failure, it executes the scripts. **Greenboot** informs that the boot failed.

The following diagram illustrates the RHEL for Edge image roll back process.



125_RHEL_1020

After booting the updated operating system, **greenboot** runs the scripts in the **required.d** and **wanted.d** directories. If any of the scripts fail in the **required.d** directory, **greenboot** runs any scripts in the **red.d** directory, and then reboots the system.

Greenboot makes 2 more attempts to boot on the upgraded system. If during the third boot attempt the scripts in **required.d** are still failing, **greenboot** runs the **red.d** scripts one last time, to ensure that the script in the **red.d** directory tried to make a corrective action to fix the issue and this was not successful. Then, **greenboot** rolls back the system from the current **rpm-ostree** deployment to the previous stable deployment.

10.4.3. Greenboot health check status

When deploying your updated system, wait until the greenboot health checks have finished before making the changes to ensure that those changes are not lost if greenboot rolls the system back to an earlier state. If you want to make configuration changes or deploy applications you must wait until the greenboot health checks have finished. This ensures that your changes are not lost if greenboot rolls your **rpm-ostree** system back to an earlier state.

The **greenboot-healthcheck** service runs once and then exits. You can check the status of the service to know if it is done, and to know the outcome, by using the following commands:

systemctl is-active greenboot-healthcheck.service

This command reports **active** when the service has exited. If the service did not even run it shows **inactive**.

systemctl show --property=SubState --value greenboot-healthcheck.service

Reports **exited** when done, **running** while still running.

systemctl show --property=Result --value greenboot-healthcheck.service

Reports **success** when the checks passed.

systemctl show --property=ExecMainStatus --value greenboot-healthcheck.service

Reports the numerical exit code of the service, 0 means success and nonzero values mean a failure occurred.

cat /run/motd.d/boot-status

Shows a message, such as "Boot Status is GREEN - Health Check SUCCESS".

10.4.4. Checking greenboot health checks statuses

Check the status of greenboot health checks before making changes to the system or during troubleshooting. You can use any of the following commands to help you ensure that greenboot scripts have finished running.

- Use one of the following options to check the statuses:

- To see a report of health check status, enter:

```
$ systemctl show --property=SubState --value greenboot-healthcheck.service
```

The following outputs are possible:

- **start** means that greenboot checks are still running.
- **exited** means that checks have passed and greenboot has exited. Greenboot runs the scripts in the **green.d** directory when the system is in a healthy state.
- **failed** means that checks have not passed. Greenboot runs the scripts in **red.d** directory when the system is in this state and might restart the system.
- To see a report showing the numerical exit code of the service, where **0** means success and nonzero values mean a failure occurred, use the following command:

```
$ systemctl show --property=ExecMainStatus --value greenboot-healthcheck.service
```

- To see a report showing a message about boot status, such as **Boot Status is GREEN - Health Check SUCCESS**, enter:

```
$ cat /run/motd.d/boot-status
```

10.4.5. Manually rolling back RHEL for Edge images

When you upgrade your operating system, a new deployment is created, and the **rpm-ostree** package also keeps the previous deployment. If there are issues on the updated version of the operating system, you can manually roll back to the previous deployment with a single **rpm-ostree** command, or by selecting the previous deployment in the GRUB boot loader.

To manually roll back to a previous version, perform the following steps.

Prerequisite

1. You updated your system and it failed.

Procedure

1. Optional: Check for the fail error message:

```
$ journalctl -u greenboot-healthcheck.service.
```

2. Run the **rollback** command:

```
# rpm-ostree rollback
```

The command output provides details about the commit ID that is being moved and indicates a completed transaction with the details of the package being removed.

3. Reboot the system.

```
# systemctl reboot
```

The command activates the previous commit with the stable content. The changes are applied and the previous version is restored.

10.4.6. Rolling back RHEL for Edge images using an automated process

Greenboot checks provides a framework that is integrated into the boot process and can trigger **rpm-ostree** rollbacks when a health check fails. For the health checks, you can create a custom script that indicates whether a health check passed or failed. Based on the result, you can decide when a rollback should be triggered. The following procedure shows how to create an health check script example:

Procedure

1. Create a script that returns a standard exit code **0**.
For example, the following script ensures that the configured DNS server is available:

```
#!/bin/bash

DNS_SERVER=$(grep ^nameserver /etc/resolv.conf | head -n 1 | cut -f2 -d" ")
COUNT=0
# check DNS server is available
ping -c1 $DNS_SERVER
while [ $? != '0' ] && [ $COUNT -lt 10 ]; do
  ((COUNT++))
  echo "Checking for DNS: Attempt $COUNT ."
  sleep 10
  ping -c 1 $DNS_SERVER
done
```

2. Include an executable file for the health checks at **/etc/greenboot/check/required.d/**.

```
chmod +x check-dns.sh
```

During the next reboot, the script is executed as part of the boot process, before the system enters the **boot-complete.target** unit. If the health checks are successful, no action is taken. If the health checks fail, the system will reboot several times, before marking the update as failed and rolling back to the previous update.

Verification steps

To check if the default gateway is reachable, run the following health check script:

1. Create a script that returns a standard exit code **0**.

```
#!/bin/bash

DEF_GW=$(ip r | awk '/^default/ {print $3}')
SCRIPT=$(basename $0)
```

```
count=10
connected=0
ping_timeout=5
interval=5

while [ $count -gt 0 -a $connected -eq 0 ]; do
    echo "$SCRIPT: Pinging default gateway $DEF_GW"
    ping -c 1 -q -W $ping_timeout $DEF_GW > /dev/null 2>&1 && connected=1 || sleep
    $interval
    ((--count))
done

if [ $connected -eq 1 ]; then
    echo "$SCRIPT: Default gateway $DEF_GW is reachable."
    exit 0
else
    echo "$SCRIPT: Failed to ping default gateway $DEF_GW!" 1>&2
    exit 1
fi
```

2. Include an executable file for the health checks at **/etc/greenboot/check/required.d/** directory.

```
chmod +x check-gw.sh
```

CHAPTER 11. CREATING AND MANAGING OSTREE IMAGE UPDATES

You can easily create and manage OSTree image updates for your RHEL for Edge systems and make them immediately available to RHEL for Edge devices. With OSTree, you can use image builder to create RHEL for Edge Commit or RHEL for Edge Container images as **.tar** files that contain OSTree commits. The OSTree update versioning system works as a “Git repository” that stores and versions the OSTree commits. The **rpm-ostree** image and package system then assembles the commits on the client device. When you create a new image with RHEL image builder to perform an update, RHEL image builder pulls updates from these repositories.

11.1. BASIC CONCEPTS FOR OSTREE

Basic terms that OSTree and **rpm-ostree** use during image updates.

rpm-ostree

The technology on the edge device that handles how the OSTree commits are assembled on the device. It works as a hybrid between an image and a package system. With the **rpm-ostree** technology, you can make atomic upgrades and rollbacks to your system.

OSTree

OSTree is a technology that enables you to create commits and download bootable file system trees. You can also use it to deploy the trees and manage the boot loader configuration.

Commit

An OSTree commit contains a full operating system that is not directly bootable. To boot the system, you must deploy it, for example, with a RHEL Installable image.

Reference

It is also known as **ref**. An OSTree ref is similar to a Git branch and it is a name. The following reference names examples are valid:

- **rhel/9/x86_64/edge**
- **ref-name**
- **app/org.gnome.Calculator/x86_64/stable**
- **ref-name-2**

By default, RHEL image builder specifies **rhel/8/\$ARCH/edge** as a path. The “\$ARCH” value is determined by the host machine.

Parent

The **parent** argument is an OSTree commit to build a new commit with RHEL image builder. It retrieves a parent commit for the new commit that you are building. You can specify the parent commit as a **ref** value to be resolved and pulled, for example **rhel/8/x86_64/edge**. You can also use the **Commit ID** that you can find in the extracted **.tar** image file.

Remote

The http or https endpoint that hosts the OSTree content. This is analogous to the baseurl for a yum repository.

Static delta

Static deltas are a collection of updates generated between two OSTree commits. This enables the

system client to fetch a smaller amount of files, which are larger in size. The static deltas updates are more network efficient because, when updating an ostree-based host, the system client will only fetch the objects from the new OSTree commit which do not exist on the system. Typically, the new OSTree commit contains many small files, which requires multiple TCP connections.

Summary

The summary file is a concise way of enumerating refs, checksums, and available static deltas in an OSTree repo. You can check the state of all the refs and static deltas available in an Ostree repo. However, you must generate the summary file every time a new ref, commit, or static-delta is added to the OSTree repo.

11.2. CREATING OSTREE REPOSITORIES

You can create OSTree repos with RHEL image builder by using either **RHEL for Edge Commit (.tar)** or **RHEL for Edge Container (.tar)** image types. These image types contain an OSTree repo that contains a single OSTree commit.

- You can extract the **RHEL for Edge Commit (.tar)** on a web server and it is ready to be served.
- You must import the **RHEL for Edge Container (.tar)** to a local container image storage or push the image to a container registry. After you start the container, it serves the commit over an integrated **nginx** web server.

Use the **RHEL for Edge Container (.tar)** on a RHEL server with Podman to create an OSTree repo:

Prerequisite

- You created a **RHEL for Edge Container (.tar)** image.

Procedure

1. Download the container image from RHEL image builder:

```
$ composer-cli compose image <UUI>
```

2. Import the container into Podman:

```
$ skopeo copy oci-archive:<UUI>-container.tar containers-storage:localhost/ostree
```

3. Start the container and make it available by using the port **8080**:

```
$ podman run -rm -p 8080:8080 ostree
```

Verification

- Check that the container is running:

```
$ podman ps -a
```

11.3. MANAGING A CENTRALIZED OSTREE MIRROR

For production environments, having a central OSTree mirror that serves all the commits has several advantages, including:

- Deduplicating and minimizing disk storage
- Optimizing the updates to clients by using static delta updates
- Pointing to a single OSTree mirror for their deployment life.

To manage a centralized OSTree mirror, you must pull each commit from RHEL image builder into the centralized repository where it will be available to your users.



NOTE

You can also automate managing an OSTree mirror by using the **osbuild.infra** Ansible collection. See [osbuild.infra Ansible](#).

To create a centralized repository you can run the following commands directly on a web server:

Procedure

1. Create an empty blueprint, customizing it to use "rhel-93" as the distro:

```
name = "minimal-rhel93"
description = "minimal blueprint for ostree commit"
version = "1.0.0"
modules = []
groups = []
distro = "rhel-93"
```

2. Push the blueprint to the server:

```
# composer-cli blueprints push minimal-rhel93.toml
```

3. Build a RHEL for Edge Commit (**.tar**) image from the blueprint you created:

```
# composer-cli compose start-ostree minimal-rhel93 edge-commit
```

4. Retrieve the **.tar file** and decompress it to the disk:

```
# composer-cli compose image _<rhel-93-uuid>
$ tar -xf <rhel-93-uuid>.tar -C /usr/share/nginx/html/
```

The **/usr/share/nginx/html/repo** location on disk will become the single OSTree repo for all refs and commits.

5. Create another empty blueprint, customizing it to use "rhel-89" as the distro:

```
name = "minimal-rhel89"
description = "minimal blueprint for ostree commit"
version = "1.0.0"
modules = []
groups = []
distro = "rhel-89"
```

6. Push the blueprint and create another RHEL for Edge Commit (**.tar**) image:

```
# *composer-cli blueprints push minimal-rhel89.toml*
# *composer-cli compose start-ostree minimal-rhel89 edge-commit*
```

1. Retrieve the **.tar file** and decompress it to the disk:

```
# composer-cli compose image <rhel-89-uuid>
$ tar -xf <rhel-89-uuid>.tar
```

2. Pull the commit to the local repo. By using **ostree pull-local**, you can copy the commit data from one local repo to another local repo.

```
# ostree --repo=/usr/share/nginx/html/repo pull-local repo
```

3. Optional: Inspect the status of the OSTree repo. The following is an output example:

```
$ ostree --repo=/usr/share/nginx/html/repo refs

rhel/8/x86_64/edge
rhel/9/x86_64/edge

$ ostree --repo=/usr/share/nginx/html/repo show rhel/8/x86_64/edge
commit f7d4d95465fbd875f6358141f39d0c573df6a321627bafde68c73850667e5443
ContentChecksum:
41bf2f8b442a770e9bf03e096a46a286f5836e0a0702b7c3516ef4e0acec2dea
Date: 2023-09-15 16:17:04 +0000
Version: 8.9
(no subject)

$ ostree --repo=/usr/share/nginx/html/repo show rhel/9/x86_64/edge
commit 89290dbfd6f749700c77cbc434c121432defb0c1c367532368eee170d9e53ea9
ContentChecksum:
70235bfb9cae82c53f856183750e809becf0b9b076122b19c40fec92fc6d74c1
Date: 2023-09-15 15:30:24 +0000
Version: 9.3
(no subject)
```

4. Update the RHEL 9.3 blueprint to include a new package and build a new commit, for example:

```
name = "minimal-rhel93"
description = "minimal blueprint for ostree commit"
version = "1.1.0"
modules = []
groups = []
distro = "rhel-93"

[[packages]]
name = "strace"
version = ""
```

5. Push the updated blueprint and create a new RHEL for Edge Commit (**.tar**) image, pointing the compose to the existing OSTree repo:

```
# composer-cli blueprints push minimal-rhel93.toml
# composer-cli compose start-ostree minimal-rhel93 edge-commit --url http://localhost/repo --
ref rhel/9/x86_64/edge
```

6. Retrieve the **.tar** file and decompress it to the disk:

```
# rm -rf repo
# composer-cli compose image <rhel-93-uuid>
# tar -xf <rhel-93-uuid>.tar
```

7. Pull the commit to repo:

```
# ostree --repo=/usr/share/nginx/html/repo pull-local repo
```

8. Optional: Inspect the OSTree repo status again:

```
$ ostree --repo=/usr/share/nginx/html/repo refs
rhel/8/x86_64/edge
rhel/9/x86_64/edge

$ ostree --repo=/usr/share/nginx/html/repo show rhel/8/x86_64/edge
commit f7d4d95465fbd875f6358141f39d0c573df6a321627bafde68c73850667e5443
ContentChecksum:
41bf2f8b442a770e9bf03e096a46a286f5836e0a0702b7c3516ef4e0acec2dea
Date: 2023-09-15 16:17:04 +0000
Version: 8.9
(no subject)

$ ostree --repo=/usr/share/nginx/html/repo show rhel/9/x86_64/edge
commit a35c3b1a9e731622f32396bb1aa84c73b16bd9b9b423e09d72efaca11b0411c9
Parent: 89290dbfd6f749700c77cbc434c121432defb0c1c367532368eee170d9e53ea9
ContentChecksum:
2335930df6551bf7808e49f8b35c45e3aa2a11a6c84d988623fd3f36df42a1f1
Date: 2023-09-15 18:21:31 +0000
Version: 9.3
(no subject)

$ ostree --repo=/usr/share/nginx/html/repo log rhel/9/x86_64/edge
commit a35c3b1a9e731622f32396bb1aa84c73b16bd9b9b423e09d72efaca11b0411c9
Parent: 89290dbfd6f749700c77cbc434c121432defb0c1c367532368eee170d9e53ea9
ContentChecksum:
2335930df6551bf7808e49f8b35c45e3aa2a11a6c84d988623fd3f36df42a1f1
Date: 2023-09-15 18:21:31 +0000
Version: 9.3
(no subject)

commit 89290dbfd6f749700c77cbc434c121432defb0c1c367532368eee170d9e53ea9
ContentChecksum:
70235bfb9cae82c53f856183750e809becf0b9b076122b19c40fec92fc6d74c1
Date: 2023-09-15 15:30:24 +0000
Version: 9.3
(no subject)

$ rpm-ostree db diff --repo=/usr/share/nginx/html/repo
89290dbfd6f749700c77cbc434c121432defb0c1c367532368eee170d9e53ea9
```

```
a35c3b1a9e731622f32396bb1aa84c73b16bd9b9b423e09d72efaca11b0411c9
ostree diff commit from:
89290dbfd6f749700c77cbc434c121432defb0c1c367532368eee170d9e53ea9
ostree diff commit to:
a35c3b1a9e731622f32396bb1aa84c73b16bd9b9b423e09d72efaca11b0411c9
Added:
elfutils-default-yama-scope-0.188-3.el9.noarch
elfutils-libs-0.188-3.el9.x86_64
strace-5.18-2.el9.x86_64
```

APPENDIX A. TERMINOLOGY AND COMMANDS

Learn more about the **rpm ostree** terminology and commands.

A.1. OSTREE AND RPM-OSTREE TERMINOLOGY

Following are some helpful terms that are used in context to OSTree and **rpm-ostree** images.

Table A.1. OSTree and **rpm-ostree** terminology

Term	Definition
OSTree	A tool used for managing Linux-based operating system versions. The OSTree tree view is similar to Git and is based on similar concepts.
rpm-ostree	A hybrid image or system package that hosts operating system updates.
Commit	A release or image version of the operating system. RHEL image builder generates an ostree commit for RHEL for Edge images. You can use these images to install or update RHEL on Edge servers.
Refs	Represents a branch in ostree. Refs always resolve to the latest commit. For example, rhel/8/x86_64/edge .
Revision (Rev)	SHA-256 for a specific commit.
Remote	The http or https endpoint that hosts the ostree content. This is analogous to the baseurl for a yum repository.
static-delta	Updates to ostree images are always delta updates. In case of RHEL for Edge images, the TCP overhead can be higher than expected due to the updates to number of files. To avoid TCP overhead, you can generate static-delta between specific commits, and send the update in a single connection. This optimization helps large deployments with constrained connectivity.

A.2. OSTREE COMMANDS

The following table provides a few OSTree commands that you can use when installing or managing OSTree images.

Table A.2. ostree commands

ostree pull	ostree pull-local --repo [path] src ostree pull-local <path> <rev> --repo=<repo-path> ostree pull <URL> <rev> --repo=<repo-path>
ostree summary	ostree summary -u --repo=<repo-path>
View refs	ostree refs --repo ~/Code/src/osbuild-iot/build/repo/ --list
View commits in repo	ostree log --repo=/home/gicmo/Code/src/osbuild-iot/build/repo/ <REV>
Inspect a commit	ostree show --repo build/repo <REV>
List remotes of a repo	ostree remote list --repo <repo-path>
Resolve a REV	ostree rev-parse --repo ~/Code/src/osbuild-iot/build/repo fedora/x86_64/osbuild-demo ostree rev-parse --repo ~/Code/src/osbuild-iot/build/repo b3a008eceeddd0cfd
Create static-delta	ostree static-delta generate --repo=[path] --from=REV --to=REV
Sign an existing ostree commit with a GPG key	ostree gpg-sign --repo=<repo-path> --gpg-homedir <gpg_home> COMMIT KEY-ID...

A.3. RPM-OSTREE COMMANDS

The following table provides a few **rpm-ostree** commands that you can use when installing or managing OSTree images.

Table A.3. rpm-ostree commands

Commands	Description
rpm-ostree --repo=/home/gicmo/Code/src/osbuild-iot/build/repo/ db list <REV>	This command lists the packages existing in the <REV> commit into the repository.

Commands	Description
rpm-ostree rollback	OSTree manages an ordered list of boot loader entries, called deployments . The entry at index 0 is the default boot loader entry. Each entry has a separate /etc directory, but all the entries share a single /var directory. You can use the boot loader to choose between entries by pressing Tab to interrupt startup. This rolls back to the previous state, that is, the default deployment changes places with the non-default one.
rpm-ostree status	This command gives information about the current deployment in use. Lists the names and refspecs of all possible deployments in order, such that the first deployment in the list is the default upon boot. The deployment marked with * is the current booted deployment, and marking with 'r' indicates the most recent upgrade.
rpm-ostree db list	Use this command to see which packages are within the commit or commits. You must specify at least one commit, but more than one or a range of commits also work.
rpm-ostree db diff	Use this command to show how the packages are different between the trees in two revs (revisions). If no revs are provided, the booted commit is compared to the pending commit. If only a single rev is provided, the booted commit is compared to that rev.
rpm-ostree upgrade	This command downloads the latest version of the current tree, and deploys it, setting up the current tree as the default for the next boot. This has no effect on your running filesystem tree. You must reboot for any changes to take effect.

Additional resources

- The **rpm-ostree** man page.

A.4. FDO AUTOMATIC ONBOARDING TERMINOLOGY

Learn more about the FDO terminology.

Table A.4. FDO terminology

Commands	Description
FDO	FIDO Device Onboarding.
Device	Any hardware, device, or computer.
Owner	The final owner of the device – a company or an IT department.

Commands	Description
Manufacturer	The device manufacturer.
Manufacturer server	Creates the device credentials for the device.
Manufacturer client	Informs the location of the manufacturing server.
Ownership Voucher (OV)	Record of ownership of an individual device. Contains the following information: * Owner (fdo-owner-onboarding-service) * Rendezvous Server - FIDO server (fdo-rendezvous-server) * Device (at least one combination) (fdo-manufacturing-service)
Device Credential (DC)	Key credential and rendezvous stored in the device at manufacture.
Keys	Keys to configure the manufacturing server * key_path * cert_path * key_type * mfg_string_type: device serial number * allowed_key_storage_types: Filesystem and Trusted Platform Module (TPM) that protects the data used to authenticate the device you are using.
Rendezvous server	Link to a server used by the device and later on, used on the process to find out who is the owner of the device

Additional resources

- [FIDO IoT spec](#)

A.5. FDO AUTOMATIC ONBOARDING TECHNOLOGIES

Following are the technologies used in context to FDO automatic onboarding.

Table A.5. OSTree and rpm-ostree terminology

Technology	Definition
UEFI	Unified Extensible Firmware Interface.
RHEL	Red Hat® Enterprise Linux® operating system
rpm-ostree	Background image-based upgrades.
Greenboot	Healthcheck framework for systemd on rpm-ostree .
Osbuild	Pipeline-based build system for operating system artifacts.
Container	A Linux® container is a set of 1 or more processes that are isolated from the rest of the system.
Coreos-installer	Assists installation of RHEL images, boots systems with UEFI.
FIDO FDO	Specification protocol to provision configuration and onboarding devices.