

Oracle Linux

Using the Cockpit Web Console



F51970-05
October 2023



Oracle Linux Using the Cockpit Web Console,

F51970-05

Copyright © 2022, 2023, Oracle and/or its affiliates.

Contents

Preface

Conventions	vii
Documentation Accessibility	vii
Access to Oracle Support for Accessibility	vii
Diversity and Inclusion	viii

1 Get Started: Cockpit Web Console

Cockpit Web Console Overview	1-1
Installation and Log In	1-3
Install and Enable Cockpit	1-3
Log in to the Cockpit Web Console	1-5
Cockpit Login Banner	1-7
Management of Multiple Hosts	1-8
Security Considerations for Multiple Host Management	1-8
Add and Connect to Secondary Host	1-9
Edit or Remove Secondary Host	1-11
Extend Cockpit's Functionality	1-12
Install and Manage Add-on Applications	1-12

2 Common Host Administration Tasks

General Host Configuration Actions	2-1
Change System Hostname	2-1
Change System Date and Time	2-2
Change Tuned Performance Profile	2-3
Restart or Power Off System	2-4
Join an Active Directory Domain	2-5
Security Management Actions	2-6
Disable SMT to Prevent CPU Security Issues	2-6
Change System-Wide Crypto Policy	2-7
Set a Timeout for Inactive Sessions	2-8
Manage SELinux Security Policy and Messages	2-9

Change SELinux Policy Mode	2-9
View SELinux Modifications and Copy Automation Script	2-10
Manage SELinux Alert Message	2-11
User Management Actions	2-11
Create or Change User Accounts	2-12
Disconnect User Sessions or Remove User Accounts	2-13
System Monitoring Actions	2-14
Health, Usage, and System Details	2-14
View and Manage Services	2-15
View and Filter Log Entries	2-17
Software Updates	2-19
Apply Pending Software Updates Manually	2-20
Schedule Automatic Software Updates	2-21

3 Network Management Tasks

Manage Firewall Zoning Properties	3-1
Change Host Firewall State	3-1
Display Firewall Zone Properties	3-2
Control Access to Zone Services	3-4
Add a New Predefined Zone	3-5
Remove an Existing Predefined Zone	3-7
Monitor or Change Interface Connections	3-7
Configure Network Bonding Properties	3-9
Configure Network Teaming Properties	3-12
Teaming and Bonding Feature Comparison	3-14
Configure Network Bridging Properties	3-15
Configure Network VLAN Properties	3-17
View Network Host Log Files	3-19

4 Storage Management Tasks

Storage Management Installation and Overview	4-1
Manage Disk Devices and Partitions	4-3
Create Physical Disk Partitions	4-3
Storage Partitioning Considerations and Prerequisites	4-6
View and Change Drive Partition Properties	4-7
Encrypt Block Devices With LUKS	4-11
Lock Disk Device With LUKs	4-11
Change Passphrase Key for LUKS Encryption	4-13
Unlock Encrypted Devices Using Tang Server Key	4-15

Create a Tang Key for Encrypted Device	4-16
Confirm Tang Key Implementation on Encrypted Device	4-18
Manage Logical Volumes With LVM	4-19
Create a Volume Group	4-20
Create a Logical Volume	4-21
Create a Thin Logical Volume	4-22
Format and Mount a Logical Volume	4-24
Resize Logical Volumes	4-25
Change Volume Group Properties	4-27
Add New Drive to a Volume Group	4-27
Remove Physical Drive From Volume Group	4-28
Rename a Volume Group	4-29
Remove Volume Group	4-30
Remove Logical Volume From Volume Group	4-31
Build and Manage Software RAID Devices	4-32
Create and Configure a New Storage Array	4-32
Software RAID Levels	4-35
Manage NFS Mounted Connections	4-36
Add NFS Server Connections	4-36
Change NFS Server Connections	4-37
Manage Connections to iSCSI Targets	4-38

5 Virtual Machine (VM) Management Tasks

Create, Import, Clone, or Migrate a VM Instance	5-1
Create a VM Instance	5-1
Install Cockpit-Machines and Enable Virtualization	5-7
Video Demonstration	5-8
Clone an Existing VM Instance	5-8
Import a VM From a Disk Image	5-9
Migrate a Running VM to Another KVM Host	5-11
Requirements to Migrate Virtual Machine	5-12
Manage Storage Pools for VM Instances	5-13
Create a Storage Pool	5-13
Manage Existing Storage Pools	5-15
Start, Shutdown, Remove, or Interrupt a VM Instance	5-15
Configure Console for VM Interaction	5-17
Configure VM Devices and Services	5-20
Edit Memory, CPU, Autostart, or Watchdog Properties	5-20
Add, Edit, or Remove Disks	5-23
Attach or Remove VM Host Devices	5-25

Add, Edit, Unplug, or Plug VM Network Interface	5-27
Attach or Remove VM Host Devices	5-29
Create, Delete, or Revert a Snapshot of VM Instance	5-31
Share a Host Directory with VM Instance	5-33

6 Podman Management Tasks

Install and Configure Cockpit-Podman	6-1
Podman Image Management	6-2
Search and Download New Images	6-2
View and Inspect Available Images	6-4
Remove Images	6-6
Podman Container Management	6-7
Create and Run Container	6-8
Special Considerations for Non Administrator Containers	6-12
Inspect Container and Access Container Logs and CLI	6-13
Rename, Pause, Stop, or Restart Container	6-14
Commit Container Changes to Create New Image	6-14
Checkpoint and Restore a System Container	6-16
Remove Container or Pod Group	6-17
Podman Pod Management	6-18
Create a Pod Group	6-19
Add a Container to a Pod Group	6-21
Inspect and Change a Pod Group	6-22

7 Debugging Tools

Evaluate System Problems Using Diagnostic Reports	7-1
Generate Diagnostic Reports	7-1
Download or Remove Generated Reports	7-3
Capture Crash Dump Details Using Kdump	7-3
Overview of Crash Recovery in Kdump	7-4
Change Kdump Service State	7-4
Change Crash Dump Target Location	7-5
Test the Kdump Configuration	7-7

Preface

Oracle Linux includes a web console you can use for system administration. The web console is called Cockpit. For non-minimal installations, Cockpit is automatically installed, although not automatically enabled. Cockpit provides a web browser interface for performing system configuration and administration tasks, either locally or remotely on multiple servers. These tasks include system resource monitoring and log review, network and firewall configuration, and package management and updates. Cockpit uses the same APIs to access system services, so any changes you make using operating system command line tools are updated in real time in Cockpit.

Conventions

The following text conventions are used in this document:

Convention	Meaning
boldface	Boldface type indicates graphical user interface elements associated with an action, or terms defined in text or the glossary.
<i>italic</i>	Italic type indicates book titles, emphasis, or placeholder variables for which you supply particular values.
monospace	Monospace type indicates commands within a paragraph, URLs, code in examples, text that appears on the screen, or text that you enter.

Documentation Accessibility

For information about Oracle's commitment to accessibility, visit the Oracle Accessibility Program website at <https://www.oracle.com/corporate/accessibility/>.

For information about the accessibility of the Oracle Help Center, see the Oracle Accessibility Conformance Report at <https://www.oracle.com/corporate/accessibility/templates/t2-11535.html>.

Access to Oracle Support for Accessibility

Oracle customers that have purchased support have access to electronic support through My Oracle Support. For information, visit <https://www.oracle.com/corporate/accessibility/learning-support.html#support-tab>.

Diversity and Inclusion

Oracle is fully committed to diversity and inclusion. Oracle respects and values having a diverse workforce that increases thought leadership and innovation. As part of our initiative to build a more inclusive culture that positively impacts our employees, customers, and partners, we are working to remove insensitive terms from our products and documentation. We are also mindful of the necessity to maintain compatibility with our customers' existing technologies and the need to ensure continuity of service as Oracle's offerings and industry standards evolve. Because of these technical constraints, our effort to remove insensitive terms is ongoing and will take time and external cooperation.

1

Get Started: Cockpit Web Console

To get started with using the Cockpit web console, administrators should review the following topics to become familiar with Cockpit features. Instructional information covered in these topics include steps to install Cockpit, as well as information to optionally create a login banner, a multi-host setup, or add additional functionality to the Cockpit web interface.

- [Cockpit Web Console Overview](#)
- [Installation and Log In](#)
- [Cockpit Login Banner](#)
- [Management of Multiple Hosts](#)
- [Extend Cockpit's Functionality](#)

Cockpit Web Console Overview

The Cockpit web console is an Oracle Linux server administration tool designed for managing and monitoring Linux systems both locally and remotely. The graphical web console for management and administration is user-friendly and accessible to administrators at all levels of experience with Linux. The Cockpit installation package arrives out of the box and ready for use in most Oracle Linux distributions.

Intuitive Server Administration

Choosing Cockpit as an Oracle Linux server administration tool can help make server administration with Oracle Linux more discoverable and intuitive. For example, the web console interface doesn't require you to remember commands to input at a command line. Although, you can still access and use the command line, it's often easier to complete those same administration tasks from the web console interface.

Standard and Add-On Server Administration Tools

System administrators can conveniently administer different areas of a host server by clicking the Cockpit navigational menu in the left side panel of the web console. This menu, by default, presents a standard set of system administration tools that are ready for use. Optionally, other add-on administration tools can be added to the web console as needed.

A list of the system administration menu options that are typically included in the Cockpit web console upon first time use are as follows:

Table 1-1 Administration Tools - Menu Navigation:

Use:	To:	Details:
Host	<ul style="list-style-type: none">• View the system hostname of the connected Oracle Linux server.• Add secondary host systems to manage from a primary single web console instance.	Management of Multiple Hosts

Table 1-1 (Cont.) Administration Tools - Menu Navigation:

Use:	To:	Details:
Overview	<ul style="list-style-type: none"> Examine host system health, usage, hardware, and general configuration details. Set the system hostname and time. View and change the generated Secure Shell fingerprint. Set system recommended performance profile settings. View system usage statistics and graphs for CPU and Memory. View the host system hardware configuration details. 	Health, Usage, and System Details
Logs	<ul style="list-style-type: none"> View and filter log entries collected on the host system. Pause or resume the log view entries captured by the system. 	View and Filter Log Entries
Networking	<ul style="list-style-type: none"> View the host system firewall setup and network interfaces. Set firewall rules and zones. Set properties to configure network bonding, teaming, bridges, and VLANs 	Network Management Tasks
Accounts	<ul style="list-style-type: none"> Manage host user access. Set properties to create, edit, lock, or suspend user access. 	User Management Actions
Services	<ul style="list-style-type: none"> Administer host System and User processes, targets, sockets, timers, and paths. Find System or User processes by filtering the view with criteria of interest. For example, filter by description, name, Active State, or File state. Change the state of a service by selecting a different action (for example, start, stop, reload, activate, or deactivate). Click Relationships to view a service's dependencies with other system services. View log entries collected for a specific service. 	View and Manage Services
Diagnostic Reports	<ul style="list-style-type: none"> Generate a diagnostic report to gain insight on the overall operation of the host system. Download a report, delete a report, or generate other reports. 	Evaluate System Problems Using Diagnostic Reports

Table 1-1 (Cont.) Administration Tools - Menu Navigation:

Use:	To:	Details:
Kernel Dump	<ul style="list-style-type: none"> View and manage the properties of the kernel crash dump configuration. Test the kernel crash dump configuration by intentionally crashing the kernel on the host system. 	Capture Crash Dump Details Using Kdump
SELinux	<ul style="list-style-type: none"> View and manage properties of the SELinux policy configuration. View recent SELinux policy modifications applied to the host system. View and manage access control policy errors generated on the host Linux system. 	Manage SELinux Security Policy and Messages
Terminal	<ul style="list-style-type: none"> Directly access and use the host system Terminal command line user interface. Sometimes, the completion of administration tasks requires the use of the CLI. 	To access the host CLI, click Terminal In the Cockpit navigation pane.

Add-On Server Administration Applications

System administrators can optionally extend the Cockpit web console functionality by installing extra administration tool packages that were developed for Cockpit. For more details, see [Extend Cockpit's Functionality](#).

Installation and Log In

The following topics step you through the process of installing the Cockpit package and logging you in to the Cockpit web console for the first time.

- [Install and Enable Cockpit](#)
- [Log in to the Cockpit Web Console](#)

Install and Enable Cockpit

The `cockpit` package is, by default, included in all non-minimal Oracle Linux software installations.

What Do You Need?

- A system with Oracle Linux installed.
- Root administrator privileges on the host Linux system.

 **Note:**

The steps in this procedure uses the `sudo` command to execute commands as the root user.

Steps

Using the Cockpit **Terminal** window, perform the following steps to install and enable the `cockpit` package on an Oracle Linux machine.

1. To verify the version of the `cockpit` package that's available for installation, type:

```
sudo dnf info cockpit
```

The versioning information and other information describing the `cockpit` package appears in the command line output.

2. To install the Cockpit package on the Oracle Linux host system, type:

```
sudo dnf install cockpit
```

3. To enable the `systemd` socket service and automatically start it upon a system restart, type:

```
sudo systemctl enable --now cockpit.socket
```

The socket web service, by default, is configured to accept connections on TCP port 9090.

4. To verify that the socket web service is enabled for Cockpit, type:

```
sudo systemctl status cockpit.socket
```

An **Active** status and the default **listening** port of 9090 appears in the command line output to indicate the service is enabled.

5. (Optional) If a firewall is enabled, perform the following to enable the Cockpit service to receive inbound connections.

- a. Open the firewall for the `cockpit` service as follows:

```
sudo firewall-cmd --add-service=cockpit --permanent
```

- b. Apply the firewall configuration change in the runtime environment by reloading the firewall configuration as follows:

```
sudo firewall-cmd --reload
```

 **Note:**

The firewall restricts access to the cockpit system. Remote authorized users can access the system securely by using port forwarding over SSH. For example, the user can type the following command from the user's local system:

```
ssh -L 9090:localhost:9090 user@cockpit-system
```

The *cockpit-system* can be the system's fully qualified domain name (FQDN), such as `myserver.example.com` or the system's IP address. Then the user can log in to `localhost`, as explained in [Log in to the Cockpit Web Console](#).

Log in to the Cockpit Web Console

Cockpit enables you to log in directly to any Oracle Linux system that permits connections to the application using TCP port 9090.

 **Note:**

If the 9090 port isn't accessible on the system, you can still use Cockpit to administer the system by adding it as a secondary host server. For details on how to add a secondary server, see [Management of Multiple Hosts](#).

What Do You Need?

- The IP address or hostname of the Oracle Linux server where Cockpit is installed.
- A valid Oracle Linux user account on the server where Cockpit is installed.

 **Note:**

Cockpit uses PAM for authentication and the configuration is available in `/etc/pam.d/cockpit`. Authentication with PAM enables you to log in with a username and password of any system account that has administrator privileges.

- (Optional) Signed Certificate by Certificate Authority on Oracle Linux host.

 **Note:**

To avoid granting a security exception each time you access the Cockpit web console, install a certificate signed by a certificate authority (CA) in the `/etc/cockpit/ws-certs.d` directory. The last file (in alphabetical order) with a `.cert` extension is used. For more details, see [Managing Certificates and Public Key Infrastructure](#) in the [Oracle Linux: Managing Certificates and Public Key Infrastructure](#).

Steps

Perform the following steps to log in to the Cockpit web console.

1. In a web browser, access the Cockpit web console using the hostname or IP address of the system at port 9090 using HTTPS. For example:

```
https://myserver.example.com:9090
```

If you're logging in on the local host, you can use:

```
https://localhost:9090
```

If you aren't using a signed security certificate, a warning message appears indicating "Your connection is not private". To continue, you can choose to bypass this error by clicking the browser option available (such as **Advanced** or **More details**.) for adding a security exception for this site.

The Cockpit **Login** page appears.

2. In the Cockpit **Login** page, enter your system user name and password.

Oracle Linux Server

User name

Password

[▶ Other options](#)

Log in

Note:

As an alternative, you can remotely connect to another Linux host system on the network by clicking: (1) **Other Options**, (2) **Connect to**, and then entering the URL of the remote host.

3. Click **Log in**.
After successful authentication, the Cockpit web console page appears.

! Important:

The Cockpit web console takes on the privileges and security context of the logged in user. Upon first time log in, **Limited access** mode is enabled by default. To elevate privileges to **Administrative access** mode, click the toggle switch to **Turn on administrative access**. Upon successful authentication of the user password, administrative privileges are granted.

Cockpit Login Banner

Administrators can choose to communicate changes to host users by displaying a banner message on the Cockpit **Login** page. Typically, a banner message is useful for enforcing a security policy or to inform users of scheduled maintenance events.

What Do You Need

- The Cockpit web console must be installed and accessible. For details, see these topics: [Install and Enable Cockpit](#) and [Log in to the Cockpit Web Console](#).
- An existing `cockpit.conf` file must be saved in the `/etc/cockpit/` directory.

! Important:

The Cockpit configuration file (`cockpit.conf`) is not a system required file. In cases where the file does not already exist, the file must be manually created before performing the steps in this procedure. For more details about the use of this file, see the `cockpit.conf(5)` manual page.

- Administrator privileges.

Steps

Using the command line, perform the following steps to add a pre-login banner message to the Cockpit configuration file.

1. Specify the text for the banner message in the `issues.cockpit` file, for example:
 - a. Open or create the `/etc/issue.cockpit` file in a text editor.
 - b. In the `issue.cockpit` file you can edit or add the content you want to display as a pre-login banner message.
Don't include any macros in this file as no reformatting is handled between the file content and the displayed content. Use line breaks and indentation to format the banner content. You can also use ASCII art.
 - c. Save the file.
2. If you haven't done so already, add the banner file properties to the Cockpit configuration file as follows:
 - a. Open the `cockpit.conf` file in the `/etc/cockpit/` directory. For example:

```
sudo vi cockpit.conf
```

- b. In the `cockpit.conf` file, add the following information:

```
[Session]
Banner=/etc/issue.cockpit
```

- c. Save the file.
3. Restart Cockpit for the configuration file changes to take effect.

```
sudo systemctl try-restart cockpit
```

Management of Multiple Hosts

Cockpit enables administrators flexibility when it comes to the number of hosts they can manage from a single Cockpit session. Administrators requiring the configuration of a multi-host setup in Cockpit must ensure that all the proper prerequisites are met first. For details on how to set up a multi-host configuration in Cockpit, see the following topics.

- [Security Considerations for Multiple Host Management](#)
- [Add and Connect to Secondary Host](#)
- [Edit or Remove Secondary Host](#)

Security Considerations for Multiple Host Management

For optimal security, consider implementing the following configurations when accessing and managing multiple host systems from a single Cockpit web console instance.

- **Optimal topology configuration over SSH connection:**
 - Install Cockpit on a bastion host and use it to connect and manage other secondary Cockpit hosts. The Cockpit bastion host should be configured with a certificate-authority-issued TLS certificate.
 - Configure all secondary hosts to communicate over an SSH connection. For example, in this scenario:
 - * All secondary Cockpit hosts are reachable through the SSH protocol (which defaults to port 22).
 - * The SSH firewall port is open on all secondary Cockpit hosts.
 - * Enabling the `cockpit.socket` service on the secondary Cockpit hosts is *not* required.
 - * A certificate-authority-issued TLS certificate isn't required on the secondary Cockpit hosts. However, the primary Cockpit bastion host must be configured with a certificate-authority-issued TLS certificate.

For SSH configuration details, see *Configuring OpenSSH Server* in [Oracle Linux: Connecting to Remote Systems With OpenSSH](#)

 **Note:**

Cockpit Project - Authentication: For additional information when managing a primary and secondary servers using Cockpit, see <https://cockpit-project.org/guide/latest/authentication.html>.

- **Use of SSH for remote host authentication:**
 - **SSH key-based authentication (preferred authentication method)** – Key-based authentication helps to prevent brute force password attacks against SSH and it provides administrators with password-less key-based authentication. If an SSH key-based authentication isn't already set up, it's easily configurable by selecting the Authorize SSH Key check box when logging in to a remote host. For details, see Step 3 in this procedure [Add and Connect to Secondary Host](#).
 - OR–
 - **SSH password authentication** – Password authentication of the SSH client requires entering the user id and password from the host on which the SSH server resides. While SSH password authentication might be convenient for some users, password authentication is discouraged because it can make accounts more susceptible to intrusion.

Add and Connect to Secondary Host

The Cockpit web console enables system administrators to manage multiple Linux host systems all from a single Cockpit web console instance.

What Do You Need?

- For optimal security access requirements, see [Security Considerations for Multiple Host Management](#).

NOT_SUPPORTED:


The Steps appearing in this topic assume that the security access considerations have been met for the primary and secondary Cockpit hosts.

- Cockpit is installed and enabled on a bastion host.
- The primary Cockpit bastion host is configured with a certificate-authority-issued TLS certificate.
- Cockpit is installed on all secondary hosts.
- All secondary Cockpit hosts are configured to support OpenSSH. For SSH configuration details, see *Configuring OpenSSH Server* in [Oracle Linux: Connecting to Remote Systems With OpenSSH](#).
- Established connection with the primary Cockpit (bastion enabled) host.
- Administrator privileges.

Steps

Follow these steps to add and connect a secondary Cockpit host to a primary Cockpit host.

1. In the primary Cockpit web console, do the following:
 - a. In the `username@hostname` section that appears above the navigation pane, expand the hidden **Host** menu by clicking the **down-arrow**.



- b. In the Host menu, select **Add new host**. The **Add new host** dialog appears.
2. In the **Add new host** dialog, configure the following properties and click **Add**.

Host	Enter a hostname or IP address of the secondary host that you want to add.
User Name	Enter the username configured on the secondary host.
Color	Select a top-border color to distinguish this secondary host from other secondary hosts.

3. If this is a first-time SSH connection to the newly added secondary host, one of the following dialogs appears.
 - **New Host** dialog - If this dialog appears, follow the instructions on the dialog to confirm the SSH key fingerprint, then click **Accept Key and Connect**.
 - **Log in to [host]** dialog. If this dialog appears, configure the following properties and click **Log in**.

 **Note:**

If SSH key authentication isn't already set up, the **Log in to: [host]** dialog appears.

Password	Enter the user password on configured on the secondary host. Note that when the Automatic Login option isn't configured, future authentication for the secondary host prompts for a user password.
-----------------	--

Automatic Login

(Optional) Authorize SSH Key
Perform these steps:

- a. Select the **Automatic Login** checkbox to configure SSH key authentication on the secondary host.
- b. In the **Key password** text box and the **Confirm key password** text box, specify the key password.

When **Automatic Login** is enabled, future logins to this secondary host are no longer prompted for a user password.

A connection is established to the secondary host. The name of the secondary host appears on the **Host** menu at the upper left corner on the page.

4. To log in to the newly added secondary host, click the **Host** menu, and then select the name of the secondary host.
A dialog prompting for a user password only appears if the secondary host isn't configured with SSH key authentication.

Edit or Remove Secondary Host

System administrators can use the **Host** menu on the primary Cockpit web console to edit or remove secondary host configurations.

What Do You Need?

- Established connection with the primary Cockpit web console.

 **Note:**

The *primary Cockpit web console* is the management web console that's configured with secondary hosts.

- Administrator privileges.

Steps

Using the primary Cockpit web console, follow these steps to edit or remove a secondary host configuration.

1. In the upper left corner of the web console, click the down arrow next to the primary host name.
The **Host** drop-down menu appears listing the name of each secondary host configuration.
2. In the **Host** drop-down menu, click **Edit hosts** and then perform one of the following:
 - Remove secondary host configuration:
 - a. Click the red [-] minus button next to the secondary host name that you want to remove. A prompt appears to confirm the removal operation, click **Yes** to proceed.
 - b. Click **Stop editing host**.
The secondary host name is removed from the primary Cockpit web console **Host** menu.
 - Edit existing secondary host configuration:

- a. Click the pencil (✎) icon next to the secondary host name that you want to edit.
- b. In the **Edit [hostname]** dialog, make the necessary changes, and click **Save**.
- c. Click **Stop editing host**.

Extend Cockpit's Functionality

To include extra server administration functionality in Cockpit, administrators can optionally install add-on applications that are included in the Oracle Linux distribution. For example:



Note:

Depending on the Oracle Linux version installed, some add-on application packages might not be available for installation.

Table 1-2 Oracle Linux Cockpit Add-on Applications

Application	Install Package	Usage
Applications and Software Updates	<code>cockpit-packagekit</code>	Add-on Application management and software updates (typically installed by default)
Diagnostic Reports	<code>cockpit-sosreport</code>	Reports to help diagnose system problems.
Kernel Dump	<code>cockpit-kdump</code>	Helps to catch stack traces,
Machines	<code>cockpit-machines</code>	Manage libvirt virtual machines
Performance Data (on-demand)	<code>cockpit-pcp</code>	Installed on demand from the Cockpit web console upon clicking View Usage metrics and history.
Podman	<code>cockpit-podman</code>	Manage Podman containers as of Oracle Linux 8.1.
SELinux	<code>cockpit-selinux</code>	View and manage SELinux exceptions.
Storage	<code>cockpit-storaged</code>	Manage host system's storage devices.

Install and Manage Add-on Applications

The Cockpit add-on applications are installable by either issuing a command in the CLI or by clicking a button on the **Applications** page of the Cockpit web console.

Initially, all Cockpit add-on applications packages must be installed through the CLI. After an add-on application is installed through the CLI, the add-on application is then manageable from the Cockpit **Applications** page by clicking **Remove** or **Install**.

What Do You Need?

- The Cockpit web console must be installed and accessible. For details, see these topics: [Install and Enable Cockpit](#) and [Log in to the Cockpit Web Console](#).
- Administrator privileges.
- Names of the add-on application packages to install. For a list of package names, see [Extend Cockpit's Functionality](#).

Steps

Follow these steps to install and manage Cockpit add-on applications in the web console.

1. Perform one of the following:

- **Install new add-on Cockpit applications.** In the Cockpit navigation pane, click **Terminal** and then use the following command syntax to install add-on application packages:

```
sudo dnf install [packagename]
```

For example, to install the `cockpit-packagekit` package, type:

```
sudo dnf install cockpit-packagekit
```

Note:

In some instances, the `cockpit-packagekit` add-on application might already be installed and available for use in the web console.

- **Manage existing add-on applications in Cockpit.** In the Cockpit navigation pane, click **Applications** to manage all existing add-on applications. For example, on the **Applications** page, you can choose to perform any of the following tasks:
 - View the add-on application names already installed.
 - Click **Remove** to remove a selected application from Cockpit.
 - Click **Install** to re-install a selected application that was previously removed from Cockpit.

2

Common Host Administration Tasks

Common host administration tasks that are often performed from the command line are easily performed from the Cockpit web console. These common tasks involve basic host configuration, security management, user management, system monitoring actions, and software updates. For more information about performing these tasks from the Cockpit web console, see the following topics:

- [General Host Configuration Actions](#)
- [Security Management Actions](#)
- [User Management Actions](#)
- [System Monitoring Actions](#)
- [Software Updates](#)

General Host Configuration Actions

Using the Configuration panel on the Overview page, Cockpit administrators can conveniently access and perform any of the following host configuration actions.

- [Change System Hostname](#)
- [Change System Date and Time](#)
- [Change Tuned Performance Profile](#)
- [Restart or Power Off System](#)
- [Join an Active Directory Domain](#)

Change System Hostname

When you access the Cockpit web console the hostname assigned to the system appears in the **Configuration** section of the **Overview** page. Cockpit administrators can choose to change the properties assigned to the **Real hostname** (Static or Transient) or the **Pretty hostname** (free-form UTF8) as needed.



Note:

A transient hostname, such as `localhost` appears by default whenever a Static or Pretty hostname isn't configured.

What Do You Need?

- The Cockpit web console must be installed and accessible.
For details, see these topics: [Install and Enable Cockpit](#) and [Log in to the Cockpit Web Console](#).
- Administrator privileges.

Steps

Using the Cockpit web console, follow these steps to change the existing hostname assigned to the host system.

1. In the Cockpit navigation pane, click **Overview**.
The **Overview** page appears.
2. In the **Overview** page, navigate to the Configuration section and find the **Hostname** field, then click **Edit**.
The **Change host name** dialog appears.
3. In the **Change host name** dialog, change either the **Pretty host name** field or the **Real host name** field, and click **Change**.

Change System Date and Time

When you access the Cockpit web console, the host system date and time appears in the **Configuration** section of the **Overview** page. As needed, Cockpit administrators can use the configurable properties on the **Overview** page to change the date, time, and timezone options set on the host system.

What Do You Need?

- The Cockpit web console must be installed and accessible on a host machine. For details, see these topics: [Install and Enable Cockpit](#) and [Log in to the Cockpit Web Console](#).
- Administrator privileges.

Steps

Using the Cockpit web console, follow these steps to change the date, time, and time zone properties set on a host system.

1. In the Cockpit navigation pane, click **Overview**.
The **Overview** pane appears.
2. In the **Overview** page, navigate to the **Configuration** section and find the System time property, then click the current date and time.
The **Change system time** dialog appears.
3. In the **Change system time** dialog, perform the following:
 - a. Change the following properties as needed:
 - **Time zone** – Set the appropriate timezone by region and city location.



Tip:

To navigate through the list, type the first few letters of the closest city.

- **Set time** – Select one of the following options to set the system time.

! Important:

Some options might not be available for configuration on all host systems.

- **Manually** – When selected, configurable properties appear to set the exact time and date specific to the time zone specified.
- **Automatically using NTP** – When selected, the system uses any available NTP service to obtain the correct time. For example, the Oracle Linux `chrony` NTP service is typically available and configured, by default, to use the `pool.ntp.org` servers.
- **Automatically using specific NTP servers** – When selected, configurable properties appear enabling you to specify a specific NTP server for synchronizing the system clock.

b. Click **Change**.

Change Tuned Performance Profile

Oracle Linux uses the `Tuned` service to monitor connected devices and dynamically tune the system performance according to a selected profile. By default, the `Tuned` service assigns a recommended performance profile. In cases where the performance profile requires modification, Cockpit administrators can use the **Overview** page in the web console to assign a different performance profile.

What Do You Need?

- The Cockpit web console must be installed and accessible. For details, see these topics: [Install and Enable Cockpit](#) and [Log in to the Cockpit Web Console](#).
- The `Tuned` service in Oracle Linux must be started and enabled on the host system. For more information about configuring the `Tuned` service, see *Working With TuneD* in [Oracle Linux 8: Monitoring and Tuning the System](#) or [Oracle Linux 9: Monitoring and Tuning the System](#).
- Administrator privileges.

Steps

Using the Cockpit web console, follow these steps to change the performance profile assigned to the host system.

1. In the **Overview** page, navigate to the **Configuration** panel and find the **Performance profile** property, then click the assigned profile link. The **Change performance profile** dialog appears.
2. In the **Change performance profile** dialog, select one of the following profiles and then click **Change profile** to apply the change.

None

None, when selected, disables the tuned service state.

Accelerator performance	This profile provides the same tuning as the throughput-performance profile and improves the performance of certain accelerators, such as GPUs.
Balanced	This profile provides a generalized performance profile and is considered suitable for systems requiring a compromise between power saving and performance.
Desktop	This profile optimizes performance for desktop environments based on the balanced profile.
hpc-compute	This profile optimizes the performance for high-performance computing (HPC) and is based on the latency-performance profile.
intel-sst	This profile is optimized for system configurations with user-defined Intel Speed Select Technology.
Latency-performance	This profile is optimized for deterministic performance at the cost of increased power consumption.
Network-latency	This profile is optimized for deterministic performance at the cost of increased power consumption and focuses on low latency network performance.
Network-throughput	This profile is optimized for streaming network throughput. This profile is only necessary on older CPUs or 40G+ networks.
Power-save	This profile is optimized for low power consumption.
Throughput-performance	This profile is optimized for increased performance across various common server workloads.
Virtual-guest	This profile is optimized when running virtual guests based on the throughput-performance profile.
Virtual-host	This profile is optimized for maximum performance when running KVM host.

Restart or Power Off System

Using the **Overview** page in the web console, Cockpit administrators can restart or power off the host, set a time delay, and send a message to warn users to save their changes prior to signing out.

What Do You Need?

- The Cockpit web console must be installed and accessible.
For details, see these topics: [Install and Enable Cockpit](#) and [Log in to the Cockpit Web Console](#).
- Administrator privileges.

Steps

Using the Cockpit web console, follow these steps to perform actions related to powering off or restarting a managed host system.

1. In the Cockpit navigation pane, click **Overview**.
2. In the **Overview** page, click the **Reboot** arrow and then select one of the following:
 - **Reboot** – To restart the host system.
 - **Shutdown** – To power off the host system.
3. In the **Reboot** or **Shutdown** dialog box, perform the following:
 - a. (Optional) In the text box, type a brief message warning users to save their work and sign out.
 - b. (Optional) Click the **Delay** arrow and select a time interval for delaying the action (reboot or shutdown).
 - c. Click **Reboot** or **Shutdown**.

Join an Active Directory Domain

For system environments where Active Directory authentication is already set up, Cockpit administrators can use the **Overview** page in the web console to join an existing Active Directory domain service.



Note:

The **Join a domain** process is similar to using the `realm join` command from the command line. More information is available on the `realm(8)` manual page.

What Do You Need?

- The Cockpit web console must be installed and accessible.
For details, see these topics: [Install and Enable Cockpit](#) and [Log in to the Cockpit Web Console](#).
- An Active Directory service must already be set up and running on the server instance.
- Administrator privileges.

Steps

Using the Cockpit web console, follow these steps to join the host system to an existing Active Directory domain service.

1. In the Cockpit navigation pane, click **Overview**.
2. In the **Overview** page, navigate to the **Configuration** panel, find the **Domain** property, and then click **Join domain**.
The **Join a domain** dialog appears.
3. In the **Join a domain** dialog, perform the following:
 - a. Enter the **Domain address**, **Domain administrator**, and the **Domain administrator password**.
 - b. Click **Join**.

Security Management Actions

To help increase system security, Cockpit administrators can perform the following security management actions from the web console.

- [Disable SMT to Prevent CPU Security Issues](#)
- [Change System-Wide Crypto Policy](#)
- [Set a Timeout for Inactive Sessions](#)
- [Manage SELinux Security Policy and Messages](#)

Disable SMT to Prevent CPU Security Issues

For host systems supporting CPU SMT, system administrators should consider disabling the use of this configuration to prevent system security vulnerabilities. The CPU SMT configuration is typically enabled by default to enhance CPU workload performance.

For more details relating to the CPU SMT usage notice, see [Oracle Linux: Simultaneous Multithreading Notice](#).

What Do You Need?

- The Cockpit web console must be installed and accessible.
For details, see these topics: [Install and Enable Cockpit](#) and [Log in to the Cockpit Web Console](#).
- Administrator privileges.

Steps

 **WARNING:**

Disabling SMT on the host requires restarting the system.

Using the Cockpit web console, follow these steps to disable the SMT configuration on the host system.

1. In the Cockpit navigation pane, click **Overview**.
2. In the **Overview** page, perform the following:
 - a. Navigate to the **System Information** panel and click **View hardware details**.
 - b. In the **Hardware information** page, find the **CPU Security** property, and, if available, click **Mitigations**.

! Important:

For system configurations where CPU SMT isn't available, the Security link for Mitigations doesn't appear. In these instances, the system configuration isn't considered vulnerable to security related attacks because of the misuse of CPU SMT.

The **CPU security toggle** dialog appears.

3. In the **CPU security toggle** dialog, perform the following:
 - a. (Optional) Click **Read** to access further details about CPU SMT configurations.
 - b. Click the toggle button to set the **Disable simultaneous multi-threading (nosmt)** property.
 - c. Click **Save** and **Reboot**.
The host system restarts and disables the CPU use of SMT.

Change System-Wide Crypto Policy

As of Oracle Linux 8 and later, a default system-wide Crypto Policy no longer permits host systems to communicate with older, insecure protocols. For system configurations that require a different level of protection, Cockpit administrators can change the assigned **Crypto Policy** level (Default, Legacy, Future, FIPs) by using the web console.

What Do You Need?

- The Cockpit web console must be installed and accessible.
For details, see these topics: [Install and Enable Cockpit](#) and [Log in to the Cockpit Web Console](#).
- Administrator privileges.

Steps

⚠ WARNING:

Changing the **Crypto Policy** on the host requires restarting the system.

Using the Cockpit web console, follow these steps to change the **Crypto Policy** configuration on the host system.

1. In the Cockpit navigation pane, click **Overview**.
2. In the **Overview** page, navigate to the **Configuration** panel, find the **Crypto policy** property, and then click **Default** (or the policy name that appears).
The **Change Crypto Policy** dialog appears with a brief description of each policy level.
3. In the **Change Crypto Policy** dialog, select a policy level that best meets the requirements of the managed system, and then click **Apply** and **Reboot**.

Set a Timeout for Inactive Sessions

Cockpit, by default, doesn't automatically expire inactive user sessions. To prevent unauthorized use of an unattended session, consider specifying a duration for when inactive user sessions expire.

What Do You Need?

- The Cockpit web console must be installed and accessible. For details, see these topics: [Install and Enable Cockpit](#) and [Log in to the Cockpit Web Console](#).
- An existing `cockpit.conf` file must be saved in the `/etc/cockpit/` directory.

Important:

The Cockpit configuration file (`cockpit.conf`) isn't a system required or provided file. In cases where it doesn't already exist, you must manually create the file before performing the steps in this procedure. For more details about the use of this file, see the `cockpit.conf(5)` manual page.

- Administrator privileges.

Steps

Using the Cockpit Terminal CLI, perform the following steps to configure a session idle timeout duration in the Cockpit configuration file.

1. In the Cockpit navigation pane, click **Terminal**.
2. In the **Terminal** page, open the `cockpit.conf` file in the `/etc/cockpit/` directory in a text editor. For example:

```
sudo vi /etc/cockpit/cockpit.conf
```

3. Add the following information to the `cockpit.conf` file:

```
[Session]
IdleTimeout=value
```

The `IdleTimeout` property specifies how many minutes until an inactive user session in Cockpit is automatically logged out.

4. To set the `IdleTimeout` property, perform one of the following:
 - To **disable** the `IdleTimeout` property, set the `value` to 0. For example:
`IdleTimeout=0`
-OR-
 - To **enable** the `IdleTimeout` property, set the value to reflect the number of minutes allowed before a user session times out. For example, to set an hour of inactivity before a user is logged out, you would set the value to 60. (`IdleTimeout=60`).
5. Save the changes to `cockpit.conf` file.

- Restart Cockpit for the configuration file changes to take effect.

```
sudo systemctl try-restart cockpit
```

Manage SELinux Security Policy and Messages

Security Enhanced Linux (SELinux) provides an additional layer of security on a host system by applying security policies that enable administrators to control who can access the system. The SELinux package is by default included at installation by most Oracle Linux distributions. Upon starting the system, SELinux automatically applies security restrictions to users, programs, processes, and files based on the access permissions defined in the policy.

Administration tasks for setting SELinux policy rules and users are administered from the command line. However, tasks for changing the SELinux policy mode and managing the SELinux alert messages and automation script are conveniently supported in the Cockpit web console. For more details on how administrators can use the Cockpit web console to perform these tasks, see:

- [Change SELinux Policy Mode](#)
- [View SELinux Modifications and Copy Automation Script](#)
- [Manage SELinux Alert Message](#)

Change SELinux Policy Mode

Cockpit administrators can choose to use **SELinux** page in the web console to change the way SELinux runs at boot by changing its policy mode. The SELinux policy mode, by default, is set to enforcing.

What Do You Need?

- The Cockpit web console must be installed and accessible. For details, see these topics: [Install and Enable Cockpit](#) and [Log in to the Cockpit Web Console](#).
- Administrator privileges.

Steps

Using the Cockpit web console, follow these steps to change the SELinux Policy mode on the local host.

- In the Cockpit navigation pane, click **SELinux**. The **SELinux** page appears.
- In the **SELinux** page, click the **SELinux Policy mode** toggle button to switch the mode (Enforcing (default) | Permissive.)

! Important:

Cockpit enables you to switch the SELinux policy mode between enforcing and permissive. Enforcing mode is the default, and the recommended mode. Permissive mode doesn't deny operations based on SELinux security policy. Permissive mode can be helpful for SELinux policy development or debugging purposes).

⚠ WARNING:

When systems run SELinux in permissive mode, users and processes might label various file-system objects incorrectly. File-system objects created while SELinux is disabled aren't labeled at all. This behavior causes problems when changing to enforcing mode because SELinux relies on correct labels of file-system objects.

To prevent incorrectly labeled and unlabeled files from causing problems, SELinux automatically relabels file systems when changing from the disabled state to permissive or enforcing mode. Use the `sudo fixfiles -F onboot` command to create the `/.autorelabel` file containing the `-F` option to ensure that files are relabeled upon next reboot.

Before rebooting the system for relabeling, ensure the system boots in permissive mode, for example by using the `enforcing=0` kernel option. This setting prevents the system from failing to boot in case the system contains unlabeled files required by `systemd` before starting the `selinux-autorelabel` service.

View SELinux Modifications and Copy Automation Script

Cockpit administrators can view SELinux system modifications made to the local system using the **SELinux** page in the web console. If required, Cockpit administrators can also use the **Automation Script** link on the **SELinux** page to copy the same settings to other hosts.

What Do You Need?

- The Cockpit web console must be installed and accessible. For details, see these topics: [Install and Enable Cockpit](#) and [Log in to the Cockpit Web Console](#).
- The `policycoreutils-python-utils` package must be installed to transfer SELinux configuration settings to other Oracle Linux hosts.
- Administrator privileges.

Steps

Using the Cockpit web console, follow these steps to view the SELinux policy modifications that were applied to the local host, and optionally, to apply the same settings to other hosts.

1. In the Cockpit navigation pane, click **SELinux**. The **SELinux** page appears.
2. In the **SELinux** page, perform any of the following:
 - To view the SELinux policy modifications, navigate to the **System Modifications** section. The SELinux policy modifications applied to the host appear in a table format.
 - To apply the same SELinux policy settings to other hosts, perform these steps:
 - a. Click **View Automation Script**. The **Automated Script** dialog appears.
 - b. To copy the automation settings, click **Copy to clipboard**.

- c. To copy and import the configuration settings to other hosts, use the command line. For example, the syntax you use to copy and import the configuration settings might look similar to:

```
scp ./<my-selinux-settings> new-system-hostname:
```

```
new-system-hostname$ sudo semanage import -f ./<my-selinux-settings>
```

For more details, see the `scp` (secure copy) command line utility and the `semanage-import(8)` manual page.

Manage SELinux Alert Message

Cockpit administrators can monitor and manage access denial alert messages generated by SELinux using the **SELinux** page in the web console.



Note:

To perform similar tasks from the command line use `setenforce` and `sealert` tools. For more information, see *Troubleshooting Access Denial Messages* in [Oracle Linux: Administering SELinux](#).

What Do You Need?

- The Cockpit web console must be installed and accessible. For details, see these topics: [Install and Enable Cockpit](#) and [Log in to the Cockpit Web Console](#).
- Administrator privileges.

Steps

Using the Cockpit web console, follow these steps to view or dismiss access denial messages that are generated by SELinux on the local host.

1. In the Cockpit navigation pane, click **SELinux**. The **SELinux** page appears.
2. In the **SELinux** page, perform any of the following:
 - To view alert-generated access denial messages, navigate to the **SELinux access control errors** section. The messages appear in a table format.
 - To dismiss one or more alert generated message, select the check box for each message that you want to dismiss, and then click **Dismiss Alert**.

User Management Actions

Administrators, as needed, can use the web console to manage user access on the host system. For example, configurable options are available on the **Accounts** page to create or change user account profiles that define access privileges, account expiration, password management, and SSH authentication. Additionally, the **Accounts** page includes configurable options that enable administrators to remove a user account or to terminate a user session.

For more details about performing user management tasks from the **Accounts** page in the Cockpit web console, see these topics:

- [Create or Change User Accounts](#)
- [Disconnect User Sessions or Remove User Accounts](#)

Create or Change User Accounts

The **Accounts** page in the web console provides configurable properties that enable Cockpit administrators to manage user access on the host system. Upon accessing the **Accounts** page, a list of user accounts that are configured with management access appear. At a minimum, the root user account appears on the **Accounts** page.

What Do You Need?

- The Cockpit web console must be installed and accessible. For details, see these topics: [Install and Enable Cockpit](#) and [Log in to the Cockpit Web Console](#).
- For SSH authentication configurations only, the generated SSH key pair must be stored on the Cockpit system, the SSH key pair password must be known, and the file contents of the RSA key must be copied to the clipboard.

Note:

For more SSH key generation details, see [Generate SSH keys](#).

- Administrator privileges.

Steps

Using the Cockpit web console, follow these steps to manage user access on the host system.

1. Click **Accounts** in the Cockpit navigation pane. The **Accounts** page appears.
2. In the **Account** page, perform one of the following:
 - **Add a New User Account** - In the **Accounts** page, perform the following:
 - a. Click **Create new account**. The **Create new account** dialog appears.
 - b. In the **Create new account** dialog, specify the following properties and then click **Create**.

Full name	Enter the full name of the user.
User Name	Enter the account log in name for the user.
Password	Enter and confirm the account password for the user.

- **Edit an Existing User Account** - In the **Accounts** page, click an existing account name and change any of the following properties as needed, then click **Save**.

User Name	Edit to change how the display name for the user appears.
------------------	---

Groups In the **Groups** drop-down list box, select the applicable group(s) that you want to assign to the user account.

 **Note:**

A group is an entity which ties together multiple user accounts for a common purpose, such as granting access to particular files.

Options Click **Edit** to set the account expiration properties. In the **Account Expiration** dialog, click one of the following:

- **Never Expire account.** Enabled by default.
- OR–
- **Expire account YYYY-MM-DD.** Set a date for when the user account should expire. This option is helpful when managing a temporary user.

Password Set the applicable **Password** properties:

- **Never expire password.** Enabled by default. Click **Edit** to require the user to change the password within a defined number of days.
- **Set Password.** Select to change the existing user password.
- **Force password.** Select to require the user to enter a new password at next sign in.

Add Keys Click **Add Keys** to assign authorized SSH public keys to the user account. For example:

- a. In the **Add Keys** dialog, paste the contents of the public key file into the text box and click **Add**. The newly added public key assigned to the user account appears in the **Authorized Public SSH Keys** section of the **Account** page.
- b. To verify or change the private key path, select **SSH Keys** from the **Session** drop down menu. Choose to accept the default path or manually select a different path and click **ADD**.
- c. Click the **SSH Key On/Off** toggle to enable the key and, when prompted, provide the SSH password.

Disconnect User Sessions or Remove User Accounts

Using the **Account** page in the web console, Cockpit administrators can disconnect a specified user session or remove a specified user account configuration.

What Do You Need?

- The Cockpit web console must be installed and accessible. For details, see these topics: [Install and Enable Cockpit](#) and [Log in to the Cockpit Web Console](#).
- Administrator privileges.

Steps

Using the Cockpit web console, follow these steps to disconnect user sessions or to remove user account profiles from the host system.

1. In the **Accounts** page of the web console, select a user account, and then perform one of the following actions:
 - **Terminate User Session** – In the selected user account page, click **Terminate** to disconnect the web console session.

 **Note:**

If the **Terminate** button is unavailable, the selected user isn't logged in to the web console

- **Delete a User Account** – In the selected user account page, click **Delete** to remove the user account from the host system.

System Monitoring Actions

Cockpit makes it easy to monitor the system health, hardware inventory, event logs, and services running on the host system. For more details on how perform these system monitoring actions from Cockpit, see the following topics:


- [Health, Usage, and System Details](#)
- [View and Manage Services](#)
- [View and Filter Log Entries](#)

Health, Usage, and System Details

Cockpit collects system health, usage, and hardware details from the host in real time. Using the panels on the **Overview** page, Cockpit administrators can find relevant information about the system's health, resource usage, and hardware configuration.

The following table provides a brief description of each panel appearing on the **Overview** page.

Panel	Description
Health	The Health panel on the Overview page provides a general assessment of the system's health, including an indication of any failed services.

Panel	Description
Usage	<p>The Usage panel on the Overview page provides a graphical presentation of how the system CPU and memory resources are used. Clicking the metrics and history link displays a further information about the CPU, Memory, Disk, and Network usage.</p> <div data-bbox="935 453 1380 688"><p> Note:</p><p>The <code>cockpit-pcp</code> package is required to display the performance metrics on the system.</p></div>
System Information	<p>The System Information panel on the Overview page identifies information about the system hardware. Clicking the view hardware link displays more information about the hardware components detected on the system.</p>
Configuration	<p>The Configuration panel on the Overview page identifies general host configuration settings (host name, system time and date, and so on). Cockpit administrators, as needed can edit these settings to meet their needs. For more details about these settings, see:</p> <ul style="list-style-type: none">• General Host Configuration Actions

View and Manage Services

Cockpit administrators can use the **Services** page to monitor and manage the services running on the host system. For example, the **Services** page includes properties for activating or deactivating services, restarting, or reloading services, or managing the automatic start-up properties at boot. The **Services** page also provides filters that help you find the service you want to change. The names of the services appear in an easy-to-read table format. Each row displays the service name, description, state, and automatic start-up behavior. The tabs at the top of the **Services** page enable you to switch the view between Services, Targets, Sockets, Timers, and Paths.



Note:

Cockpit uses `systemd` to manage host service processes.

What Do You Need?

- The Cockpit web console must be installed and accessible. For details, see these topics: [Install and Enable Cockpit](#) and [Log in to the Cockpit Web Console](#).
- Administrator privileges.

Steps

Using the Cockpit web console, follow these steps to view and manage the behavior of system processes on the host system:

1. In the Cockpit navigation pane, click **Services**.
The **Services** page appears.
2. In the **Services** page, click a `systemd` service type tab of interest. For example, **Service**, **Targets**, **Sockets**, **Timer**, and **Path**
A list of service processes appear in a table format as follows:

Column	Displays
1	Name of service process.
2	Description of service process.
3	Status of service process.
4	Operating state of service process.

3. In the selected service type tab page, perform any of the following operations:

Filter page output

Filter the page output by:

- **User or system processes** – Click **System** or **User** (at top right side of page).
- **Process name or description** – In the text box, enter a process name or process description of interest.
- **Active process state or file process state** – Select the state of interest from the **Active State** drop-down list or **File State** drop-down list.

Change the operating state

Perform these steps:

- a. In the first column of the table, select a service name of interest.
The *service-name* page appears.
- b. In *server-process-name* page, perform the applicable action(s):
 - To activate or deactivate the process, turn on or off the **Account Service** toggle switch.
 - To start, stop, or reload the process, click the **Account Service Additional Actions** [□] menu.

View system logs

Perform these steps:

- a. In the first column of the table, select the service name of interest.
The *service-name* page appears.
- b. In *service-name* page, navigate to the Service Logs panel and then click **View all logs**.

Create Timer (Timer tab only)Perform these steps in the **Timer** tab:

- a. Click **Create timer** at the upper right top of the page.
A **Create timer** dialog appears.
- b. In the **Create timer** dialog, configure the required properties.
For more details about creating `systemd` timers, see *Work with Systemd Timer Units* in [Use systemd on Oracle Linux](#).

View and Filter Log Entries

The **Logs** page in Cockpit web console identify events and messages generated by the kernel, applications, and users signed in to the system. Cockpit administrators can exclude or include log entries appearing on the page by using predefined filter options or a free-form text search. In addition to the log filtering options, Cockpit administrators can also pause and resume the on-going reporting of log entries appearing on the **Logs** page as needed.

What Do You Need?

- The Cockpit web console must be installed and accessible.
For details, see these topics: [Install and Enable Cockpit](#) and [Log in to the Cockpit Web Console](#).

Steps



Using the Cockpit web console, follow these steps to manage the event and messages appearing on the **Logs** page.

1. In the Cockpit navigation pane, click **Logs**.
The **Logs** page appears displaying a listing of logged events.

 **Note:**

On this page, the text field, by default, contains a text string expression that filters the list of displayed events. To see all the logged events, clear the text field of the filter.

2. Perform any of the following tasks on the **Logs** page. You might need to clear the filter to see all the logged events.
 - **View event details:** In the list, click an event to access further details about that event. A journal page appears listing the details about the event, such as the priority level, syslog facility, syslog identifier, and the audit login UID and session.
 - **Pause or resume the log view:** Click **Pause** to stop new log entries from appearing on the page. Click **Resume** to display all on-going log entries, including the entries that were paused earlier from displaying.
 - **Change log view with predefined filter options:** Filter the log view by using the following predefined filters options:

Filter	Options
Time	<p>Select a relative time period for the events you want to display in the log view. For example:</p> <ul style="list-style-type: none"> - Current boot - Previous boot - Last 24 hours (default) - Last 7 days
Priority	<p>Select a priority level for the events that you want to display in log view. For example:</p> <ul style="list-style-type: none"> - Only emergency - Alert and above - Critical and above - Error and above (default) - Warning and above - Notice and above - Info and above - Debug and above
	<div style="border: 1px solid #0070c0; padding: 10px; background-color: #e6f2ff;"> <p> Note:</p> <p>The Debug and above option displays the most expansive list of events.</p> </div>
Identifier	<p>Select a syslog identifier for the type of events you want to include in the log view. For example:</p> <ul style="list-style-type: none"> - All (default) - cockpit-session - kernel - password - sshd - sudo - systemd
	<div style="border: 1px solid #0070c0; padding: 10px; background-color: #e6f2ff;"> <p> Note:</p> <p>The syslog identifier options correspond to the <code>journalctl --identifier</code> options.</p> </div>

- **Change log view with free-form search expression:**
 - a. Click the down-arrow in the `text` field to apply extra qualifiers to the search, or to create a free-form search expression.

Search Parameters	Description
Since and Until	Use the following date specification format to filter log entries by a specific date or time. <ul style="list-style-type: none"> YYYY-MM-DD HH:MM:SS Or, you can: <ul style="list-style-type: none"> Apply the following search strings: now, today, tomorrow, and yesterday. Express relative times by prefixing "-" or "+" to the search string.
Boot	By default, when a boot ID is omitted the logs for the current boot appear. To display logs from a specific boot, specify a boot ID, for example: <ul style="list-style-type: none"> boot=[ID] Where <i>ID</i> equals the boot order number (1, 2, and so on) found in the journal.
Unit	Filter log entries by a <code>systemd</code> unit (such as a service unit), or for units matching a specific pattern, for example: <ul style="list-style-type: none"> unit=UNIT PATTERN Note that when a pattern is specified, a list of unit names in the journal are compared with the specified pattern to display a list of all units matching the specified pattern.
Free-form search	Filter log entries by entering: <ul style="list-style-type: none"> Any text string OR A search string expression, for example: priority:err identifier:kernel

Important:

Each search parameter corresponds to a `journalctl` command parameter.

- b. Click **Search** to apply the search parameters to the log view; or, click **Reset** to clear the search fields.

Software Updates

Cockpit administrators can conveniently use the web console to keep their system software up-to-date with new improved functionality while addressing fixes for known issues. For more information on how to apply pending software updates from the Cockpit web console, see these topics:

- [Apply Pending Software Updates Manually](#)
- [Schedule Automatic Software Updates](#)

Apply Pending Software Updates Manually

Cockpit administrators can use the **Software** updates page in the web console to manually apply software updates to the host system. The **Software** updates page includes a **Status** section that tracks when the system was last checked for updates. If updates are available, Cockpit administrators can choose to manually apply them. If the system is up-to-date with the latest software, a green check mark appears with an up-to-date status message.

Note:

Alternatively, administrators can apply pending software updates from the command line using the `dnf upgrade` command. For more information on how to perform this task from the command line, see *Updating Software on Oracle Linux* in [Oracle Linux: Managing Software on Oracle Linux](#)

What Do You Need?

- The Cockpit web console must be installed and accessible. For details, see these topics: [Install and Enable Cockpit](#) and [Log in to the Cockpit Web Console](#).
- The Cockpit Software Update package (`cockpit-packagekit`) must be installed. In the case where the Software Update module doesn't appear in the web console navigation panel, see [Install and Manage Add-on Applications](#).
- Administrator privileges.

Steps

Using the Cockpit web console, follow these instructions to manually apply pending software updates on the host system.

1. In the Cockpit navigation pane, click **Software** updates. The **Software** updates page appears.
2. In the **Status** section of the **Software updates** page, perform any of the following:
 - **Verify when last check for updates occurred.** A status message appears indicating the last time the system was checked for pending software updates.
-OR-
 - **Manually check and install pending updates:**
 - a. Click **Check for updates** (curved blue arrow icon).
The system checks for pending updates. In the case that software updates are pending, a list of the pending updates appear.
 - b. To apply the software updates:
 - i. Perform one of the following:
 - Click **Install all updates** to apply all pending updates listed.
 - Click **Install security updates** to apply the pending security updates listed.

- Click **Install kpatch updates** to apply the pending `kpatch` updates listed.
- ii. Turn on the **Reboot** toggle switch to automatically reboot the system after applying the selected software updates.

 **Important:**

Avoid restarting the system until the updates are applied. If the **Reboot** toggle mode is turned off, a message appears to **Restart** the system. In message dialog, click **Ignore** to postpone the system restart until the updates are applied. After the updates are applied, a message appears to restart the system, click **Restart Now**.

- c. To ensure that the selected updates were successfully applied, log in to the Cockpit web console and verify the **Status** on the **Software updates** page.

Schedule Automatic Software Updates

Cockpit administrators can use the Software updates page in the web console to automatically schedule when software updates occur on the host system. Configurable properties for scheduling an automatic software update include selecting the update type (none, security, or all) and the frequency for how often the automatic update occurs.

 **Note:**

Alternatively, administrators can use the command line to configure automatic software updates. For more information about using the command line to perform this task, see *Updating Software Automatically* in [Oracle Linux: Managing Software on Oracle Linux](#)

What Do You Need?

- The Cockpit web console must be installed and accessible. For details, see these topics: [Install and Enable Cockpit](#) and [Log in to the Cockpit Web Console](#).
- The Cockpit Software Update package (`cockpit-packagekit`) must be installed. In the case where the Software Update module does not appear in the web console navigation panel, see [Install and Manage Add-on Applications](#).

 **Note:**

A first time set up is also required for automatic updates after installing the `cockpit-package kit`. The steps for performing the first time set up is covered in the procedure that follows this section.

- Administrator privileges.

Steps

Using the Cockpit web console, follow these instructions to configure automatic software updates on the host system.

1. In the Cockpit navigation pane, click **Software** updates.
The **Software updates** page appears.
2. (First time set up only) In the **Software updates** page, perform the following actions to set up the automatic software update option for the first time:

! Important:

If these actions (install and enable) were previously performed, proceed to the next Step.

Settings: Automatic updates - Not set up

If the automatic software update service was not previously enabled, an option to enable the automatic software service appears.
Click **Enable** to enable the automatic software update service.

Install Software: dnf-automatic

If the automatic software was not previously installed, an option to install the software appears.
Click **Install** to install the required software.

An **Edit** option appears in the **Automatic updates** section after installing the required software and enabling the automatic update service.

3. In the **Settings** section of the **Software updates** page, perform any of the following actions:
 - **View schedule for automatic updates** – A message appears indicating the current status of the **Automatic updates**, for example:
 - If an automatic update schedule is active, an **Updates will be applied ...** message appears indicating how often and when the automatic updates occur.
 - If the automatic update schedule is disabled, an **Automatic updates: Disabled** message appears.
 - **Edit the automatic update schedule:**
 - a. Click **Edit**. The **Automatic updates** dialog appears.
 - b. In the **Automatic updates** dialog, specify the following properties and then click **Save changes**.

Type	Select the update type that you want to automate (or disable). For example: <ul style="list-style-type: none">– None. Select to disable all automatic updates.– Security updates. Select to automate security related updates only.– All updates. Select to automate all updates (patches, bug fixes, security updates and so on).
When	In the drop-down lists, specify the frequency (daily or a specific day of week) and the time for the automatic updates to occur.

The system applies the automatic updates per the properties selected on the **Automatic Updates** dialog. Upon completing the updates, the system automatically restarts.

3

Network Management Tasks

Cockpit administrators can use the **Networking** page in the web console to help optimize network performance and prevent network disruptions. For more details about performing network management tasks from the Cockpit web console, see the following topics:

- [Manage Firewall Zoning Properties](#)
- [Monitor or Change Interface Connections](#)
- [Configure Network Bonding Properties](#)
- [Configure Network Teaming Properties](#)
- [Configure Network Bridging Properties](#)
- [Configure Network VLAN Properties](#)
- [View Network Host Log Files](#)

Manage Firewall Zoning Properties

Using the **Networking** page in the web console, Cockpit administrators can monitor and manage properties related to firewall zoning rules.



Note:

The firewall properties in the Cockpit web console work directly with the Oracle Linux `firewalld` management service.

- [Change Host Firewall State](#)
- [Display Firewall Zone Properties](#)
- [Control Access to Zone Services](#)
- [Add a New Predefined Zone](#)
- [Remove an Existing Predefined Zone](#)

Change Host Firewall State

Using the **Networking** page, Cockpit administrators can enable or disable the firewall state on the host system.

 **Note:**

By default, the firewall management service (`firewalld`) is enabled on the host system. When this service is enabled, all incoming network traffic is blocked with the exception where firewall zoning rules are set to enable incoming traffic for services and their ports.

For more information about a zone-based firewall implementation and the firewall management service in Oracle Linux, see [Oracle Linux 8: Configuring the Firewall](#) or [Oracle Linux 9: Configuring the Firewall](#).

What Do You Need?

- The Cockpit web console must be installed and accessible. For details, see these topics: [Install and Enable Cockpit](#) and [Log in to the Cockpit Web Console](#).
- The `cockpit-networkmanager` package is installed. Typically, this package is installed by default. For package installation details, see [Install and Manage Add-on Applications](#)
- Administrator privileges.

Steps

Using the Cockpit web console, follow these steps to enable or disable the firewall management service (`firewalld`) state on the host system.

1. In the Cockpit navigation pane, click **Networking**.
The Networking page appears.
2. In the **Firewall** panel of the Networking page, find the Firewall (Enabled|Disabled) toggle switch and then perform one of the following:
 - **To disable the `firewalld` service** – Click the **Firewall Enabled** toggle switch.
-OR-
 - **To enable the `firewalld` service** – Click the **Firewall Disabled** toggle switch.

Display Firewall Zone Properties

The `firewalld` management service filters all incoming interface traffic into one or more predefined zones. Each predefined zone has its own set of firewall rules for accepting or denying packets.

A default zone called, *public*, is automatically assigned to the host system during the installation of Oracle Linux. In cases, where a host system is configured as a multi-zoned system, other predefined zones are available to view in addition to the default public zone.

Using the Networking page in the web console, Cockpit administrators can view the firewall management rules associated with each zone.

What You Need

- The Cockpit web console must be installed and accessible.

For details, see these topics: [Install and Enable Cockpit](#) and [Log in to the Cockpit Web Console](#).

- The `cockpit-networkmanager` package is installed. Typically, this package is installed by default. For package installation details, see [Install and Manage Add-on Applications](#).
- Administrator privileges.

Steps

Using the Cockpit web console, follow these steps to review the host system's current Zone information:

1. In Cockpit navigation pane, click **Networking**. The **Networking** page appears.
2. In the **Networking** page, find the Firewall panel and perform one of the following to access and view the predefined zoning information:
 - Click the zone link appearing under the Firewall heading.

Important:

The name of the zone link indicates the number of *active zones*. A zone is only active if it has at least one interface or source assigned.

- OR -

- Click **Edit rules and zones** in the Firewall panel.

Note:

For information on how to edit the firewall management rules associated with a predefined zone, see [Control Access to Zone Services](#).

Information about each predefined zone appears in tables, for example:

- **firewalld predefined zone name.** The name of the predefined zone appears. For example: Public, External, DMZ, Work, Home, or Internal.
- **Interfaces and source addresses.** The names of the interfaces and source addresses that are allowed access through the predefined zone appear.

Important:

`firewalld` doesn't automatically pair the *interface source IP address ranges* to the default public zone. It does, however, automatically pair all the interface names to the default public zone. *Interface names* are the host names for the physical and virtual network interfaces that are configured on the system.

- **Services and ports.** The names of the access-allowed services and ports associated with the predefined zone appear.

Control Access to Zone Services

Cockpit administrators can control access to zone services by either adding access to a new service or removing access from an existing service. Configuration properties for adding or removing access to zone services are easily configurable using the **Networking** page in the web console.

What Do You Need?

- The Cockpit web console must be installed and accessible. For details, see these topics: [Install and Enable Cockpit](#) and [Log in to the Cockpit Web Console](#).
- The `cockpit-networkmanager` package is installed. Typically, this package is installed by default. For package installation details, see [Install and Manage Add-on Applications](#).
- Administrator privileges.

Steps

Using the Cockpit web console, follow these steps to add or remove access to zone services.

1. In the **Networking** page, find the Firewall panel and click **Edit rules and zones**. The names for the allowed services appear in a table under each Zone name.

 **Note:**

For single-zone systems, only one zone name appears.

2. In the table, find the zone configuration you want to edit, and perform any of the following actions:
 - **Remove existing service access**
 - a. In the allowed services section of the zone information, find the name of the service.
 - b. Click the actions menu (associated with the service name to be removed) and select **Delete**.
The selected service is removed from the allowed service list that appears under the Zone name.
 - **Add access to a new service**
 - a. Click **Add Services**.
The dialog for Add services to zone appears. This dialog provides a list services that are available to add to the current zone configuration.
 - b. In the **Add services to zone** dialog, perform one of the following:
 - Click **Services** to add services using standard ports.
Select the individual check boxes for the host system services that you want to add.

 **Note:**

Zone services assigned to standard ports are, by default, opened to accept traffic.

- OR -

- Click **Custom ports** to add a service using custom port. Enter the following information:

Property	Description
TCP or UDP	Enter comma separated ports, ranges, and service accepted. Example: 22, SSH, 80:80, 5900-5910
ID	The ID field automatically generates a custom ID based on the information entered in the TCP or UDP fields. Example: custom--ssh-ssh-5900-5910
Description	Enter a description for the accepted service and its custom port numbers.

 **Caution:**

Adding a service with custom ports can automatically reload the `firewalld` service, which can result in the loss of the runtime configurations.

- Click **Add**.
The selected services and their associated ports appear in the allowed service access list under the Zone name.

Add a New Predefined Zone

In addition to the default *public* predefined zone, the `firewalld` service provides several other predefined zones for configuration. Configuration properties for adding other predefined zones are easily configurable using the **Networking** page in the Cockpit web console.

What Do You Need?


- The Cockpit web console must be installed and accessible.
For details, see these topics: [Install and Enable Cockpit](#) and [Log in to the Cockpit Web Console](#).
- The `cockpit-networkmanager` package is installed. Typically, this package is installed by default. For package installation details, see [Install and Manage Add-on Applications](#).

- Administrator privileges.

Steps

Using the Cockpit web console, follow these steps to add other `firewalld` predefined zones to the host system.

1. In the **Networking** page, find the Firewall panel and click **Edit rules and zones**. A Firewall page appears listing information for the current zone configurations.
2. In the **Firewall** page, click **Add new zone**. The **Add zone** dialog appears.
3. In the **Add zone** dialog, perform the following:
 - a. Specify the following information:

Property	Description
Trust Level	Select a predefined zone from the list. Upon selecting a predefined zone, the Description property and Service included property identify information about the selected predefined zone and the <code>firewalld</code> services included.
Interfaces	Assign host interfaces to the predefined zone. Select the names of the available interfaces from the host interface list.
	<div style="border: 1px solid #0070C0; padding: 10px; background-color: #E6F2FF;"> <p> Note:</p> <p>A host interface can't be assigned to more than one zone at a time. By default, <code>firewalld</code> pairs all interfaces with the public zone. Therefore, the public zone is the only active zone. A zone is only active if it has at least one interface or source assigned. The <code>firewalld</code> service doesn't automatically pair sources (interface IP address ranges) to the public zone.</p> </div>
Allowed addresses	<p>Choose to assign one of the following:</p> <ul style="list-style-type: none"> • Entire subnet. Click subnet to allow firewall access to the entire subnet. -OR- • Range. Click Range to enter a specific range of IP addresses that are allowed access through the firewall.

- b.** Click **Add**.
The name of the newly added zone appears on the **Firewall** page.

Remove an Existing Predefined Zone


Cockpit administrators, as needed, can remove the configuration an existing `firewalld` predefined zone using the Networking page in the web console.

What Do You Need?

- The Cockpit web console must be installed and accessible.
For details, see these topics: [Install and Enable Cockpit](#) and [Log in to the Cockpit Web Console](#).
- The `cockpit-networkmanager` package is installed. Typically, this package is installed by default. For package installation details, see [Install and Manage Add-on Applications](#).
- Administrator privileges.

Steps

Using the Cockpit web console, follow these steps to remove an existing zone-based firewall configuration on the host system.

1. In the **Networking** page, find the Firewall panel and click **Edit rules and zones**.
A **Firewall** page appears listing information for the current zone configuration(s).
2. In the **Firewall** page, perform the following steps:
 - a.** Find the name of the zone (for example, work zone).
 - b.** Click the actions  menu (associated with the zone to be removed) and select **Delete**.
The selected zone is removed from the **Firewall** page.

Monitor or Change Interface Connections

The **Interface** section, on the **Networking** page of the web console, displays host network information about each configured network interface. Selecting an interface name provides details about the interface, and configurable properties that Cockpit administrators can choose to change. These properties include IP address, MAC address, connection state at reboot, and the interface status state (active or inactive).


What Do You Need?

- The Cockpit web console must be installed and accessible.
For details, see these topics: [Install and Enable Cockpit](#) and [Log in to the Cockpit Web Console](#).
- The `cockpit-networkmanager` package is installed. Typically, this package is installed by default. For package installation details, see [Install and Manage Add-on Applications](#).
- Administrator privileges.

Steps

Use the Cockpit web console, follow these steps to monitor the host interface network traffic, and to change their configuration.

1. In the **Networking** page, navigate to the Interface panel and view the following information about each host interface.
 - Interface name, IP address, and inbound and outbound traffic load.
 - At the top of the page, a set line charts appear. These line charts depict the total traffic load seen across all host interfaces. Click the drop-down menu to change the interval values appearing on the chart axis labels. By default, each chart depicts the data in 5-minute increments.
2. Click the name of an interface and perform any of the following tasks:
 - **Monitor interface performance or review device related details.** For example, the following interface-specific information appears.
 - Real-time line charts displaying the inbound and outbound traffic load. Click the drop-down menu to change the interval values on the line chart axis labels. By default, each chart depicts the data in 1-hour increments.
 - Connection name, device manufacturer model name, device status, and carrier speed.
 - Configured property values for: device MAC address, IP4 address, IP6 address, MTU packet or frame size, and the current toggle switch setting depicting the current interface state (inactive or active).
 - **Edit the interface configuration properties.** Choose to change any of the following properties:

Property	Description
MAC address <div style="border: 1px solid #0070C0; padding: 10px; margin-top: 10px;"> <p> Note:</p> <p>This address appears as a 12 digit hexadecimal number with a colon or hyphen separating every two digits (2C:54:91:28:C7:E3).</p> </div>	Click the mac-address link, edit the mac-address shown by selecting an option from the drop-down text box, then click Save .
Automatic Connection	Activate or deactivate the network interface connection after a reboot. Do one of the following: <ul style="list-style-type: none"> – Select the checkbox to automatically activate the connection after a reboot. -OR- – Clear the checkbox to automatically deactivate the connection after a reboot.
IPv4 or IPv6	Click edit to view and change the IPv4 or IPv6 configuration properties. Choose to set automatic generated IP property values or specify user-defined IP property values.

Property	Description
MTU	Click edit to change the interface maximum transfer unit size (MTU). Choose to set an automatic generated MTU value or specify a user-defined MTU property value.
Toggle switch	Click the toggle switch to control the network interface state (inactive or active). By default, the toggle switch appears blue with a check mark to indicate that the state is active. Clearing the toggle switch deactivates the state.

Configure Network Bonding Properties

Network bonding is the method of joining multiple interfaces together on a host system to create a bonded interface. A bonded interface can improve network throughput, and also provide a redundancy plan in the event of a failed interface. The behavior of a bonded interface is determined by the bonding mode. For example, different modes implement different levels of bonding for features like load balancing, fault tolerance, and failsafe.

Cockpit administrators can easily create, change, or delete a bonded interface by using the bonding configuration properties available on the **Networking** page.

What Do You Need?

- The Cockpit web console must be installed and accessible. For details, see these topics: [Install and Enable Cockpit](#) and [Log in to the Cockpit Web Console](#).
- Two or more physical or virtual network devices are installed on the host system.
- The `cockpit-networkmanager` package is installed. Typically, this package is installed by default. For package installation details, see [Install and Manage Add-on Applications](#).
- Administrator privileges.

Steps

Using the Cockpit web console, follow these steps to create, change, or remove a bonded interface on the host system.

1. In the **Networking** page, navigate to the Interface panel and perform any of the following:
 - **Create a bonded interface configuration.**
 - a. Click **Add bond**. The **Bond Settings** dialog appears.
 - b. In the **Bond Settings** dialog, specify the following properties and then click **Save**.

Property	Description
Name	Use the default name (for example, Bond0) or type in a user-specified bond name such as bond0 (round robin).

Property	Description
Interfaces	Select the interfaces you want bond from the available interface list.
MAC	In the drop-down list menu, select the MAC address for the bonded interfaces, or type a user-specified MAC address.
Mode	<p>In the drop-down list menu, select one of the following bonding modes for the selected bonded interfaces.</p> <ul style="list-style-type: none"> – Active Backup: A single interface is configured as the primary interface and other interfaces in the bond act as backups if the primary interface fails. – Round Robin: Network traffic is balanced by transmitting in sequential order beginning with the first available interface. If an interface fails, it's skipped in the round-robin selection. – XOR: Network traffic is balanced based on a hash policy derived from interface MAC addresses. This mode ensures that network traffic destined for specific peers always comes from the same physical interface. – Broadcast: All network traffic is sent on all network interfaces. This option includes fault tolerance, but it doesn't include load balancing. – 802.3ad: Uses the IEEE 802.3ad dynamic link aggregation policy and requires an 802.3ad capable switch. Traffic is broadcast in aggregation groups to maximize fault tolerance and to provide load balancing functionality. – Adaptive transmit load balancing: Outgoing traffic is balanced across interfaces within the bond based on each interface's current load. Incoming traffic is delivered to the current active interface. – Adaptive load balancing: This option is similar to dynamic link aggregation, but it requires the use an 802.3ad capable switch. Outgoing traffic is handled in the same manner as adaptive transmit load balancing. Incoming traffic is balanced based on ARP negotiation.
Primary	The Primary property only appears when the Mode is set to Active Backup. In the drop-down menu, select the primary active interface device.

Property	Description
Link Monitoring	<p>In the drop-down list box, select the applicable link monitoring option. For example:</p> <ul style="list-style-type: none"> – MII (Recommended): This default option uses the MII (Media Independent Interface) monitor, which detects carrier signal for each interface using the local device driver or MII registers to determine carrier state. You can set the <code>Monitoring Interval</code>, which determines how often the carrier state is checked in milliseconds; the <code>Link up delay</code> to determine how long to wait, in milliseconds, until using an interface that's up; and the <code>Link down delay</code> which determines how long to wait before switching to another interface if the interface is marked as down. – ARP: The ARP monitor sends ARP queries to peer systems on the network and uses the response to indicate whether an interface is up. The ARP monitor relies on the device driver to track the last transmission and receipt times. If the information isn't updated by the device driver, the interface is marked as down.
Monitoring Level; Link Up; and Link Down	<p>Edit these properties as required. Typically, the default values for these properties are only changed for troubleshooting purposes.</p>

The name of the newly bonded interface appears in the Interface panel of the **Networking** page.

- **Edit, disable, or remove existing bonded interface properties.**
 - a. In the **Interface** panel, click the name of the bonded interface that you want edit. The **Networking > [bond name]** page appears.
 - b. In the **Networking > [bond name]** page, edit the configurable bonding properties as needed. For example:

Toggle the connection state of a bonded interface.	Turn off the toggle switch to deactivate the bonded link state or turn on the toggle switch to activate the bonded link state.
Delete a bonded configuration	Click Delete (next to the bonded interface name) to remove the bonded interface configuration.
Toggle the connection state of an interface	Turn off the toggle switch to deactivate an interface that is part of a bonded interface. Turn on the toggle switch to activate the interface that's part of a bonded interface.

Automatically connect a bonded interface (after reboot)	Select the check box to enable automatic connection after reboot, or clear the check box to disable automatic connection after reboot.
Change MAC, IPv4, IPv6 addresses or the MTU size	In the Interface table, turn off the toggle switch to deactivate the interface state, or turn on the toggle switch to activate the interface state.
Bond: Modename	Click the link to edit any of the applicable properties appearing on the Bond Settings dialog. For example, mode option, interface assignment, and so on.
Edit interface-specific properties	In the Interface table, click the name of one of interfaces that's part of the bonded interface, and then as needed, edit any of properties appearing on the Bond Settings dialog.

Configure Network Teaming Properties

Network interface teaming is a feature-rich alternative to network interface bonding. Teaming, like bonding, is another way to implement network link aggregation, where one or more interfaces are bundled together to act as one logical link. Network teaming functionality is provided by the small kernel driver while network bonding functionality is provided by the bonding driver. For a comparison of the features, see [Teaming and Bonding Feature Comparison](#).

Cockpit administrators can easily create, change, or delete a network teaming interface by using the teaming configuration properties provided on the **Networking page**.

! Important:

In Oracle Linux 9, network teaming has been deprecated as a feature in favor of network bonding. If you're using Oracle Linux 9, [configure network bonding](#) instead.

What Do You Need?

- The Cockpit web console must be installed and accessible. For details, see these topics: [Install and Enable Cockpit](#) and [Log in to the Cockpit Web Console](#).
- Two or more physical or virtual network devices are installed on the host system.
- The `cockpit-networkmanager` package is installed. Typically, this package is installed by default. For package installation details, see [Install and Manage Add-on Applications](#).
- Administrator privileges.

Steps

Using the Cockpit web console, follow these steps to create, change, or remove a network teaming interface configuration on the host system.

1. In the **Networking** page, navigate to the Interface panel and perform any of the following:
 - **Create a bonded interface.**
 - a. Click **Add Team**. The **Team Settings** dialog appears.
 - b. In the **Team Settings** dialog, specify the following properties and then click **Save**.

Property	Description
Name	Use the default name (for example, team0) or type in a user-specified team name such as team0 (round robin).
Interface Ports	Select the interfaces that you want bundle together from the available port list.
Runner	<p>Runner modes determine the load balancing and fail-over schemes applied to the teamed port interfaces. In the drop-down list menu, select one of the following runner modes for the selected teamed port interfaces.</p> <ul style="list-style-type: none"> – Broadcast. All network traffic is sent across all port interfaces. – Round Robin. Network traffic is balanced by transmitting in sequential order beginning with the first available port interface. If an interface fails, it's skipped in the round-robin selection. – Active Backup. A single port interface or link is used while other port interfaces act as backups. – LACP. Uses the IEEE 802.3ad dynamic link aggregation policy and requires an 802.3ad capable switch. Traffic is broadcast in aggregation groups to maximize fault tolerance and to provide load balancing functionality.
Link watch	<p>Link watch monitors the state of the teamed port interfaces. In the drop-down list box, select the applicable link watch option. For example:</p> <ul style="list-style-type: none"> – Ethtool (Default). Link watch uses the Oracle Linux ethtool utility for managing the teamed port interfaces. – ARP. Link watch uses the Oracle Linux arp_ping utility and the Address Resolution Protocol (ARP) at the data layer to send ARP requests to destination hosts.

Property	Description
Link Up and Link Down	<p>Set the link up or link down timer delays in milliseconds. Delays enable the teamed port interfaces to synchronize. For <code>Link up delay</code>, you might want to specify a time between when a device link is reestablished and when the device can be used to service network traffic.</p> <p>For <code>Link down delay</code>, some devices and switches might take some time before their backup mode becomes activated. Specifying a delay prevents a fail-over to immediately occur before those backup devices are ready to be used.</p>

The name of the newly teamed port interface appears in **Interface** section of the **Networking** page.

- **Edit, disable, or remove existing team interface properties.**
 - a. In the **Interface** table, click the name of a team interface that you want edit.
The **Networking [team name]** page appears.
 - b. In the **Networking [team name]** page, edit the configurable teaming properties as needed. For example:

Toggle the connection state of the team interface.	Turn off the toggle switch to deactivate the team interface state or turn on the toggle switch to activate the team interface state.
Delete the team interface configuration	Click Delete (next to the team interface name) to remove the bonded interface configuration.
Toggle the connection state of an interface port	In the Ports table, turn off the toggle switch to deactivate a port or turn on the toggle switch to activate the port.
Automatically connect team interface (after reboot)	Select the checkbox to enable or clear the check box to disable.
Edit the Team: <i>name</i> properties	Click the link to edit any of the applicable properties appearing on the Bond Settings dialog. For example, runner option, interface assignment, and so on.
Edit port-specific properties	In the Port table, click the name of one of ports that's part of the team interface, and then as needed, edit any of properties appearing on the Team Settings dialog

Teaming and Bonding Feature Comparison

The following table provides a comparison of the features in Bonding and Teaming.

Feature	Bonding	Teaming
broadcast Tx policy	Yes	Yes
round-robin Tx policy	Yes	Yes
active-backup Tx policy	Yes	Yes
LACP (802.3ad)	Yes (active only)	Yes
Hash-based Tx policy	Yes	Yes
User can set hash function	Yes	Yes
Tx load-balancing (TLB)	No	Yes
LACP hash port select	Yes	Yes
load-balancing for LACP	No	Yes
Ethtool link monitoring	Yes	Yes
ARP link monitoring	Yes	Yes
NS/NA (IPv6) link monitoring	No	Yes
ports up/down delays	Yes	Yes
port priorities and stickiness ("primary" option enhancement)	No	Yes
separate per-port link monitoring setup	No	Yes
multiple link monitoring setup	Limited	Yes
lockless Tx/Rx path	No (rwlock)	Yes (RCU)
VLAN	Yes	Yes
user-space runtime control	Limited	Full
Logic in user-space	No	Yes
Extensibility	Hard	Easy
Modular design	No	Yes
Performance overhead	Low	Very low
D-Bus interface	No	Yes
multiple device stacking	Yes	Yes
zero config using LLDP	No	(in planning)
Network Manager integration	Yes	Yes

Configure Network Bridging Properties

A *network bridge* joins two or more network segments and enables them to work as a single network. Bridging is implemented at the datalink layer (L2) of the networking stack. Bridges use a packet-forwarding mechanism based on MAC addresses to connect subnetworks together.

Using the bridge properties on the **Networking** page in the web console, Cockpit administrators can easily create and maintain host bridge configurations.

What Do You Need?

- The Cockpit web console must be installed and accessible.
For details, see these topics: [Install and Enable Cockpit](#) and [Log in to the Cockpit Web Console](#).
- Two or more physical or virtual network devices are installed on the host system.

- The `cockpit-networkmanager` package is installed. Typically, this package is installed by default. For package installation details, see [Install and Manage Add-on Applications](#).
- Administrator privileges.

Steps

Using the Cockpit web console, follow these steps to create, change, or remove a network bridged configuration on the host system.

1. In the **Networking** page, navigate to the Interface panel and perform any of the following:
 - **Create Linux bridge configuration.** Click **Add Bridge**, specify the following properties on the **Bridge Settings** dialog, and then click **Save**.

Property	Description
Name	In the Name text box, choose to use the default bridge name provided (for example, <code>bridge0</code>) or enter a user-specific bridge name.
Ports	In the Ports checkbox list, add interfaces to the bridge by selecting the applicable check boxes.
Option	Select or clear the checkbox for Spanning Tree Protocol support.

 **Note:**

For multiple or redundant bridge configurations, enabling the Spanning Tree Protocol helps to prevent multiple hops and cyclic routes.

The name of the new bridge configuration appears in **Interface** pane of the **Networking** page.

- **Edit, disable, or remove existing Linux bridge properties.**
 - a. In the **Interface** panel, click the name of a bridge (for example `bridge0`) that you want edit.
The **Networking [bridge name]** page appears.
 - b. In the **Networking [bridge name]** page, edit the configurable bridge properties as needed. For example:

Toggle the state of a bridge connection	Turn off the toggle switch to deactivate the bridge connection state or turn on the toggle switch to activate the bridge connection state.
---	--

Delete a bridge configuration	Click Delete (next to the bridge name) to remove the bridge configuration on the host.
Toggle the state of a bridge port connection	In the Ports table, turn off the toggle switch to deactivate a port or turn on the toggle switch to activate the port.
Automatically connect bridge configuration (after reboot) IPv4 or IPv6 Network Address	Select the checkbox to enable or clear the check box to disable. Click the applicable Automatic link to change the IPv4 or IPv6 network addressing.
Edit bridge port: <i>bridge name</i> properties	Click the link to edit any of the applicable properties appearing on the Bridge Settings dialog. For example, bridge name, assigned ports, and to set the spanning tree protocol option.
Edit port-specific bridge properties	In the Ports table, click a port name that's part of the bridge configuration, and then as needed, edit any of the following port specific property values for MTU , bridge port , and Automatic Connection .

Configure Network VLAN Properties

A *Virtual Local Area Network* (VLAN) is a logical network within a physical network. Cockpit administrators, as needed, can build a VLAN configuration from any of the available network interfaces on the host system. Properties for creating and maintaining a VLAN configuration are available on the **Networking page** of the Cockpit web console.



Note:

For more information about implementing a VLAN connection on the host system, see [Create VLANs on Oracle Linux](#).

What Do You Need?

- The Cockpit web console must be installed and accessible.
For details, see these topics: [Install and Enable Cockpit](#) and [Log in to the Cockpit Web Console](#).
- The parent VLAN network interface must support VLAN tags.
- The network switch connected to the host system must support VLAN. For more details, see the documentation provided with the switch.
- If you configure the VLAN on top of a bonded interface, the following requirements apply:
 - The port connections must be up to support a VLAN configuration over a bonded interface.
 - The bond `fail_over_mac=follow` option isn't supported for VLAN configurations over a bonded interface.

- IPv4 from a DHCP server and IPv6 auto configuration options aren't supported on a VLAN over a bonded interface.
- The `cockpit-networkmanager` package is installed. Typically, this package is installed by default. For package installation details, see [Install and Manage Add-on Applications](#).
- Administrator privileges.

Steps

Using the Cockpit web console, follow these steps to create, change, or remove a VLAN configuration on the host system.

1. In the **Networking** page, navigate to the **Interface** panel and perform any of the following:
 - **Create VLAN configuration.** Click **Add VLAN**, specify the following properties on the **VLAN Settings** dialog, and click **Save**.

Property	Description
Parent	In the Parent drop-down list, assign a parent interface to the VLAN configuration. For example, <code>enp0s3</code>
VLAN ID	In the VLAN ID text box, use the default VLAN ID provided (for example, 1) or enter a user-specified VLAN ID.
Name	In the Name text box, keep the default name provided (for example, <code>enp0s3.1</code>) or enter a user-specified VLAN name.

The name of the new VLAN configuration appears in **Interface** panel on the **Networking** page.

- **Edit or remove existing VLAN properties.**
 - a. In the **Interface** panel, click the name of a VLAN (for example, `enp0s3.1`) that you want edit.
The **Networking [VLAN name]** page appears.
 - b. In the **Networking [VLAN name]** page, edit the configurable VLAN properties as needed. For example:

Toggle the state of a VLAN connection	Turn off the toggle switch to deactivate the VLAN connection state or turn on the toggle switch to activate VLAN connection state.
Delete a VLAN configuration	Click Delete (next to the VLAN name) to remove the VLAN configuration.
Automatically connect the VLAN configuration (after reboot)	Select the checkbox to enable automatic connection after reboot, or clear the check box to disable automatic connection after reboot.

Edit IPv4 or IPv6 Network Address	Click the applicable Automatic link to change the IPv4 or IPv6 network addressing properties.
Set MTU size	Click the link to change the set MTU size.
Edit VLAN: Parent: <i>name</i> properties	Click the link to edit any of the applicable properties appearing on the VLAN Settings dialog. For example, VLAN name, interface assignment, and VLAN ID.

View Network Host Log Files

The network logs appearing on the **Networking** page are specific to the Network Manager service. By default, the 10 most recent log entries appear. The list is identical to the output of the following command:

```
sudo journalctl -u NetworkManager -n 10S
```

Using the **Networking** page in the web console, Cockpit users can click the **View all logs** option to view a full list of network-related log events. Clicking a log entry in the **Networking** page or Logs page provides further details about the logged event.

4

Storage Management Tasks

Cockpit administrators can use the **Storage** page in the web console to help plan and maintain their host system storage needs. The **Storage** page provides configurable options for disk partitioning management, data security management, virtual volume management, and data compression. The **Storage** page also provides configurable options to manage NFS and iSCSI storage connections.

For further details about how to use the storage management options from the Cockpit web console, see the following topics:

- [Storage Management Installation and Overview](#)
- [Manage Disk Devices and Partitions](#)
- [Encrypt Block Devices With LUKS](#)
- [Unlock Encrypted Devices Using Tang Server Key](#)
- [Manage Logical Volumes With LVM](#)
- [Build and Manage Software RAID Devices](#)
- [Manage NFS Mounted Connections](#)
- [Manage Connections to iSCSI Targets](#)

Storage Management Installation and Overview

The Storage application enables Cockpit administrators to perform routine storage management tasks on a host system. Before you can access the **Storage** page in the web console, the cockpit storage application must be installed, for example:

```
sudo dnf install cockpit-storaged
```



Note:

For more information about adding applications to the Cockpit web console, see [Install and Manage Add-on Applications](#).

Upon accessing the **Storage** page, administrators can easily monitor and manage all things to do with disk storage. For example, at the top of the page, two line graphs appear depicting the disk's read and write performance and following the line graphs is the disk's local file system information. In addition, configuration properties are easily accessible for adding or changing RAID devices, volume groups, iSCSI devices, and storage drives. Further down the page, a summary appears in the **Log** section displaying real-time storage events generated by the host system.

For further details about the storage management functionality available from the Cockpit web console, see the following table:

 **Note:**

Administrator privileges are required to view and use the configurable storage properties. Read-only properties for viewing storage information are accessible for users with limited access.

Filesystems	The Filesystem table on the main Storage page identifies the physical partitions and logical volumes configured on a host storage device. Clicking an entry in the table displays configuration details about the storage device and properties for managing its physical partitions or logical volumes.
NFS mounts	When available, the NFS mounts table on the main Storage page identifies the NFS mount connections configured on the host system. Clicking the entry in the NFS mount table displays connection information and configurable properties for that NFS mount. When NFS mounts are available but not configured on the host, the table displays a message indicating "No NFS Mounts are set up". Clicking the plus [+] icon displays a dialog for configuring an NFS mount connection.
Storage Logs	The Storage Log section on the main Storage page identifies storage related events for observation. Clicking a log entry displays more information about the selected storage event.
Devices	The Devices section on the main Storage page identifies RAID devices or logical volume groups configured on the host system. Clicking the Devices menu displays options to create a RAID device or logical volume (LVM) group.
Drives	The Drives section on the main Storage page, identifies the block storage type drives that are configured on the host system. Clicking a storage drive entry in this section displays information about that storage drive's configuration, as well as configurable properties to manage its storage partitions.
iSCSI target	When available, the iSCSI target section on the main Storage page identifies the iSCSI server connections configured on the host system. Clicking the edit icon displays properties for changing a configured connection. Clicking the Plus [+] icon displays properties for configuring a new iSCSI server connection. For instances where iSCSI targets aren't configured on the host, a message appears indicating "No iSCSI targets" are set up.
Other Devices	When available, the Other Devices section on the main Storage page identifies attached data storage host configurations. For example, attached OCI block data volume configurations appear in this section. Clicking on a block data volume entry in this section displays the file directory path, the storage capacity size and the data volume's contents. Configurable properties for managing the block storage data volume content are also available.

 **Note:**

The host system storage configuration appearing on the **Storage** page reflects the storage devices discovered by the `cockpit-storaged` application. For specific details about the `cockpit-storaged` application, see the Cockpit project website (<https://cockpit-project.org/guide/latest/feature-storaged>).

Manage Disk Devices and Partitions

Cockpit administrators, as needed, can easily manage host configured drives and disk partitions using the **Storage** page in the web console. For more details, see the following topics:

- [Create Physical Disk Partitions](#)
- [Storage Partitioning Considerations and Prerequisites](#)
- [View and Change Drive Partition Properties](#)

Create Physical Disk Partitions

Cockpit administrators can use the **Storage** page to partition storage regions of a new or an existing physical disk device (for example, hard drive, solid state drive, RAID device, a disk device that's part of an existing volume group, and so on).

Creating a disk partition involves:

1. **Creating a new partition table** – The **Initialize Drive** dialog in Cockpit defines the partitioning format and erases all existing data on the drive. The newly created partition table tracks the size and location of each partition on the drive.
2. **Creating a new partition** - The **Create New Partition** dialog in Cockpit enables you to define the partition size, file system type, and properties for mounting and encrypting a partition. Free space must be available on the drive to perform this operation.

What Do You Need?

- The Cockpit web console must be installed and accessible. For details, see these topics: [Install and Enable Cockpit](#) and [Log in to the Cockpit Web Console](#).
- The `cockpit-storaged` application package must be installed.

 **Note:**

If the `cockpit-storaged` package isn't installed, see [Install and Manage Add-on Applications](#)

- Review [Storage Partitioning Considerations and Prerequisites](#) prior to performing the steps in this procedure.
- Administrator privileges.

Steps

Using the Cockpit web console, follow these steps to create disk partitions on a host disk device.

1. In Cockpit navigation pane, click **Storage**.
The **Storage** page appears.
2. In the **Storage** page, perform one of the following:
 - **To partition a new disk device**, select the device from the **Devices** table. The **Storage [device name]** page appears. Proceed to Step 3.
-OR-
 - **To repartition an existing disk**, select the device from the **Drives** table. The **Storage [device name]** page appears. Proceed to Step 3.
-OR-
 - **To add a new partition to disk device and using existing partition table**, select the device from the **Drives** table. The **Storage [device name]** page appears. Proceed to Step 4.
3. In the **Storage [device name]** page, perform the following:
 - a. Click **Create New Partition Table** and specify the following properties:




Partitioning	In the Partitioning drop-down list box, select the appropriate partitioning format for creating and organizing disk partitions.
Overwrite	Select or clear the Overwrite checkbox. When selected, the deleted data is overwritten with zeros, making the deleted data unrecoverable.

- b. Review the change summary.
 - c. Click **Initialize** to prepare the disk device with a new partitioning format. All data is removed on the disk device at this time.
 - d. Upon completing the disk initialization, proceed to Step 4 to create partition.
4. In the **Storage [device name]** page, perform the following:
 - a. Navigate to the table below the device information and then find and click the option to **Create a new partition**.

! Important:

The table title on this page might appear with a Partition heading or Content heading. The option to **Create a new partition** appears only when free space is available on the device to support the creation a new partition.

- b. In the **Create new partition** dialog, specify the following properties:

<p>Size</p>	<p>Specify the size of the new partition.</p> <div data-bbox="816 247 1458 504" style="border: 1px solid #0070C0; padding: 10px; margin-top: 10px;"> <p> Note:</p> <p>For LVM configurations, consider creating a single partition of the whole disk and label it as an LVM physical volume. Tracking system hardware is easier if each real disk only appears once.</p> </div>
<p>Type</p>	<p>In the Type drop-down box, select a file system format type, for example:</p> <ul style="list-style-type: none"> • XFS (recommended). The default high performance scalable file system format for Oracle Linux systems disk devices. This file system type caters to large logical volumes, switching physical drives online without outage, and growing an existing file system. • EXT4. A scalable extension of the ext3 file system. This file system caters to logical volumes, switching physical drives online without outage, growing and shrinking an existing file system. • No File system. Creates a partition without creating a file system.
<p>Overwrite data</p>	<p>Perform one of the following:</p> <ul style="list-style-type: none"> • Clear the Overwrite checkbox. This option doesn't overwrite existing data, it only overwrites the header. <div data-bbox="865 1119 1458 1318" style="border: 1px solid #0070C0; padding: 10px; margin-top: 10px;"> <p> Note:</p> <p>Overwriting the header only is considered a faster process but less secure because the deleted data might be recoverable.</p> </div> <ul style="list-style-type: none"> • Select the Overwrite checkbox. When selected, the deleted data is overwritten with zeros, making the deleted data unrecoverable. <div data-bbox="865 1444 1458 1648" style="border: 1px solid #0070C0; padding: 10px; margin-top: 10px;"> <p> Note:</p> <p>The Overwrite process is slower but is considered a more secure option as all the deleted data is overwritten with zeros.</p> </div>
<p>Mount point and options</p> <p>Encryption (if option is available for this storage device configuration)</p>	<p>Specify the appropriate mount point directory and applicable options for when to mount the partition.</p> <p>In the Encryption drop-down list box, select the appropriate encryption option. For more details, see Encrypt Block Devices With LUKS.</p>

- c. Click **Create**.

Formatting a partition with a new file system can take several minutes and depends on the partition size and format type selected. When complete, the new partition is available for system use (per the mounting options specified).

- d. (Optional) If enough space is available on the disk device to accommodate other partitions, you can repeat Step 4 to define additional partitions.

Storage Partitioning Considerations and Prerequisites

Consider the following information when partitioning storage space on host disk devices.

- Oracle Linux, at installation, provides a default disk partition layout. When the default partition layout is selected, the installer typically assigns 100 MB for `/boot`, 2 GB for swap, and the remainder is assigned to the root (`/`) partition.
- Oracle Linux requires a minimum of one partition for the root (`/`) file system.
- For hard disks with a *Master Boot Record* (MBR), the partitioning scheme permits up to 4 primary partitions. In turn, a primary partition can further be divided up to 11 logical partitions. The primary partition that contains the logical partitions is known as an extended partition. The MBR scheme functions for disks up to 2 TB in size.
- For hard disks with a *GUID Partition Table* (GPT), you can configure up to 128 partitions. The GPT partition scheme doesn't use the concept of extended or logical partitions. If the disk's size is larger than 2 TB, you can use GPT to configure the device partitions.
- When partitioning most *block storage devices* (for example: hard disk drives, solid state drives (SSD), LUNs on storage arrays and host RAID adapters), align primary and logical partitions on one-megabyte (1048576 bytes) boundaries. If partitions, file system blocks, or RAID stripes are incorrectly aligned and overlap the boundaries of the underlying storage's sectors or pages, the device controller must change twice as many sectors or pages than when the **correct** alignment is used.
- **Prerequisites:**
 - Backup data on the disk device that you don't want to lose.

 **Note:**

If the device is empty, Cockpit describes the content of the storage device as unknown

- Disk devices in which you want to configure partitions must be:
 - * Unmounted. For more information about unmounting a partition, see *Unmount or mount a disk partition* in [View and Change Drive Partition Properties](#).
 - * Visible to Cockpit and selectable in the applicable storage type table (Drives, Devices, or Other Devices).
- Disk devices must have enough free space to support creating a new partition.

For more information about partitioning Oracle Linux storage devices, see *Using Disk Partitions* in [Oracle Linux 8: Managing Storage Devices](#) or [Oracle Linux 9: Managing Storage Devices](#).

View and Change Drive Partition Properties

The Storage page enables Cockpit administrators to view and change host storage device information and partition properties as needed. Configurable properties available for managing partitions include: mount, unmount, format, and delete. Other properties are also available for initializing a drive and changing the current disk partitioning format (MBR, GPT). For information about how to create a disk partition, see [Create Physical Disk Partitions](#) and [Storage Partitioning Considerations and Prerequisites](#).

Note:

System administrators can manage storage partitions using the command line. For more information about managing Oracle Linux storage devices, see *Using Disk Partitions* in the [Oracle Linux 9: Managing Storage Devices](#).

What Do You Need?

- The Cockpit web console must be installed and accessible. For details, see these topics: [Install and Enable Cockpit](#) and [Log in to the Cockpit Web Console](#).
- The `cockpit-storaged` package must be installed.

Note:

If the `cockpit-storaged` package isn't installed, see [Install and Manage Add-on Applications](#)

- An available drive and partitions on host that you want to view or manage.
- Administrator privileges.

Steps

Using the Cockpit web console, follow these steps to view and manage drive and partition configuration properties on the host system.

1. In Cockpit navigation pane, click **Storage**. The **Storage** page appears.
2. In the **Storage** page, navigate to the **Drives** table and then click a drive entry that you want to view. The **Storage [drive model name]** page appears providing information about the drive and its configured partitions or contents.
3. In the **Storage [drive model name]** page, perform any of the following:
 - **View drive details and partition or content information:**
 - a. In the **Storage [media model name]** page, review the **Drive details** and the applicable **Partitions** table or **Contents** table.

For example, the information on this page might appear as follows:

- **Drive panel:** This panel typically displays the: 1) Model [examples: harddisk, VBox harddisk, blockvolume, CD-ROM, and so on]; 2) Firmware, 3) Serial No; 4) Capacity; and 5) Device file system directory path.
- **Partitions table:** If the drive media type is a hard disk drive or block storage device, a partitions table appears following the Drive table. This **Partitions** table typically identifies: 1) the *standard partitions* and *logical volume*, 2) the partition file system directory path, 3) the partition file system type (EXT 4 or XFS), and 4) if the partition is mounted, the mount point directory appears.

 **Note:**

The *logical volumes* only appear in the partition table when the drive is part of a volume group.

Partition configurable options include:

- * For standard disk partitions, the standard configurable properties typically include **Create New Partition Table**, **Mount**, **Unmount**, **Delete**, **Format**, and **Create Partition**. For encrypted partitions, extra configurable properties are provided for managing the encryption keys and passphrases.
- * For logical volumes, the standard configurable properties typically include **Lock**, **Format**, and **Delete**. For encrypted partitions, extra configurable properties are provided for managing the encryption keys and passphrases.

 **Note:**

For information about logical storage groups, see [Manage Logical Volumes With LVM](#).

- **Contents table:** If the drive media is for example a CD-ROM device, a **Contents** table appears. The Contents table typically identifies: 1) the file system directory path, 2) the file system type, and 3) if the media is mounted, the mounted directory path.
- b. To expand entries appearing in the **Partition** table or **Contents** table, navigate to an entry in a row and then click the **down arrow** icon. One of the following occurs:
- For *standard disk partition* entries, the table row expands displaying tabs for viewing: 1) the partition details, 2) the file system information, and 3) for encrypted partitions, the partition encryption properties.
 - For *logical volume partition* entries, the table row expands displaying: 1) the logical volume management partition details, and 2) for encrypted logical partitions, a second tab appears displaying the logical volume encryption properties.

- For *content* entries, the table row expands displaying the storage media file system directory path and configurable properties for editing the Name of the file system and the Mount point directory configuration.
- c. (Optional) In the expanded **File system** tab, click the **edit** links to change either the file system **Name** or the **Mount point** configuration information.

 **Note:**

The properties for **Name** and **Mount point** are only configurable for standard physical files systems.

- d. (Optional) In the expanded **Encryption partition** tab, click the **edit** links to change the encryption passphrase properties or click the Plus [+] or Minus [-] icons to add or remove encryption keys. For more information about configuring encryption properties, see [Encrypt Block Devices With LUKS](#) or [Unlock Encrypted Devices Using Tang Server Key](#).
- **Initialize drive with new partition table properties.**

 **WARNING:**

Initializing a drive erases all data on the disk.

- a. In the **Storage [drive model name]** page, click **Create Partition Table**. The **Initialize [drive file directory path]** dialog appears.
- b. In **Initialize [drive file directory path]** dialog, specify the following information:

Partitioning	In the Partitioning drop-down list box, select the appropriate partitioning format for creating and organizing disk partitions.
Overwrite	Select or clear the Overwrite checkbox. When selected, the deleted data is overwritten with zeros, making the deleted data unrecoverable.

A summary of changes appears for review.

- c. Review the changes and then click one of the following:
 - **Initialize:** To proceed with initializing the drive.
 - **Cancel:** To dismiss the Initialize [drive file directory path] dialog.

– **NOT_SUPPORTED:**

To create new partitions on an initialized drive, see Step 4 in [Create Physical Disk Partitions](#) For information about the default partition layout created at installation or the number of disk partitions available by each partitioning format, see [Storage Partitioning Considerations and Prerequisites](#)

- **Unmount or mount a disk partition:**

! Important:

Before unmounting a partition, ensure that the file system for that partition isn't in use by any system process; otherwise, the unmount operation fails and an error message appears indicating the partition file system is in use.

- a. In the **Partition** table, find the partition entry and then select **Unmount** or **Mount** from the actions [⋮] menu.
When the mounting a partition, the partitioned file system becomes accessible and attaches it to the existing host directory structure.
- **Format a physical standard partition or logical volume partition:**

⚠ WARNING:

Formatting a partition deletes all the data and sets up a file system. Create a back up copy of any important data before wiping the data on the partition.

- a. In the **Partition** table, find a partition entry that you want to format and then select **Format** from the adjacent Actions [⋮] menu.
The **Format [partition file name]** dialog appears.
- b. In the **Format [partition file name]**, enter the following information and then click **Format**.

Name	In the Name text box, enter a name for the newly created file system.
Type	In the Type drop-down box, select a file system format type, for example: <ul style="list-style-type: none"> – XFS (recommended). The default high performance scalable file system format for Oracle Linux systems disk devices. For more details, see Oracle Linux XFS file system technical details https://www.oracle.com/linux/technologies/xfs-overview.html – EXT4. A scalable extension of the ext3 file system, which was the earlier default file system of Oracle Linux 5 – No File system. Data is saved as one large body of data with no way to tell where any piece of data is found or how to later view and retrieve it.
Overwrite	Select or clear the Overwrite checkbox. When selected, the deleted data is overwritten with zeros, making the deleted data unrecoverable.
Mount options	Select one or more mount point options.
Encryption	In the Encryption drop-down list box, select the appropriate encryption option. For more information, see Lock Disk Device With LUKs .

- **Delete an unmounted physical partition:**

 **WARNING:**

The partition must be unmounted before deleting it. Deleting a partition removes the allocated partition space and its data. Backup any disk partitioned data that you don't want to lose.

- a. In the **Partition** table, find the unmounted physical partition entry that you want to delete and then select **Delete** from the adjacent Actions [⋮] menu.
- b. A confirmation dialog appears indicating that partition and its data will be removed. Click **Yes** to continue the deletion process.
At the completion of the deletion process, both the partition data and the allocated space are removed.

Encrypt Block Devices With LUKS

Oracle Linux uses LUKs to perform block device encryption. By default, the option to encrypt a disk with LUKs is disabled at installation. If the encryption option at installation is enabled, the system prompts for a passphrase every time you boot or mount the device. The passphrase is an encryption key that decrypts the partition or volume and makes the file system accessible.

Post installation, Cockpit administrators can use the web console to change the encryption passphrase or to format a disk partition or logical volume with or without LUKs encryption. For more information, see these topics:

- [Lock Disk Device With LUKs](#)
- [Change Passphrase Key for LUKS Encryption](#)

Lock Disk Device With LUKs

Cockpit administrators can use the **Storage** page in the web console to format a partition or volume with LUKs encryption.

What Do You Need?

- The Cockpit web console must be installed and accessible.
For details, see these topics: [Install and Enable Cockpit](#) and [Log in to the Cockpit Web Console](#).
- The `cockpit-storaged` package must be installed.

 **Note:**

If the `cockpit-storaged` package isn't installed, see this section [Install and Manage Add-on Applications](#)

- Unmount all file systems on the device that you plan to encrypt

 **Note:**

You can re-encrypt encrypted devices while the devices are in use (change encryption key or algorithm) using the LUKS2 format. The LUKS1 format doesn't provide online re-encryption.

- Block storage device must have a file system.
- The disk name in which you want to encrypt a partition must be visible to Cockpit and appear in the Drives table on the Storage page.
- The volume name in which you want to encrypt a logical volume must be visible to Cockpit and appear in the Devices table on the Storage page.
- Backup the data on the partition or logical volume in which you want to encrypt using LUKS. Formatting a partition or volume deletes all the data and sets up a new file system.
- Administrator privileges.

Steps

Using the Cockpit web console, follow these steps to format and encrypt a host partition or logical volume with LUKS.

1. In the Cockpit navigation pane, click **Storage**.
The **Storage** page appears.
2. In the Storage page, perform one of the following:
 - **Select a drive from the Drives table:**
 - a. In the **Storage [model name]** page, navigate to the **Partitions** table.
 - b. In the **Partitions** table, find the partition you want to format and then select **Format** from the actions [⋮] menu.
 - OR-
 - **Select a volume from the Devices table:**
 - a. In the **Storage [volume group name]** page, navigate to the **Logical Volumes** table.
 - b. In the **Logical Volumes** table, find the volume entry that you want to format and then select **Format** from the actions [⋮] menu.
3. In the **Format** dialog, specify the following properties and then Click **Format**.

 **WARNING:**

Formatting deletes all the data and sets up a new file system.

Name	In the Name text box, enter a partition label to help users identify a partition.
-------------	--

Type	In the Type drop-down box, select a file system format type, for example: <ul style="list-style-type: none"> • XFS (recommended) – XFS is considered the high performance scalable file system format for Oracle Linux systems disk devices. • EXT4 – EXT4 is a scalable extension of the EXT3 file system. • No File system – A no file system format causes the system to save data as one large body of data with no way to tell where any piece of data is found or how to review and retrieve it.
Overwrite	Select or clear the Overwrite checkbox. When selected, the deleted data is overwritten with zeros, making the deleted data unrecoverable.
Mount options	Select one or more mount point options.
Encryption (LUKs)	In the Encryption drop-down list box, select one of the following options: <ul style="list-style-type: none"> • LUKs1 - LUKs1 provides compatible format for earlier release of Oracle Linux. • LUKs2 (recommended) - LUKS2 offers more flexible unlocking policies, stronger cryptography, and better compatibility with future enhancements. • No Encryption - Encryption protection isn't implemented.
Passphrase	In the Passphrase text box, specify a passphrase to be used to decrypt the partition and make the content accessible.
Confirm	In the Confirm text box, enter the passphrase that you entered in the Passphrase text box.
Store Passphrase	Select or clear the checkbox for Store Passphrase .
Encryption Options	In the Encryption Options text box, specify the required encryption options.

Change Passphrase Key for LUKS Encryption

Cockpit administrators can use the **Storage** page in the web console to change the LUKS passphrase key.

What Do You Need?

- The Cockpit web console must be installed and accessible. For details, see these topics: [Install and Enable Cockpit](#) and [Log in to the Cockpit Web Console](#).
- The `cockpit-storaged` package must be installed.

 **Note:**

If the `cockpit-storaged` package isn't installed, see this section [Install and Manage Add-on Applications](#)

- An unmounted LUKS1 formatted file system.

! Important:

You can re-encrypt encrypted devices while the devices are in use (change encryption key or algorithm) using the LUKS2 format. The LUKS1 format doesn't provide online re-encryption. In this case, devices encrypted with LUKS1 format might require you to unmount the file system to apply encryption property changes.

- Administrator privileges.

Steps

Using the Cockpit web console, follow these steps to change the LUKS primary or slot passphrases assigned to a host encrypted partition or logical volume.

1. In the **Storage** page, select a drive from the **Drives** table.
2. In the **Storage [model name]** page, navigate to the **Partitions** table.
3. In the **Partitions** table, find the row with the encrypted partition or volume and then click the down arrow icon to expand the table information. An **Encryption** tab appears.
4. Click the **Encryption** tab and configure the applicable passphrase properties as needed.

Stored passphrase

Click the **Edit** link to change the LUKs primary encryption passphrase.

Keys - passphrase

Important:

For information about configuring Tang server keys, see [Unlock Encrypted Devices Using Tang Server Key](#).

LUKS provides the ability for users to configure multiple passphrase keys per slots (up to 8 slots). Any one of the configured keys can open the encrypted partition.

Note:

LUKS encryption passphrases are stored in slots in the header of the partition.

Click the **Edit** icon to change an existing encryption passphrase that's assigned to a specific slot #.

-OR-

Click the plus [+] icon to add a new storage slot and assign a passphrase.

-OR-

Click the minus [-] icon to remove a passphrase storage slot configuration.

Unlock Encrypted Devices Using Tang Server Key

Using the **Storage** page in the web console, Cockpit administrators can automatically unlock an encrypted storage device using a key from a Tang server. For more information, see the following topics on how to create and confirm the implementation of a Tang Key on an encrypted device.

- [Create a Tang Key for Encrypted Device](#)
- [Confirm Tang Key Implementation on Encrypted Device](#)

Note:

The steps in this section are part of a wider task of implementing Policy-Based Decryption (PBD) by configuring Network-Bound Disk Encryption (NBDE) that features Tang and Clevis server and client components. For more information about implementing Policy-Based Decryption (PBD), see [Oracle Linux: Enabling Network-Bound Disk Encryption](#). For a tutorial in installing and configuring a Tang server, see [Use Network Bound Disk Encryption on Oracle Linux 8](#).

Create a Tang Key for Encrypted Device

For host systems that support a network-bound disk encryption (NBDE) environment, Cockpit administrators can use the **Storage** page to configure a **Tang server key address** and **passphrase**.

What Do You Need?

- The Cockpit web console must be installed and accessible. For details, see these topics: [Install and Enable Cockpit](#) and [Log in to the Cockpit Web Console](#).
- The `cockpit-storaged` package must be installed.

 **Note:**

If the `cockpit-storaged` package isn't installed, see this section [Install and Manage Add-on Applications](#)

- Configuration requirements for adding a Tang server key include:
 - A LUKS encrypted partition or volume on the host system.
 - A Tang server configuration on the host network.
 - A host with decryption client software (Clevis) installed.

 **Note:**

LUKS provides the facility to store multiple keys in different slots to decrypt data. System administrators can maintain the primary passphrase that locks a disk or a volume alongside the NBDE Tang server key. For host systems with LUKs encryption, users are prompted for a passphrase entry to gain access to data stored on the encrypted disk partition. For hosts systems with LUKs and an NBDE environment, user input isn't required to gain access to data stored on the encrypted partition or volume. The LUKs passphrase prompt closes automatically after the Tang key is decrypted and users are granted access to the data.

For more information about how to set up network-bound disk encryption, see [Oracle Linux: Enabling Network-Bound Disk Encryption](#) For more information about configuring LUKS encryption using the Cockpit web console, see [Encrypt Block Devices With LUKS](#).

- An unmounted LUKS1 formatted file system.

! Important:

You can re-encrypt encrypted devices while the devices are in use (change encryption key or algorithm) using the LUKS2 format. The LUKS1 format doesn't provide online re-encryption. In this case, devices encrypted with LUKS1 format might require you to unmount the file system to apply encryption property changes.

- Administrator privileges.

Steps

Using the Cockpit web console, follow these steps to add Tang key properties to an existing encrypted partition or volume configuration.

1. In the **Storage** page, select a drive from the **Drives** table.
2. In the **Storage [model name]** page, navigate to the **Partitions** table.
3. In the **Partitions** table, find the row with the encrypted partition or volume and then click the down arrow icon to expand the table information. An **Encryption** tab appears.
4. Click the **Encryption** tab and navigate to Key section.
5. In the **Keys** section, click the plus **[+]** icon to add a Tang key. The **Add Key** dialog appears.
6. In the **Add Key** dialog, specify the following information and then click **Add**.

! Important:

For information about configuring multiple passphrase keys for LUKS encryption, see [Change Passphrase Key for LUKS Encryption](#).

Key Source**Keyserver address**

Select the **Tang keyserver** radio button.

In the **Keyserver address** text box, specify the fully qualified domain name (FQDN) or IP address of the Tang server including the port number that the server uses, for example, `tangserver.example.com:7500`.

Note:

By default, the Tang server uses port 80. However, you can configure the server to use a different port number. For instructions, see [Oracle Linux: Enabling Network-Bound Disk Encryption](#)

Disk passphrase

In the **Disk passphrase** text box, if the LUKS passphrase isn't already entered, specify the current LUKS passphrase for the encrypted device.

The **Verify key** dialog appears displaying a generated key hash with instructions for verifying the hash key.

7. Perform these steps to verify the hash key:
 - a. Click **Terminal**. A host terminal window appears.
 - b. Obtain the key hash that the Tang server provided by typing the following command.

```
sudo curl -s tangserver.example.com:7500/adv | jose fmt -j- -g
payload -y -o- | jose jwk use -i- -r -u verify -o- | jose jwk
thp -i-
```

Verify that the key hash that's generated matches the key that's displayed in the Verify key window.

8. In the **Verify key** dialog, click **Trust key**.
9. Access the **Terminal** window from the Cockpit web console and enable early boot decryption.

```
sudo dracut -fv --regenerate-all
```

10. Confirm that the Tang server key configuration is successful, see [Confirm Tang Key Implementation on Encrypted Device](#)

Confirm Tang Key Implementation on Encrypted Device

Follow these steps to confirm that the Tang Key configuration applied to a host encrypted partition or volume was successfully implemented.

What Do You Need?

- Success completion of the Tang server key configuration on a encrypted host device. For details, see [Create a Tang Key for Encrypted Device](#)
- Administrator privileges.

Steps

1. In the Cockpit web console **Storage** page, select a drive from the **Drives** table.
2. In the **Storage [model name]** page, navigate to the **Partitions** table.
3. In the **Partitions** table, find the row with the encrypted partition or volume and then click the down arrow icon to expand the table information. An **Encryption** tab appears.
4. Click the **Encryption** tab and navigate to **Key** section.

5. In the **Keys** section, verify that the Tang server properties appear in the **Keys** list, for example:

```
Keys
Passphrase                               Slot 0
Keyserver: tangserver.example.com:7500   Slot 1
```

6. Perform the following steps to verify that the bindings are available for early boot:
 - a. In the Cockpit web console, click **Terminal** to access the terminal window.
 - b. Use the `lsinitrd` command to verify that the host Clevis bindings are available for early boot, for example:

```
sudo lsinitrd | grep clevis
```

Output similar to the following appears:

```
clevis
clevis-pin-sss
clevis-pin-tang
clevis-pin-tpm2
-rwxr-xr-x  1 root  root   1600 May 3 16:30 usr/bin/clevis
-rwxr-xr-x  1 root  root   1654 May 3 16:30 usr/bin/clevis-decrypt
...
-rwxr-xr-x  2 root  root    45 May 3 16:30 usr/lib/dracut/hooks/
initqueue/settled/60-clevis-hook.sh
-rwxr-xr-x  1 root  root   2257 May 3 16:30 usr/libexec/clevis-luks-
askpass
```

Manage Logical Volumes With LVM

Oracle Linux installations, by default, include Logical Volume Manager (LVM).

LVM is a storage management solution that enables administrators to combine any number of physical disk or partitions into a single file system or multiple file systems. The file system layouts are flexible, and can be used to resize allocated space. For example, administrators can choose to shrink a volume when space is no longer needed or grow a volume when more space is needed. Finally, administrators can mount a file system created on top of an LVM almost anywhere except for at `/boot`.

! Important:

For systems running AMD, Intel, or ARM architecture, the boot loader is unable to read LVM volumes. Therefore, the `/boot` partition, which contains the boot loader software to load the Linux kernel, should always remain as a standard disk partition and not become part of a LVM group.

For information on how to use the Cockpit web console to configure storage volumes managed by LVM, see these topics:

- [Create a Volume Group](#)

- [Create a Logical Volume](#)
- [Create a Thin Logical Volume](#)
- [Format and Mount a Logical Volume](#)
- [Resize Logical Volumes](#)
- [Change Volume Group Properties](#)

For information on how to manage LVM tasks from a command line, see [Configure Logical Volumes on Oracle Linux](#) or *Working With Logical Volume Manager* in [Oracle Linux 8: Managing Storage Devices](#) or [Oracle Linux 9: Managing Storage Devices](#).

Create a Volume Group

Using the **Storage** page, Cockpit administrators can create a volume group from one or more physical storage devices (hard drive, solid state drive, disk partition, raid device, and so on).

What Do You Need?

- The Cockpit web console must be installed and accessible. For details, see these topics: [Install and Enable Cockpit](#) and [Log in to the Cockpit Web Console](#).
- The `cockpit-storaged` package must be installed.

Note:

If the `cockpit-storaged` package isn't installed, see this section [Install and Manage Add-on Applications](#)

- Unused block devices in which you want to add to a volume group. Unused block devices are unformatted and unmounted storage devices.
- Administrator privileges.

Steps

Using the Cockpit web console, follow these steps to create a volume group.

1. In the **Storage** page, find the three-line menu icon in the **Devices** table and select **Create LVM2 volume group**. The **Create LVM2 volume group** dialog appears.
2. In the **Create LVM2 Volume Group** dialog, specify the following properties and then click **Create**.

Name	Enter a volume group name. For example, vg1 , vg2 , or provide a name that identifies the purpose of the volume group. For example, if you're creating the volume group to configure "thin" logical volumes, you might want to use a volume name such as <i>Thin Volume Group</i> . The LVM group name is limited to 128 characters.
-------------	--

Disks

In the **Disks** drop down list box, select the drives you want to include in the volume group.

 **Note:**

A message "No disks available" appears when the host system doesn't have an unallocated storage device (such as a partition with no file system) to add.

The name of the new volume group appears in the **Devices** table on the **Storage** page. You can click the **New volume group** to create logical volumes. For details, see [Create a Logical Volume](#).

Create a Logical Volume

Using the **Storage** page, Cockpit administrators can create a logical volume from a volume group.

What Do You Need?

- The Cockpit web console must be installed and accessible. For details, see these topics: [Install and Enable Cockpit](#) and [Log in to the Cockpit Web Console](#).
- The `cockpit-storaged` package must be installed.

 **Note:**

If the `cockpit-storaged` package is not installed, see this section [Install and Manage Add-on Applications](#)

- A created volume group on the host system. For details, see [Create a Volume Group](#)
- Administrator privileges.

Steps

Using the Cockpit web console, follow these steps to create a logical volume from a volume group.

1. In the **Storage** page, select the volume group in which you want to create a logical volume. The **LVM2 group [name]** page appears.
2. In the **LVM2 group [name]** page, navigate to the Logical volumes table and click **Create new logical volume**. The **Create logical volume** dialog appears.
3. In the **Create logical volume** dialog, specify the following properties and click **Create**.

Name

In the **Name** text box, enter a meaningful name (without spaces) for the new logical volume.

Purpose	In the Purpose drop down list box, select the applicable for option. For example: <ul style="list-style-type: none">• To create a logical volume for a file system, select Block device for filesystem.• To create a thin pool volume to add thin logical volumes, select Pool for thin logical volume.
Size	Use the Size slider to define the size of the logical volume. Consider the following: <ul style="list-style-type: none">• The amount of space required for the logical volume.• The number of logical volumes you plan to create from the specified volume group.

The logical volume is created and the name of the new logical volume appears in the **Logical Volumes** table.

Next Steps:

- If the selected **Purpose** of the logical volume was for *block device file system*, proceed to [Format and Mount a Logical Volume](#)

 **Note:**

In order for the logical volume to act as a physical disk, you must 1) create a file system by formatting the logical volume, and 2) attach the file system to the host by mounting it.

-OR-

- If the selected **Purpose** of the logical volume was to create a *thin volume pool*, you can add thin logical volumes to the newly created thin volume pool. For instructions, see Step 3 in [Create a Thin Logical Volume](#).

Create a Thin Logical Volume

Thin provisioning in LVM enables administrators to create larger logical volume(s) than the available physical disk space.

Creating a thin logical volume involves:

1. Create a volume group.
2. Create a thin volume pool from the volume group.
3. Add thin logical volumes to the thin volume pool.

What Do You Need?

- The Cockpit web console must be installed and accessible. For details, see these topics: [Install and Enable Cockpit](#) and [Log in to the Cockpit Web Console](#).
- The `cockpit-storaged` package must be installed.

 **Note:**

If the `cockpit-storaged` package isn't installed, see this section [Install and Manage Add-on Applications](#).

- Administrator privileges.

Steps

Using the Cockpit web console, follow these steps to create a thinly provisioned logical volume.

1. Create a **Volume Group**. For instructions, see: [Create a Volume Group](#).
2. Create a thin pool from the volume group created in Step 1. For instructions, see [Create a Logical Volume](#)

 **Important:**

In the Create a Logical Volume procedure, select *Pool for thin logical volume* as the Purpose.

3. Add thin logical volumes to the thin volume pool:
 - a. On the **Storage** page, click the volume group that you created in Step 1.

 **Note:**

The volume group name appears in the **Devices** table.

- b. In the **Storage [logical volume name]** page, click the thin pool volume that you created in Step 2.

 **Note:**

The thin volume appears in the Volumes table.

The **Thin volume** page appears.

- c. Click **Create thin volume**. The **Create thin volume** dialog appears.
- d. In the **Create thin volume** dialog, perform the following:
 - Define the size of the volume.
 - Click **Create** to create a thin volume.

The thin logical volume is created.

- e. Format and mount the thin logical volume. For instructions, see [Format and Mount a Logical Volume](#).

Format and Mount a Logical Volume

Logical volumes act as physical drives. A newly created logical volume must be formatted and mounted to act as a physical drive. Using the **Storage** page in the web console, Cockpit administrators can format and mount a logical volume.

What Do You Need?

- The Cockpit web console must be installed and accessible. For details, see these topics: [Install and Enable Cockpit](#) and [Log in to the Cockpit Web Console](#).
- The `cockpit-storaged` package must be installed.

 **Note:**

If the `cockpit-storaged` package isn't installed, see this section [Install and Manage Add-on Applications](#).

- A created logical volume on the host system. For details, see [Create a Logical Volume](#).
- Administrator privileges.

Steps


Using the Cockpit web console, follow these steps to format and mount the file system of an existing logical volume.

1. In the **Storage** page, find the **Devices** table and select the volume group in which the logical volume exists. The **LVM2 group [name]** page appears.
2. In the **LVM2 group [name]** page, navigate to the **Logical volumes** table.
3. In the **Logical volumes** table, find the volume entry that you want to format and select **Format** from the actions [⋮] menu. The **Format** dialog appears.
4. In the **Format** dialog, specify the following properties and click **Format**.

 **WARNING:**

Formatting a logical volume deletes all the data and sets up a file system.

Name	In the Name text box, enter a name for the volume file system.
-------------	---

Type	<p>In the Type drop-down box, select a file system format type, for example:</p> <ul style="list-style-type: none"> • XFS (recommended) – The XFS file system caters for large logical volumes, switching physical drives online without outage, and growing an existing file system. <div style="border: 1px solid #0070c0; padding: 10px; margin: 10px 0;"> <p> Note:</p> <p>The XFS file system doesn't support reducing the size of a logical volume. For more details, see Resize Logical Volumes.</p> </div> <ul style="list-style-type: none"> • EXT4 – A scalable extension of the ext3 file system. The EXT 4 file system caters for logical volumes, switching physical drives online without outage, and growing and shrinking a file system.
Overwrite	Select or clear the Overwrite checkbox. When selected, the deleted data is overwritten with zeros, making the deleted data unrecoverable.
Mount point and options	In the Mount Point text box, specify the mount point path. In the Mount options , select the applicable checkbox options.
Encryption	In the Encryption drop-down list box, select the appropriate encryption option. For more information, see Lock Disk Device With LUKs .

Formatting can take several minutes. When complete, the formatted logical volume appears in the **Logical volumes** table.

Resize Logical Volumes

Use the **Storage** page in the web console, Cockpit administrators can resize the storage size of a logical volume.

Note:

Logical Volume Manager (LVM) caters for the ability to increase the space of a logical volume and its file system while it's in use. The ability to reduce the space of a logical volume depends on the file system type. For example, an XFS file system doesn't support the ability to reduce the volume size.

What Do You Need?

- The Cockpit web console must be installed and accessible. For details, see these topics: [Install and Enable Cockpit](#) and [Log in to the Cockpit Web Console](#).
- The `cockpit-storaged` package must be installed.

 **Note:**

If the `cockpit-storaged` package isn't installed, see this section [Install and Manage Add-on Applications](#).

- When increasing a logical volume size, the following considerations apply:
 - A mounted active logical volume exists on the host.

 **Note:**

LVM caters for the ability to add space to a logical volume and its file system while the file system is mounted and in use.

- Free space must be available in the volume group in which the logical volume resides.
- When reducing a logical volume size, the following considerations apply:
 - The logical volume file system must support the ability to reduce the volume space. EXT 3 and 4 file systems support offline resizing when reducing the size. While XFS and GFS2 file systems don't allow you to reduce the volume size offline or online.
 - For EXT3 and EXT4 file systems, the logical volume must be unmounted before reducing its size.
 - An EXT3 or EXT4 file system must be reduced in size before the logical volume in which it resides can be reduced.
- Administrator privileges.

Steps

Using the Cockpit web console, follow these steps to resize a logical volume.

1. In the **Storage** page, find the table for **Devices** and select a volume group that contains the logical volume that you want to size.
The **LVM2 group [name]** page appears.
2. In the **LVM2 group [name]** page, navigate to the Logical volumes table.
3. In the **Logical volumes** table, find the volume entry that you want to resize and then click the down arrow to expand the volume information.
A **Volume** tab appears.
4. In the **Volume** tab, click **Grow** or **Shrink**. One of the following dialogs appears.

 **Note:**

An error message might appear for any of these reasons: (1) the volume group doesn't have free space available; (2) the logical volume content is unreadable; or (3) the logical file system doesn't support the ability for reducing the volume size.

- **Grow Logical Volume** dialog.

- a. Click and drag the **Grow** slider to define the increased volume size.
- b. Click **Grow**.
LVM grows the size of the logical volume space.
- **Shrink Logical Volume** dialog.
 - a. Click and drag the **Shrink** slider to define the reduced volume size.
 - b. Click **Shrink**.
LVM shrinks the size of the logical volume space.

Change Volume Group Properties

Cockpit administrators, as needed, can use the web console to change volume group properties. For more details, see the following topics:

- [Add New Drive to a Volume Group](#)
- [Remove Physical Drive From Volume Group](#)
- [Rename a Volume Group](#)
- [Rename a Volume Group](#)
- [Remove Logical Volume From Volume Group](#)

For further information about LVM, see *Working With Logical Volume Manager* in [Oracle Linux 8: Managing Storage Devices](#).

Add New Drive to a Volume Group

Using the **Storage** page in the web console, Cockpit administrators can add a new drive to a volume group.

What Do You Need?

- The Cockpit web console must be installed and accessible.
For details, see these topics: [Install and Enable Cockpit](#) and [Log in to the Cockpit Web Console](#).
- The `cockpit-storaged` package must be installed.

 **Note:**

If the `cockpit-storaged` package isn't installed, see this section [Install and Manage Add-on Applications](#)

- Unmounted and unpartitioned new disk drive attached to the host system. For instructions on how to partition a new disk or a disk with partitions", see [Create Physical Disk Partitions](#).
- An established volume group on the host system. For details on how to add a volume group, see [Create a Volume Group](#)
- Administrator privileges.

Steps

Using the Cockpit web console, follow these steps to add a new drive to an established volume group.

1. In the **Storage** page, find the **Devices** table and select the volume group in which you want to add a new drive.
The **LVM2 group [name]** page appears.
2. In the **LVM2 group [name]** page, navigate to the **Physical volumes** table.
3. In the **Physical volumes** table, click the plus [+] sign icon.
The **Add Disks** dialog appears.
4. In the **Add Disks** dialog, perform the following:
 - a. In the **Disks** drop-down list box, select the check boxes for the disk that you want to add.

 **Note:**

If LVM doesn't detect an unformatted drive attached to the host system, a message "No disks are available" appears.

- b. Click **Add**.
LVM adds the selected drives to the volume group and the space on the new drive is ready to be allocated to a logical volume within that volume. For example, in the that volume group, you can choose to increase the size of an existing logical volume (see, [Resize Logical Volumes](#)) or allocate the space to a new logical volume (see, [Create a Logical Volume](#)).

Remove Physical Drive From Volume Group

When a physical disk is no longer in use, Cockpit administrators can use the **Storage** page in the web console to remove the disk volume from the volume group.

What Do You Need?

- The Cockpit web console must be installed and accessible.
For details, see these topics: [Install and Enable Cockpit](#) and [Log in to the Cockpit Web Console](#).
- The `cockpit-storaged` package must be installed.

 **Note:**

If the `cockpit-storaged` package isn't installed, see this section [Install and Manage Add-on Applications](#)

- An existing volume group configuration on the host system that has more than one physical drive.

 **Note:**

You can't remove the last physical drive of a volume group.

For more information about removing physical volumes, see [Creating and Managing Volume Groups](#) in [Oracle Linux 9: Managing Storage Devices](#).

- A physical drive volume that isn't in use by any of the logical volumes in the volume group.
- Administrator privileges.

Steps

Using the Cockpit web console, follow these steps to rename a volume group.

1. In the **Storage** page, find the Devices table and select the volume group that contains an unused physical drive volume.
The **LVM2 group [name]** page appears.
2. In the **LVM2 group [name]** page, navigate to the **Physical volumes** table.
3. In the **Physical volumes** table, click the minus [-] sign icon next to the unused physical disk volume.
One of the following occurs:
 - If LVM detects it's the last disk volume in the volume group, a message appears indicating that it's the last disk volume in the volume group and it can't be removed.
-OR-
 - If LVM detects the disk volume is in use, a message appears indicating that the disk volume is in use, and it can't be removed.
-OR-
 - The **Remove Physical Volume** dialog appears.
4. In the **Remove Physical Volume** dialog, click **Remove**.
LVM removes the unused disk volume and shrinks the storage capacity of the volume group without loss of data.

Rename a Volume Group

Use the **Storage** page in the web console, Cockpit administrators can rename existing volume groups on the host system.

What Do You Need?

- The Cockpit web console must be installed and accessible.
For details, see these topics: [Install and Enable Cockpit](#) and [Log in to the Cockpit Web Console](#).
- The `cockpit-storaged` package must be installed.

 **Note:**

If the `cockpit-storaged` package isn't installed, see this section [Install and Manage Add-on Applications](#).

- An existing volume group configuration on the host system.
- Administrator privileges.

Steps

Using the Cockpit web console, follow these steps to rename a volume group.

1. In the **Storage** page, find the table for **Devices** and then select the volume group that you want to rename.
The **LVM2 volume group [name]** dialog appears.
2. In the **LVM2 Volume Group [name]** dialog, click **Rename**.
The **Rename volume group** dialog appears.
3. In the **Rename volume group** dialog, enter the new name, and then click **Rename**.

LVM renames the volume group and propagates the name change to all logical volumes in the volume group. Note that the renaming of the volume group doesn't affect the naming of thin pools.

Remove Volume Group

When a logical volume is no longer in use, Cockpit administrators can use the **Storage** page in the web console to remove an unused volume group.

What Do You Need?

- The Cockpit web console must be installed and accessible.
For details, see these topics: [Install and Enable Cockpit](#) and [Log in to the Cockpit Web Console](#).
- The `cockpit-storaged` package must be installed.

Note:

If the `cockpit-storaged` package isn't installed, see this section [Install and Manage Add-on Applications](#).

- Unused volume group without existing logical volumes. For information on how to remove a logical volume, see [Remove Volume Group](#).
- Administrator privileges.

Steps

Using the Cockpit web console, follow these steps to remove an unused volume group.

1. In the **Storage** page, find the table for **Devices** and then select the volume group you want to remove.
The **LVM2 group [name]** page appears.
2. In top section of the **LVM2 group [name]** page, click **Delete**.
The **Delete Volume Group** dialog appears.

 **Note:**

LVM prevents you from deleting the volume group if the volume group is in use or if it contains logical volumes.

3. In the **Delete Volume Group** dialog, click **Delete**.
LVM removes the unused volume group.

Remove Logical Volume From Volume Group

When a logical volume is no longer in use, Cockpit administrators can use the **Storage** page in the web console to remove a logical volume from a volume group.

What Do You Need?

- The Cockpit web console must be installed and accessible.
For details, see these topics: [Install and Enable Cockpit](#) and [Log in to the Cockpit Web Console](#).
- The `cockpit-storaged` package must be installed.

 **Note:**

If the `cockpit-storaged` package isn't installed, see this section [Install and Manage Add-on Applications](#).

- Unmounted and unused logical volume in a volume group.

 **WARNING:**

Removing a logical volume from a group, deletes all the data associated with the logical volume. To avoid losing critical data, back up the data residing on the logical volume.

- A deactivated logical volume if logical volume is part of a clustered environment.

 **Note:**

A clustered environment exists when volume groups are shared among multiple hosts.

 **Note:**

Logical volumes can be deactivated or activated from the Eclipse (3 vertical dots) menu in the Logical volumes table.

- Administrator privileges.

Steps

Using the Cockpit web console, follow these steps to remove an unused logical volume from a group.

1. In the **Storage** page, find the **Devices** table and select the volume group that contains an unused logical volume.
The **LVM2 group [name]** page appears.
2. In the **LVM2 group [name]** page, navigate to the Logical volumes table.
3. In the **Logical volumes** table, find the unused logical volume and then select Delete from the actions [⋮] menu.
The **Delete Logical Volume** dialog appears.
4. In the **Delete Logical Volume** dialog, click **Delete**.
LVM removes all the data on the logical volume, and it also removes the logical volume instance from the volume group.

Build and Manage Software RAID Devices

RAID stands for Redundant Array of Independent Drives. In a RAID configuration, two or more separate drives are combined to act as a single logical storage unit or drive. The benefits of a RAID is to keep data safe and to increase data access performance. In a RAID environment, data is distributed across volumes in different ways depending on the selected RAID level. A total of 6 different RAID levels are available, with each level offering different pros and cons for balancing speed and security.

Oracle Linux kernel uses the Multiple Device (MD) driver to support Linux software RAID configurations. Administrators requiring to build and manage software RAID devices on a host system can do so by using the Cockpit web console. For further details, see these topics.

- [Create and Configure a New Storage Array](#)
- [Software RAID Levels](#)

Note:

Alternatively, administrators can choose to use the command line to create and manage RAID devices. For further details about RAID configurations, see Working With Software RAID in [Oracle Linux 8: Managing Storage Devices](#).

Create and Configure a New Storage Array

Using the **Storage** page in the web console, Cockpit administrators can create a redundant array of independent disks.

Creating a storage array for system use involves the following process:

1. Create the RAID device.
2. Format the RAID device without partitions -OR- Partition the RAID device -OR- Format a RAID device part of a volume group.

3. Mount the RAID device.

What Do You Need?

- The Cockpit web console must be installed and accessible. For details, see these topics: [Install and Enable Cockpit](#) and [Log in to the Cockpit Web Console](#).
- The `cockpit-storaged` package must be installed.

 **Note:**

If the `cockpit-storaged` package isn't installed, see this section [Install and Manage Add-on Applications](#).

- Physical drives (HDDs or SSDs) connected to the host system. Note that each RAID level requires either two or more drives. To determine the number of drives required, review the [Software RAID Levels](#).
- Administrator privileges.

Steps

Follow these steps to create a storage array of independent disks on the host system.

1. In the **Storage** page, find the Devices table and then select **Create RAID device** from the hamburger icon menu (3 stacked horizontal lines). The **Create RAID device** dialog appears.
2. In the **Create RAID device** dialog, perform the following.
 - a. Specify the following properties:

Name	In the Name text box, specify a meaningful name to describe the storage array.
RAID level	In the RAID level drop down list box, select the applicable software RAID level. For more details, Software RAID Levels .
Chunk size	In the chunk size drop down list box, select the default predefined value (512 KiB).

 **Note:**

The **Chunk Size** value specifies block size for data writing. If the chunk size is 512 KiB, the system writes the first 512 KiB to the first disk, the second 512 KiB is written to the second disk, and the third chunk is written to the third disk. If you have three disks in the RAID configuration, the fourth 512 KiB is written to the first disk again.

Disks

In the **Disks** drop down list box, select checkboxes for the disks that you want to include in the RAID set.

 **Note:**

The number of disks required depends on the RAID level assigned.

- b. Click **Create**.
The newly created RAID device appears in the **Devices** table.
3. Perform one of the following to format and mount the newly created RAID device.
 - **Partition and mount the RAID device:**
 - Follow the steps in this procedure [Create Physical Disk Partitions](#).
 - OR-
 - **Format and mount the RAID device without partitions:**
 - a. In the **Storage** page, select the newly created RAID device from the **Devices** table.
The **Storage [RAID name]** page appears.
 - b. In the **Storage [RAID name]** page, in the **Content** table following the RAID device details click the newly created RAID device and then select **Format**.
The **Format** dialog appears.
 - c. In the **Format** dialog, specify the following properties and then click **Format**.

Name	In the Name text box, enter a name for the newly created RAID file system.
Type	In the Type drop-down box, select a file system format type, for example: <ul style="list-style-type: none"> – XFS (recommended) – The default high performance scalable file system format for Oracle Linux systems disk devices. – EXT4 – A scalable extension of the ext3 file system. – No File system – Data is saved as one large body of data with no way to tell where any piece of data is found or how to later view and retrieve it.

Overwrite

Perform one of the following:

- **Clear the `Overwritecheckbox`.** This option doesn't overwrite existing data, it only overwrites the header.

 **Note:**

Overwriting the header only is considered a faster process but less secure because the deleted data might be recoverable.

- **Select the `Overwritecheckbox`.** When selected, the deleted data is overwritten with zeros, making the deleted data unrecoverable.

 **Note:**

The Overwrite process is slower but is considered a more secure option as all the deleted data is overwritten with zeros.

Mount point and mounting options

Perform the following:

- i. In the **Mount point** text box, specify the path to mount the RAID device.
- ii. Select the check boxes for all applicable mounting options.

Encryption (if available) In the **Encryption** drop-down list box, select the appropriate encryption option. For more information, see [Lock Disk Device With LUKs](#).

Formatting can take several minutes depending on the file system type and RAID size specified.

After the format process completes, the file system details for the newly formatted RAID device appears in the **Filesystem** tab.

The RAID device is ready for system use (per the mounting options specified).

-OR-

- **Format and mount a RAID device as part of a volume group:**
 - a. Add the RAID device to a new volume group, see [Create a Volume Group](#).
 - b. Create a logical volume from the RAID volume group, see [Create a Logical Volume](#).
 - c. Format and mount the logical volume within the RAID volume group, see [Format and Mount a Logical Volume](#).

Software RAID Levels

The following software RAID levels are available.

RAID Level	Description	Disks Required	Redundancy
0	Striping	2-36	No
RAID- 1	Mirroring	2	Yes
RAID 0+1	Mirroring of striped disks	2-36	
RAID 1+0	Striping of mirrored disks or RAID-10	4-36 (even number only)	Yes
3	Striping with dedicated parity.	3-31	Yes
5	Striping with distributed parity.	3-31	Yes
6	striping with double distributed parity	2-36	Yes

For more information about software RAID levels, see [Working with Software RAID in Oracle Linux 8: Managing Storage Devices](#).

Manage NFS Mounted Connections

This section describes how to use the Cockpit web console to mount remote directories by using the Network File System (NFS) protocol.



Note:

NFS makes it possible to reach and mount files residing on storage devices across the network and work with them as if the files reside on the local physical drive.

- [Add NFS Server Connections](#)
- [Change NFS Server Connections](#)

Add NFS Server Connections

Use the **Storage** page in the Cockpit web console to mount a directory on a NFS server.

What Do You Need?

- The Cockpit web console must be installed and accessible. For details, see these topics: [Install and Enable Cockpit](#) and [Log in to the Cockpit Web Console](#).
- The `cockpit-storaged` package must be installed.



Note:

If the `cockpit-storaged` package isn't installed, see this section [Install and Manage Add-on Applications](#).

- NFS server host name or IP address.
- Path to the directory on the remote NFS server.
- Administrator privileges.

Steps

Using the Cockpit web console, follow these steps to mount a directory on an NFS server.

1. In the **Storage** page, find the table for NFS then click the plus [+] sign icon to add a NFS mount.

The New NFS Mount dialog appears.

2. In the **New NFS Mount** dialog, specify the following information and then click **Add**.

Server address	In the Server address text box, enter the NFS server host name or IP address.
Path on server	In the Path on server text box, type the path to a shared directory on the NFS server. Example: <code>/shares/nfs</code>
Local mount point	In the Local mount point text box, type the path to the directory location on the local system. Example: <code>/mounts</code>
Mount options	Select the applicable Mount option checkboxes.

Note:

To ensure that the shared directory is reachable after restarting the system, select the **Mount at boot** checkbox.

Change NFS Server Connections

Cockpit administrators can use the **Storage** page in the web console to change existing NFS mount properties. Configurable properties for unmounting, removing, and editing NFS server connections are available.

What Do You Need?

- The Cockpit web console must be installed and accessible. For details, see these topics: [Install and Enable Cockpit](#) and [Log in to the Cockpit Web Console](#).
- The `cockpit-storaged` package must be installed.

Note:

If the `cockpit-storaged` package isn't installed, see this section [Install and Manage Add-on Applications](#).

- An existing NFS shared directory connection on the host.

- The NFS server connection must be unmounted to delete or edit the NFS server configuration.
- Administrator privileges.

Steps

Using the Cockpit web console, follow these steps to define custom mount options for an existing unmounted NFS shared directory.

1. In the **Storage** page, find the table for NFS and select the NFS server that you want to edit.
The **Storage [NFS server name]** page appears.
2. In the **Storage [NFS server name]** page, perform one of the following:
 - **Unmount an inactive NFS server connection:**
 - a. Click **Unmount**. An **Unmount** dialog appears.
 - b. In the **Unmount** dialog, confirm the process to unmount the connection.

-OR-
 - **Delete an unmounted NFS server connection:**
 - a. Click **Delete**. A **Delete** dialog appears.
 - b. In the **Delete** dialog, confirm the process to delete the connection by clicking **Delete**.

-OR-
 - **Edit an unmounted NFS server connection:**
 - a. Click **Edit**. An **Edit** dialog appears.
 - b. In the **Edit** dialog, change the applicable properties and then click **Custom mount options** to specify properties that can help resolve NFS connection if an issue occurs.
To configure the custom mount option, enter following properties separated by a comma.

<code>nfsvers=<i>n</i></code>	where <i>n</i> equals the NFS protocol version number.
<code>soft</code> or <code>soft</code>	where <i>soft</i> or <i>hard</i> equals the recovery behavior after an NFS request times out. For more information, open a Terminal window and access the NFS manual pages by typing <code>man nfs</code> on the command line.
<code>sec=<i>krb5</i></code>	where <i>krb5</i> equals the kerberos authentication mode for securing the files on the NFS server.

- c. Click **Save**.

Manage Connections to iSCSI Targets

Cockpit administrators can use the **Storage** page in the web console to manage iSCSI storage target connections. The Storage page provides configurable options to add, mount, and remove an iSCSI storage connection.

What Do You Need?

- The Cockpit web console must be installed and accessible. For details, see these topics: [Install and Enable Cockpit](#) and [Log in to the Cockpit Web Console](#).
- The `cockpit-storaged` package must be installed.


 **Note:**

If the `cockpit-storaged` package isn't installed, see this section [Install and Manage Add-on Applications](#).

- – iSCSI target server IP address and user credentials (*username* and *password*).
- iSCSI host client initiator name.
- The `iscsid` service is started and enabled on host client system. For more information, see [Configure an iSCSI Target and Initiator on Oracle Cloud Infrastructure](#).
- Administrator privileges.

Steps

Using the Cockpit web console, follow these steps to add, mount, or remove an iSCSI target connection.

1. In the **Storage** page, find the table for iSCSI Targets and perform one of the following:
 - **Add and mount iSCSI storage connection:**
 - a. Click the pencil () icon to verify the correct initiator name is assigned to the host client.
The **Change iSCSI [initiator name]** dialog appears.
 - b. In the **Change iSCSI [initiator name]** dialog, change or keep the host client initiator name, and then click one of the following:
 - Change – To apply the changes made and dismiss the dialog.
 - OR-
 - Cancel – To dismiss the dialog without saving the changes.
 - c. Click the plus [+] icon to define the iSCSI server target IP address and user credentials.
The **Add iSCSI portal** dialog appears.
 - d. In the **Add iSCSI portal** dialog, enter the following information and then click Next.

Property	Description
Server	In the Server text box, enter the IP address of the iSCSI Linux server.
User name	In the Username text box, enter the username assigned to the administrator user account.

Property	Description
Password	In the Password text box, enter the secret phrase assigned to the administrator user account.

The Available targets on [server_ip_address] dialog.

- e. In the **Available target** dialog, select one of the storage targets listed and then click **Add**.
The **Storage [target name]** appears in the **Drives** table on the **Storage** page.
- f. In the **Drives** table, click the iSCSI storage target name.
The **Storage [target name]** page appears.
- g. In the **Storage [target name]** page, navigate to **Content** table and click **Mount** to mount the iSCSI file system on the host client.

--OR--

- **Remove existing iSCSI storage connection:**
 1. In the iSCSI target table, perform the following:
 - a. If more than one connection appears, select the name of the iSCSI storage connection that you want to remove.
 - b. Click the check mark icon above the connection name to disable the connection and to reveal the **Delete** icon.
 - c. Click the **Delete** icon to remove the configured iSCSI storage connection from the host.

5

Virtual Machine (VM) Management Tasks

Using the Virtual machines page in the web console, Cockpit administrators can create and manage KVM-based virtual machines on the host system.

Before creating and managing VMs on the managed host system, the `cockpit-machines` add-on application must be installed and the Oracle Linux Virtualization packages must be installed and enabled. For instructions, see [Install Cockpit-Machines and Enable Virtualization](#).

For more information about using Cockpit to create and manage VM instances on a host system, see the following topics:

- [Create, Import, Clone, or Migrate a VM Instance](#)
- [Start, Shutdown, Remove, or Interrupt a VM Instance](#)
- [Create, Delete, or Revert a Snapshot of VM Instance](#)
- [Manage Storage Pools for VM Instances](#)
- [Configure VM Devices and Services](#)

Create, Import, Clone, or Migrate a VM Instance

Using the **Virtual machines** page in the web console, Cockpit administrators can choose to create, import, clone, or migrate a VM instance as needed. For more information about using Cockpit to perform these tasks, see these topics:

- [Create a VM Instance](#)
- [Import a VM From a Disk Image](#)
- [Clone an Existing VM Instance](#)
- [Migrate a Running VM to Another KVM Host](#)

Create a VM Instance

Using the **Virtual Machines** page in the web console, Cockpit administrators can create virtual machine (VM) instances as needed.



Note:

Alternatively, administrators can create VM instances by using the Terminal CLI in the Cockpit web console. For command line instructions, see [Creating a New Virtual Machine](#).

What Do You Need?

- The Cockpit web console must be installed and accessible.

For details, see these topics: [Install and Enable Cockpit](#) and [Log in to the Cockpit Web Console](#).

- The Cockpit-machines add-on application must be installed in the web console and the Oracle Linux virtualization packages must be installed and enabled on the host system. For more information, see [Install Cockpit-Machines and Enable Virtualization](#).
- A minimum of one virtual storage pool must exist to create a virtual machine. Note that a virtualization storage pool is automatically provided in the `/var/lib/libvirt/images` directory. For information on how to add a new storage pool, see [Create a Storage Pool](#).
- A compatible guest OS is required to create a virtual machine. For more details, see *Guest Operating System Requirements* in the [Oracle Linux: KVM User's Guide](#).
- All minimum KVM system requirements must be met on the Cockpit host system to create, run, and manage VM instances. For more information, see *System Requirements and Recommendations* in the [Oracle Linux: KVM User's Guide](#).
- Internet access is required on host system when the selected installation type for installing a guest OS is **Download OS** or **URL**.

Steps

Using the Cockpit web console, follow these steps to create a virtual machine on the host system.

1. In the Cockpit navigation pane, click **Virtual machines**.
The **Virtual machines** page appears.
2. In the **Virtual machines** page, click **Create VM**.
The **Create new virtual machine** dialog appears.
3. In the **Create new virtual machine** dialog, perform the following:
 - a. Specify the applicable properties:

Name	In the Name text box, enter a unique name for the VM.
Details tab: Connection	Click either the System or User session radio button. For more details about these options, click the question mark [?] icon (next to the Connection property title).

Details tab: **Installation type**

In the **Installation type** drop down list box, select one of the following options:

- **Download an OS**
Internet access on host required to download OS.
- **Cloud Base Image**
You must have access to a Cloud base image on host system.
- **Local Install Media**
You must have access on host system to either ISO image or the distribution installation tree.
- **URL**
You must have Internet access on host system to specify a URL installation
- **Network Boot (PXE)**
You must have a PXE boot configuration available on host system.

Details tab: **Installation source**

 **Note:**

This property isn't applicable (not shown) when installation type is **Download OS**.

In the **Installation source** text box, specify the applicable path to the installation source.

Details tab: **Operating system**


In the **Operating system** list box, select the name of the guest OS.

Details tab: **Storage**

In the **Storage** list box, select one of the following options:

- **Create New Volume.** When selected, you can choose to accept the default storage size or specify a different storage size in MiBs or GiBs.
- **No Storage.**
- **Storage pool.** Select an existing storage pool from the list box and then choose a volume.

Details tab: **Storage limit**

 **Note:**

The **Storage Limit** property is only configurable when the **Create New Volume** option is selected.

Choose to accept the default storage limit properties or specify a different storage limit size.

Details tab: **Memory**

Choose to accept the default memory size property or change the memory size property as needed.

Automation tab: **Unattended install properties** (Optional)

 **Note:**

This option is only available for configuration when a **Cloud base image** installation type is specified.

If applicable, specify the following properties to deploy an unattended install:

- **Password** (optional). Enter a root password for the VM or leave this field blank. Click the question mark [?] icon for more details.
- **User Log-in** (required). Specify a username to enable an unattended installation.
- **User password** (optional). Enter a user account password or leave this field blank. Click the questions mark [?] icon for more details.

b. Click one of the following:

- **Create and Run.** Upon creating the VM instance, the VM guest OS installation starts automatically. The newly created VM appears in a Running state on the VM page.

 **Note:**

If the option for unattended installation is enabled, no user input is required. If the option for unattended installation is disabled, complete the installation by responding to the installation prompts displaying in the VM console. For more details, see the OS specific installation documentation.

- **Create and Edit.** Upon creating the VM instance, the VM guest OS installation is loaded and ready for installation. The VM state appears in a Shut off state on the VM page.
- c. In the VM page of the newly created VM, you can perform any of the following actions as needed:

Install

 **Note:**

Install appears only when the option for **Create and Edit** was selected to create the VM.

Click **Install** to run the guest OS installation program on the VM instance. If an unattended installation was enabled, the installation process proceeds without user input. If the unattended installation was disabled, respond to the installation prompts appearing in the VM console. For information on how to respond to the installation prompts, see the manufacturer's OS installation documentation.

Shutdown

 **Note:**

This option appears only when the VM guest is in a running or paused state.

Click **Shutdown** to exit the OS program on the VM instance.

Run

 **Note:**

Run appears after shutting down the guest installation. The VM instance appears in a shut off (inactive) state.

Click **Run** to start the guest OS program on the VM instance.

Actions menu [⌵]

Click actions [⌵] menu to select a VM management action.

Overview section

Navigate to the **Overview** section to view hypervisor details associated with the VM instance.

Configurable options are available for editing the CPU, Memory, AutoStart, and Watchdog properties.

For more information about changing these properties, see [Edit Memory, CPU, Autostart, or Watchdog Properties](#).

 **Note:**

The Firmware property is only editable after selecting the option to **Create and Edit**.

Console section

Navigate to the **Console** section to view and interact with VM guest OS.

Click the list box to select another console type (serial, VNC, and desktop viewer). For more information on how to set up the different consoles, see [Configure Console for VM Interaction](#).

Usage section

Navigate to the **Usage** section to view the VM CPU and memory resource consumption.

Device sections: **Disks, Network, Hosts**

Navigate to the applicable device section to view, add, or change a virtual device configuration associated with the VM instance.

For more information about these properties, see the following topics:

- [Add, Edit, Unplug, or Plug VM Network Interface](#)
- [Attach or Remove VM Host Devices](#)
- [Add, Edit, or Remove Disks](#)

Snapshot section

Navigate to the **Snapshot** section to take a snapshot of the VM instance.

The VM instance must be powered down to create a snapshot. For more information, see [Create, Delete, or Revert a Snapshot of VM Instance](#).

Shared directories section

Navigate to the **Shared directories** section to view and manage shared host directories mounted to the VM instance. Click the question mark icon [?] for information on how to manually mount a shared directory. For more details, see [Share a Host Directory with VM Instance](#).

 **Note:**

If the VM guest installation fails, remove the VM instance and then create another instance.

Install Cockpit-Machines and Enable Virtualization

Before Cockpit administrators can access and use the Virtual Machine functionality in the web console, the following tasks must be completed:

- Install the Cockpit-Machines add-on application in the web console.
- Install the Virtualization base packages on the Cockpit host system.

 **Note:**

The Virtualization base packages are typically installed on the host system during the Oracle Linux installation. In the case that these packages aren't installed, they must be installed to use the Virtual Machine module in the Cockpit web console.

- Verify the Virtualization service on the host system is started.

What Do You Need?

- The Cockpit web console must be installed and accessible.
For details, see these topics: [Install and Enable Cockpit](#) and [Log in to the Cockpit Web Console](#).

Steps

Follow these steps to ensure that the host system is properly configured with the Cockpit Virtual Machine add-on application and the Oracle Linux Virtualization packages.

1. In the Cockpit web console, click **Terminal**. The Terminal command-line interface page appears.
2. In the Terminal command-line interface, perform the following:
 - a. Install the Cockpit Virtual Machine add-on application. Type:

```
sudo dnf install cockpit-machines
```

- b. Install Oracle Linux Virtualization packages (if not already installed on host system).

For Oracle Linux 8 installations:

Type:

```
sudo dnf module install virt
sudo dnf install virt-install
virt-viewer
```

For Oracle Linux 9 installations:

Type:

```
sudo dnf group install  
"Virtualization Host"  
sudo dnf install qemu-kvm virt-  
install virt-viewer
```

For more information, see *Installing Virtualization Packages on an Existing System* in the [Oracle Linux: KVM User's Guide](#).

- c. Verify that the Virtualization service is started on host system. Type:

```
sudo virt-host-validate qemu
```

For more information, see *Validating the Host System* in the [Oracle Linux: KVM User's Guide](#).

Video Demonstration

The video demonstration and tutorial provided at <https://www.youtube.com/watch?v=daHQeCY13s8> might also be useful if you need more information on using Cockpit to create virtual machine.

You can also see the video demonstration at <https://www.youtube.com/watch?v=-Z3AwP2HPa4> for more information on setting up Cockpit to manage virtual machines:

Clone an Existing VM Instance

Cockpit administrators can use the **Virtual machines** page in the web console to clone existing virtual machine (VM) instances as needed. VM cloning is useful when deploying identical VMs to a group of users.



Note:

When a VM is cloned, the system makes a copy of the source VM XML configuration and its disk images. A new name is assigned to the newly cloned VM and the data stored on the cloned virtual disk is identical to the data stored on the source VM disk.

What Do You Need?

- The Cockpit web console must be installed and accessible. For details, see these topics: [Install and Enable Cockpit](#) and [Log in to the Cockpit Web Console](#).
- The Cockpit-machines add-on application must be installed in the web console and the Oracle Linux virtualization packages must be installed and enabled on the host system. For more information, see [Install Cockpit-Machines and Enable Virtualization](#).
- A source VM instance for cloning must exist on the host system.
- The source VM instance must **not** contain these configuration properties:

- A persistent network MAC address other persistent properties that could prevent the clone from working correctly.
- Sensitive data, such as SSH keys and password files.

For more information on how to prepare a VM instance for cloning, see *Cloning Virtual Machines* in the [Oracle Linux: KVM User's Guide](#).

- All minimum KVM system requirements must be met on the Cockpit host system to create, run, and manage VM instances. For more information, see *System Requirements and Recommendations* in the [Oracle Linux: KVM User's Guide](#).
- The source VM instance must be shut down before the cloning process. For more information, see [Start, Shutdown, Remove, or Interrupt a VM Instance](#).

Steps

Using the Cockpit web console, follow these steps to create a clone of an existing VM instance on the host system.

1. In the Cockpit navigation pane, click **Virtual machines**.
The **Virtual machines** page appears.
2. In the **Virtual machines** page, navigate to the VM instance that you want to clone and select **Clone** from the actions menu [⋮].
The **Create a clone VM based on [name]** dialog appears.
3. In the **Create a clone** dialog, perform one of the following:
 - Accept the system-assigned name for the clone VM and click **Clone**.
-OR-
 - Edit the name provided for the clone VM and click **Clone**.

A new VM instance is created based on the source VM configuration. The newly created VM appears in the **Virtual machine** list on the Virtual machines page.

Import a VM From a Disk Image

Using the **Virtual Machines** page in the web console, administrators can import other preconfigured VM instances on the host system to Cockpit for management using the disk file image. The Cockpit virtual machine import process preserves all guest properties associated with the imported disk image.

What Do You Need?

- The Cockpit web console must be installed and accessible.
For details, see these topics: [Install and Enable Cockpit](#) and [Log in to the Cockpit Web Console](#).
- The Cockpit-machines add-on application must be installed in the web console and the Oracle Linux virtualization packages must be installed and enabled on the host system.
For more information, see [Install Cockpit-Machines and Enable Virtualization](#).
- The import VM instance must be shut down before importing it.
- A VM disk image with a compatible guest OS is required.
For more details, see *Guest Operating System Requirements* in the [Oracle Linux: KVM User's Guide](#).

- All minimum KVM system requirements must be met on the Cockpit host system. For more information, see *System Requirements and Recommendations* in the [Oracle Linux: KVM User's Guide](#).

Steps

Using the Cockpit web console, follow these steps to import a host VM disk image to Cockpit for management.

1. In the Cockpit navigation pane, click **Virtual machines**. The **Virtual machines** page appears.
2. In the **Virtual machines** page, click **Import VM**. The **Import a virtual machine** dialog appears.
3. In the **Import a virtual machine** dialog, perform the following:
 - a. Specify the applicable properties:

Name	In the Name text box, enter a unique name for the imported VM guest image.
Connection	Click either the System or User session radio button. For more details about these options, click the question mark icon (next to the Connection property title).
Disk Image	In the Disk Image list box, specify the host file path to the disk image.
Operating system	In the Operating system list box, select the disk image OS name.
Memory	Choose to accept the default memory size provided or change the memory size property as needed.

- b. Click one of the following:
 - **Import and Run.** Upon completing the import operation the VM guest OS is started and the VM instance appears in a Running state. The VM [import VM name] page appears.
 - **Import and Edit.** Upon completing the import operation, the VM guest OS is loaded and the VM instance appears in a Shut down state. The VM [import name] page appears.

Note:

To change the state of the newly imported VM, see [Start, Shutdown, Remove, or Interrupt a VM Instance](#)

Note:

To view and interact with the newly imported VM, see [Configure Console for VM Interaction](#).

 **Note:**

If the import operation fails, see the virtualization host log files to help diagnose the issue.

Migrate a Running VM to Another KVM Host

Using the **Virtual machines** page in the web console, Cockpit administrators can migrate a live running virtual machine (VM) instance on the managed host to another KVM host system.

What Do You Need?

- The Cockpit web console must be installed and accessible. For details, see these topics: [Install and Enable Cockpit](#) and [Log in to the Cockpit Web Console](#).
- The Cockpit-machines add-on application must be installed in the web console and the Oracle Linux virtualization packages must be installed and enabled on the host system. For more information, see [Install Cockpit-Machines and Enable Virtualization](#).
- Internet access is required on the source and destination host systems.
- All requirements for migrating a virtual machine must be met. For more information, see [Requirements to Migrate Virtual Machine](#).

Steps

Using the Cockpit web console, follow these steps to migrate a live running VM instance on the managed host system to another KVM host system.

1. In the Cockpit navigation pane, click **Virtual machines**. The **Virtual machines** page appears.
2. In the **Virtual machines** page, navigate to the VM instance that you want to migrate and select **Migrate** from the actions menu [⋮]. The **Migrate VM to another host** dialog appears.
3. In the **Migrate VM to another host** dialog, specify the applicable properties and then click **Migrate**.

Destination URI

Enter the URI of the destination host

Duration

Specify the duration of the migration:

- **Permanent** - (default) Clear the checkbox to migrate the VM permanently. Permanent migration removes all the VM configuration from the source host.
 - **Temporary** - Select the checkbox to migrate the VM temporary. A temporary migration makes a copy of the VM on the destination host. This copy is deleted from the destination host when the VM is shut down. The original VM remains on the source host.
-

The selected VM instance on the managed host system is migrated to the specified (URI) destination host system.

Requirements to Migrate Virtual Machine

The following requirements must be met before performing a live migration of a virtual machine to another KVM host.

- The source and destination hosts are running.
- Ensure the following ports are open on the destination host.
 - Port 22 is needed for connecting to the destination host by using SSH.
 - Port 16509 is needed for connecting to the destination host by using TLS.
 - Port 16514 is needed for connecting to the destination host by using TCP.
 - Ports 49152-49215 are needed by QEMU for transferring the memory and disk migration data.
- The VM must be compatible with the CPU features of the destination host.
- The VM's disk image is accessible to the source host and the destination host.
- When migrating a running VM, the network bandwidth rate must be higher than the rate in which the VM generates dirty memory pages.

To obtain the VM dirty page rate before you start the live migration, using the terminal CLI perform the following:

1. Monitor memory generation rate for a short duration.

```
sudo virsh domdirtyrate-calc vm-name 30
```

2. After the monitoring process completes, obtain its results:

```
sudo virsh domstats vm-name --dirtyrate
Domain: 'vm-name'
dirtyrate.calc_status=2
dirtyrate.calc_start_time=200942
dirtyrate.calc_period=30
dirtyrate.megabytes_per_second=2
```

In this example, the VM is generating 2 MB of dirty memory pages per second. A live-migration of this VM on a network with a bandwidth of 2 MB/s or less prohibits the migration to proceed unless you pause the VM or decrease the VM's workload.

To ensure that the live migration completes successfully, the network bandwidth must be much greater than the VM's dirty page generation rate.

WARNING:

For VM instances with processing tasks that changes memory pages faster than the KVM can transfer them, such as heavy I/O load tasks, the recommendation is to avoid performing a live migration on these type of VM instances.

Manage Storage Pools for VM Instances

A *storage pool* is a quantity of storage set aside for virtual machines. By default, a storage pool is automatically available for use upon installing the Oracle Linux Virtualization packages.

Storage pools enable administrators to better organize their virtual machines. Cockpit administrators can use the Virtual machines page in the web console to create, activate, deactivate, or remove storage pools as needed.

For more information about using Cockpit to manage storage, see these topics:

- [Create a Storage Pool](#)
- [Manage Existing Storage Pools](#)

For more information about managing virtual machine storage pools, see *Working with Storage for KVM Guests* in the [Oracle Linux: KVM User's Guide](#).

Create a Storage Pool

Using the **Virtual machines** page in the web console, Cockpit administrators can create a storage pool as needed on the managed host system. By default, a storage pool is provided in the `/var/lib/libvirt/images` directory.



Note:

For information about how to create storage pools using the `virsh` commands, see [Working With Storage for KVM Guests](#).


What Do You Need?

- The Cockpit web console must be installed and accessible. For details, see these topics: [Install and Enable Cockpit](#) and [Log in to the Cockpit Web Console](#).
- The `Cockpit-machines` add-on application must be installed in the web console and the Oracle Linux virtualization packages must be installed and enabled on the host system. For more information, see [Install Cockpit-Machines and Enable Virtualization](#).

Steps

Using the Cockpit web console, follow these steps to create a storage pool on the managed host system.

1. In the navigation pane, click **Virtual machines**, and then on the Virtual Machines page click **Storage Pools**. The **Storage Pools** page appears.
2. In the **Storage Pools** page, click **Create Storage Pool**. The **Storage Pool** dialog appears.
3. In the **Storage Pool** dialog, specify the following properties and then click **Create**.

<p>Connection</p> <p>Name</p> <p>Type</p>	<p>Select either a System or User session storage pool connection.</p> <p>In the Name text box, enter a unique name for the storage pool.</p> <p>In the Type drop down list box, select one of the following storage pool types.</p> <ul style="list-style-type: none"> • File system directory Requires path on host's file system. • Network file system Requires path on host's file system, host name, and the directory on the server being exported. • ISCSI target Requires path on host's file system, host name, and iSCSI target IQN. • Physical disk device Requires path on host's file system, physical device disk on host, and format. • LVM volume group. Requires host name, iSCSI target IQN, and iSCSI initiator IQN. <p>For more information about each type of storage pool, see https://libvirt.org/storage.html.</p>
<p>Target path</p>	<div style="border: 1px solid #0070C0; padding: 10px; background-color: #E6F2FF;"> <p> Note:</p> <p>The Target path property isn't required for LVM volume group storage pool types.</p> </div>
<p>Start pool when host boots</p>	<p>In the Target path text box, specify the file path for the storage pool on the host system.</p> <p>Do one of the following:</p> <ul style="list-style-type: none"> • (Default) Select the checkbox to automatically start the storage pools upon powering on the host. • Clear the checkbox to prevent starting the storage pool upon powering on the host.

The newly added storage pool instance appears on the **Storage Pools** page in a deactivated state.

4. In **Storage Pools** page, navigate to the newly added storage pool and perform any of the following:
 - Click the [>] icon to view the new added storage pool properties.
 - Click **Activate** to activate the newly added storage pool.

Manage Existing Storage Pools

Using the **Virtual machines** page in the web console, Cockpit administrator can choose to activate, deactivate, or remove existing host storage pools as needed.

What Do You Need?

- The Cockpit web console must be installed and accessible.
For details, see these topics: [Install and Enable Cockpit](#) and [Log in to the Cockpit Web Console](#).
- The `Cockpit-machines` add-on application must be installed in the web console and the Oracle Linux virtualization packages must be installed and enabled on the host system.
For more information, see [Install Cockpit-Machines and Enable Virtualization](#).
- One or more storage pools must already exist on the host system.

Steps

Using the Cockpit web console, follow these steps to manage existing storage pools on the host.

1. In the navigation pane, click **Virtual machines**, and then on the Virtual Machines page click **Storage Pools**.
The **Storage Pools** page appears.
2. In **Storage Pools** page, navigate to the storage pool of interest and perform any of the following:
 - View storage pool properties: Click the > icon to view the storage pool properties.
 - Activate a deactivated storage pool: Click **Activate**.
 - Deactivate an activated storage pool: Click **Deactivate**.
 - Remove a storage pool: Select **Delete** from the actions menu [⌵].

Start, Shutdown, Remove, or Interrupt a VM Instance

Using the Virtual machines (VM) page in the web console, Cockpit administrators can choose to start, shutdown, reboot, pause, remove, or interrupt an existing VM instance on the managed host system.



What Do You Need?

- The Cockpit web console must be installed and accessible.
For details, see these topics: [Install and Enable Cockpit](#) and [Log in to the Cockpit Web Console](#).
- The `Cockpit-machines` add-on application must be installed in the web console and the Oracle Linux virtualization packages must be installed and enabled on the host system.
For more information, see [Install Cockpit-Machines and Enable Virtualization](#).
- One or more virtual machine instances must exist on the host system and appear in the Virtual machines page.

Steps

Using the Cockpit web console, follow these steps to remove a VM instance or to change its operating state.

1. In the navigation pane, click **Virtual machines**.
The **Virtual machines** page appears.
2. In the **Virtual machines** page, navigate to the name of the virtual machine of interest and click the actions menu [⋮], and then select one of the following:

Start	When the VM is shut down, you can click Start to start the VM.
Pause	When a VM is running, you can click Pause to suspend the VM state. In this state, the VM continues to consume system RAM, but the VM disk and network I/O processes are suspended.
Resume	When a VM is paused, you can click Resume to reset the guest VM state without invoking a shut down.
<div style="border-left: 2px solid #0070C0; border-right: 2px solid #0070C0; padding: 10px; background-color: #E6F2FF;"> <p> Note:</p> <p>Resuming a virtual machine doesn't apply pending VM configuration changes. Pending changes only take effect after a complete shut down and restart of the VM.</p> </div>	
Shut down or Force Shut down	<p>Perform one of the following:</p> <ul style="list-style-type: none"> • Click Shut down to exit the VM instance. • Click Force Shut down to forcibly shut down the VM instance as if you had pulled the power plug.
Reboot or Force Reboot	<div style="border-left: 2px solid #0070C0; border-right: 2px solid #0070C0; padding: 10px; background-color: #E6F2FF;"> <p> Note:</p> <p>VM must be in a running state.</p> </div> <p>Perform one of the following:</p> <ul style="list-style-type: none"> • Click Reboot to restart the VM guest. • Click Force Reboot to forcibly restart the VM guest if it appears to be unresponsive.
Send Non Maskable Interrupt (NMI)	When a VM instance is unresponsive, select Send (NMI) to promote the system to respond or shut down.

Delete**Requirements:**

- The VM instance must be shut down before removing it.
- Any existing snapshots of the VM instance must be removed before removing the VM instance.

Perform these steps:

- a. Click **Delete**.
- b. In the confirmation dialog, click **Delete** to remove a VM instance from the host system.

Related Information:

- [Start, Shutdown, Remove, or Interrupt a VM Instance](#)
- [Create, Delete, or Revert a Snapshot of VM Instance](#)

Configure Console for VM Interaction

On the VM *[name]* page, Cockpit administrators can configure up to three different consoles for VM interaction. These consoles include: VNC console, serial console, and a desktop viewer. All three consoles enable administrators to interact with the VM instance in the same manner as they would with a physical machine.

What Do You Need?

- The Cockpit web console must be installed and accessible. For details, see these topics: [Install and Enable Cockpit](#) and [Log in to the Cockpit Web Console](#).
- The Cockpit-machines add-on application must be installed in the web console and the Oracle Linux virtualization packages must be installed and enabled on the host system. For more information, see [Install Cockpit-Machines and Enable Virtualization](#).
- The virtual machine instance must be running on the host system.

Steps

Using the Cockpit web console, follow these steps to interact with the VM instance using one of three consoles.

1. In the navigation pane, click **Virtual machines**, and then click the name of the virtual machine of interest. The **Virtual machines [VM name]** page appears.
2. In **Virtual machines [VM name]** page, navigate to the **Console** section and then in the drop-down list box, select one of the following:

VNC console (default)

When selected, the graphical VNC console appears in the Console section. You can interact the VM instance from the VNC console window using your keyboard and mouse.

Perform any of the following as needed:

- Click **Expand** [] link to resize the VNC console window to full page.
- Click **Send Key** to pass key combinations on to the VM instance that would normally be blocked. For example, to send `Ctrl+Alt-Del`, select the `Ctrl+Alt-Del` key combination from the list.
- Click **Disconnect** or **Reconnect** as needed to disconnect the VNC console or to reconnect the VNC console.

Desktop viewer

Note:

The Desktop viewer is available with most OSes. Some browser extensions and plugins prevent the web console to open Virt Viewer.

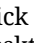
First time use: Install virt-viewer utility:

```
sudo dnf install virt-viewer
```

When selected, the Desktop viewer appears in the Console section.

You can interact the VM instance from the Desktop viewer using your keyboard and mouse.

Perform any of the following as needed:

- Click **Expand**  icon to resize the Desktop viewer window to full page.
- Click **Send Key** to pass key combinations on to the VM instance that would normally be blocked. For example, to send Ctrl+Alt-Del, select the Ctrl+Alt-Del key combination from the list box.
- Click **Disconnect** or **Reconnect** as needed to disconnect the Desktop viewer or to reconnect the Desktop viewer.

Note:

You can manually start the desktop viewer. Follow the Manual instructions provided in the Console section.

Serial console

Note:

The serial console option is useful when the host machine or the VM isn't configured with a graphical interface.

When selected, the serial console window appears in the Console section. You can interact the VM instance from the serial console window using your keyboard and mouse.

Perform any the following as needed:

- Click **Expand icon** [] to resize the serial console to full page.
- Click **Disconnect** or **Reconnect** as needed to disconnect the serial console or to reconnect the serial console.

Configure VM Devices and Services

Using the Virtual machines page in the web console, Cockpit administrators can manage devices and services associated with VM instances. For more information, see the following topics:

- [Edit Memory, CPU, Autostart, or Watchdog Properties](#)
- [Add, Edit, or Remove Disks](#)
- [Attach or Remove VM Host Devices](#)
- [Add, Edit, Unplug, or Plug VM Network Interface](#)

Edit Memory, CPU, Autostart, or Watchdog Properties

Using the **Vm [name]** page, Cockpit administrators can access and edit the Memory, CPU, Autostart, or Watchdog properties associated with a VM instance.

What Do You Need?

- The Cockpit web console must be installed and accessible. For details, see these topics: [Install and Enable Cockpit](#) and [Log in to the Cockpit Web Console](#).
- The Cockpit-machines add-on application must be installed in the web console and the Oracle Linux virtualization packages must be installed and enabled on the host system. For more information, see [Install Cockpit-Machines and Enable Virtualization](#).
- One or more virtual machine instances must already exist on the host system.

Steps

Using the Cockpit web console, follow these steps to view or edit the CPU, Memory, or Autostart properties associated with a VM instance.

1. In the navigation pane, click **Virtual machines**, and then click the name of the virtual machine of interest.
The **Virtual machines [VM name]** page appears.
2. In **Virtual machines [VM name]** page, navigate to the **Overview** section and view or edit any of the following configurable properties:

Memory

To view or edit the memory allocation properties for the selected VM instance, perform the following:

- a. Click the **Memory edit** link.
- b. In the Memory adjustment dialog, view, or edit the memory allocation properties as needed.
- c. To save changes, click **Save**, and then restart the VM guest to apply the saved changes.

Related Information:

- [Start, Shutdown, Remove, or Interrupt a VM Instance](#)
- [Configuring Memory Allocation in the Oracle Linux: KVM User's Guide.](#)

vCPUs

To view or edit the allocated virtual CPU count for the selected VM instance, perform the following:

- a. Click the **vCPU edit** link.
- b. In the vCPU details dialog, view, or edit the vCPU count allocation properties as needed.
- c. To save changes, click **Save**, and then restart the VM guest to apply the saved changes.

Related Information:

- [Start, Shutdown, Remove, or Interrupt a VM Instance](#)
- [Configuring Virtual CPU Count in the Oracle Linux: KVM User's Guide.](#)
- [Understanding CPU topology in Oracle Linux KVM in the Hard Partitioning with Oracle Linux KVM](#) <https://www.oracle.com/a/ocom/docs/linux/ol-kvm-hard-partitioning.pdf>

CPU type

To edit the **CPU type** property for the selected VM instance, perform the following:

- a. Click the **CPU edit** link.
- b. In the **CPU type** details dialog, set a virtual CPU mode that's compatible to the KVM host. For example:
 - Host.** When selected, this mode copies the model of the host CPU into the VM Guest definition. For more information, see the [libvirt upstream website](#).
 - Host-passthrough.** When selected, this mode presents the VM Guest with a CPU that's the same as the host CPU.
 - Custom.** Select an applicable custom mode that defines a normalized CPU that can be migrated throughout dissimilar hosts in a cluster.

 **Note:**

When performing a live VM migration, the CPU type must be compatible with the target KVM host.

- c. Click **Save** to save the changes, and then restart the VM guest to apply the saved changes. For information on how to restart the VM guest in Cockpit, see [Start, Shutdown, Remove, or Interrupt a VM Instance](#)

Autostart

To edit the **Autostart** property for the selected VM instance, perform the following:

- Enable or disable the **Autostart** option. When enabled, the VM guest OS automatically starts upon starting host system.

Watchdog

 **Note:**

This feature is only functional when the KVM stack fully caters for USB redirection or PCI passthrough.

Requirement: This option requires: 1) installation of watchdog device driver; and 2) the watchdog service enabled and started. For more information, see *Configuring the Watchdog Service* in the [Oracle Linux 8: Managing Core System Configuration](#) or in the [Oracle Linux 9: Managing Core System Configuration](#).

To specify Watchdog actions for the selected VM instance, perform the following:

- a. Click the **Watchdog Add** link.
- b. In the Watchdog dialog, specify the action to take if the system stops responding.
- c. Click **Add**.

Add, Edit, or Remove Disks

Using the **VM [name]** page, Cockpit administrators can manage virtual disks associated with a VM instance. Available disk management actions include: adding a new disk instance, editing properties associated with an attached disk instance, and removing a disk instance.

What Do You Need?

- The Cockpit web console must be installed and accessible. For details, see these topics: [Install and Enable Cockpit](#) and [Log in to the Cockpit Web Console](#).
- The Cockpit-machines add-on application must be installed in the web console and the Oracle Linux virtualization packages must be installed and enabled on the host system. For more information, see [Install Cockpit-Machines and Enable Virtualization](#).
- One or more virtual machine instances must already exist on the host system.
- For more information about managing virtual disks, see *Managing Virtual Disks* in the [Oracle Linux: KVM User's Guide](#).

Steps

Using the Cockpit web console, follow these steps to manage virtual disks associated with a VM instance.

1. In the Cockpit navigation pane, click **Virtual machines**. The **Virtual machines** page appears.
2. In the **Virtual machines** page, click the name of the virtual machine of interest. The **Virtual machines [VM name]** page appears.

3. In **Virtual machines [VM name]** page, navigate to the **Disk** section and perform any of the following actions:

<p>Add disk (new, existing, custom)</p>	<p>Perform these steps:</p> <ol style="list-style-type: none"> a. Click Add disk. b. In the dialog, select a source option and then configure the required properties: Source options: <ul style="list-style-type: none"> • Create. Select to add new virtual disk. • Existing. Select to attach existing virtual disk. • Custom. Select to add disk image or CD/DVD disk. Configuration properties: <ul style="list-style-type: none"> • Pool. In the list box, select the storage pool from which the virtual disk is to be attached. • Volume. In the list box, select a storage volume for the virtual disk to be attached • Persistence. This option is available only when the VM is running. To make the virtual disk persistent, select the Always attach check box. Otherwise clear the check box to make the virtual disk transient. • Cache. In the list box, select the applicable cache mode for the virtual network interface. • Bus. In the list box, select the applicable bus type for the virtual network adapter. c. Click Add to attach the disk to the VM instance.
<p>Edit an attached disk</p>	<p>Perform these steps:</p> <ol style="list-style-type: none"> a. Click Edit. b. In the dialog, change the applicable properties. c. Click Save to save the changes.
<p>Remove an attached disk</p>	<p>Perform these steps:</p> <ul style="list-style-type: none"> • Click Remove. • In the confirmation dialog, click Remove to detach the disk from the VM instance.

Attach or Remove VM Host Devices

Using the **VM [name]** page, Cockpit administrators can add host devices or remove them on a VM instance.

Host devices are physical devices connected to the managed host system, such as:

- SCSI tape drives, disks, and changers
- PCI NICs, GPUs, and HBAs
- USB mice, cameras, and disks

What Do You Need?

- The Cockpit web console must be installed and accessible.
For details, see these topics: [Install and Enable Cockpit](#) and [Log in to the Cockpit Web Console](#).
- The Cockpit-machines add-on application must be installed in the web console and the Oracle Linux virtualization packages must be installed and enabled on the host system.
For more information, see [Install Cockpit-Machines and Enable Virtualization](#).
- One or more virtual machine instances exist on the managed host system and appear in the Virtual machines page.
- The host device is directly attached and operational on the managed host system.

Steps

Using the Cockpit web console, follow these steps to add or remove a VM host device.

1. In the Cockpit navigation pane, click **Virtual machines**.
The **Virtual machines** page appears.
2. In the **Virtual machines** page, click the name of the virtual machine of interest.
The **Virtual machines [VM name]** page appears.
3. In **Virtual machines [VM name]** page, navigate to the **Host devices** section and perform any of the following actions:

Add host device

Perform these steps:

- a. Click **Add host device**.
- b. In the **Add** dialog, configure the following properties.
 - **Type:** Select **USB** or **PCI** to populate the Device product list.
 - **Device:** Select the check box for the host device you want to add.
- c. Click **Add**. The host device appears in the Host devices section.
- d. Restart the VM instance to detect the newly added host device.
For details, see [Start, Shutdown, Remove, or Interrupt a VM Instance](#).
Upon restarting the VM, the newly added device appears in the Host devices section of the VM instance page.

Troubleshooting: If an operational error appears when restarting the VM, you can fix the error by removing the host device from the virtual machine, or correcting the issues identified in the error message.

Remove

Optional Requirement:

- Before detaching a host device, consider using the `virsh dumpxml` command to create a backup of the VM XML configuration file:

```
virsh dumpxml VM_name >
VM_name.xml
cat VM_name.xml
<domain type='kvm'
xmlns:qemu='http://libvirt.org/
schemas/domain/qemu/1.0'>
  <name>VM_name</name>
  <uuid>ede29304-fe0c-4ca4-abcd-
d246481acd18</uuid>
  [...]
```

Perform these steps:

- Click **Remove**.
- In the confirmation dialog, click **Remove** to detach the host device from the VM instance.

Troubleshooting:

- If removing a host device causes the VM instance to become unbootable, use the `virsh define` utility to restore the backed up XML configuration.

```
virsh define VM_name.xml
```

- The removal of an attached USB host device might fail because of an incorrect correlation between the device and bus numbers of the USB device. As a workaround, edit VM's XML configuration file to remove the USB device entry [hostdev]. For example:

```
virsh edit VM_name_xml_file
```

Add, Edit, Unplug, or Plug VM Network Interface

Using the **VM [name]** page, Cockpit administrators can choose add, edit, unplug, or plug a virtual network interface as needed.

By default, all VM instances on the host are connected to a NAT mode virtual network environment. This default network (`libvirt0`) is available upon installing and enabling the Oracle Linux virtualization packages.

Note:

For more details about setting up and managing a virtual network, see *Setting up and Managing Virtual Networks* in the [Oracle Linux: KVM User's Guide](#).

What Do You Need?

- The Cockpit web console must be installed and accessible.

For details, see these topics: [Install and Enable Cockpit](#) and [Log in to the Cockpit Web Console](#).

- The Cockpit-machines add-on application must be installed in the web console and the Oracle Linux virtualization packages must be installed and enabled on the host system.
For more information, see [Install Cockpit-Machines and Enable Virtualization](#).
- One or more virtual machine instances must exist on the managed host system and appear in the Virtual machines page.

Steps

Using the Cockpit web console, follow these steps to configure, connect, or disconnect a virtual network interface.

1. In the Cockpit navigation pane, click **Virtual machines**.
The **Virtual machines** page appears.
2. In the **Virtual machines** page, click the name of the virtual machine of interest.
The **Virtual machines [VM name]** page appears.
3. In **Virtual machines [VM name]** page, navigate to the **Network Interface** section and perform any of the following actions:

Add Network Interface

Perform these steps:

- a. Click **Add Network Interface**.
- b. In the Add Network dialog, specify the required properties. For more details, see the configurable properties description that follows.
- c. Click **Add**.
The network interface is added and appears in the Network Interface section of the page.

Configurable properties:

- **Interface type:** In the list box, select one of the following: **Virtual** (default), **Bridge to LAN**, or **Direct attachment**.
Click question mark icon for a description of each option.
- **Source:** In the list box, select the applicable source destination for the network interface type.
- **Model:** In the list box, select the applicable model for the network interface type.
- **MAC Address:** Select to either have the MAC address for the interface automatically generated or to enter the MAC Address for the interface manually.
- **Persistence** (always attached): Clear the check box to disable this option, or select the check box to enable this option.

Remove Network Interface

Perform these steps:

- a. **Unplug** (disconnect) the network interface.
- b. Click **Remove** to delete the network interface configuration. The Network interface configuration is removed from the Network Interface section.

Unplug or Plug Network Interface

Perform one of the following:

- Click **Unplug** to disconnect the network interface configuration on the VM instance.
- OR-
- Click **Plug** to connect the network interface configuration on the VM instance.

Edit Network Interface**Requirement:**

- The VM instance must be shut down before editing the network interface configuration. For instructions, see [Start, Shutdown, Remove, or Interrupt a VM Instance](#).

Perform these steps:

- a. Click **Edit**.
- b. In the **Edit** dialog, change the applicable properties (type, source, model, MAC address).
- c. Click **Save** to save the change.
- d. Restart VM instance to apply the network changes.

Attach or Remove VM Host Devices

Using the **VM [name]** page, Cockpit administrators can add host devices or remove them on a VM instance.

Host devices are physical devices connected to the managed host system, such as:

- SCSI tape drives, disks, and changers
- PCI NICs, GPUs, and HBAs
- USB mice, cameras, and disks

What Do You Need?

- The Cockpit web console must be installed and accessible. For details, see these topics: [Install and Enable Cockpit](#) and [Log in to the Cockpit Web Console](#).
- The Cockpit-machines add-on application must be installed in the web console and the Oracle Linux virtualization packages must be installed and enabled on the host system. For more information, see [Install Cockpit-Machines and Enable Virtualization](#).

- One or more virtual machine instances exist on the managed host system and appear in the Virtual machines page.
- The host device is directly attached and operational on the managed host system.

Steps

Using the Cockpit web console, follow these steps to add or remove a VM host device.

1. In the Cockpit navigation pane, click **Virtual machines**.
The **Virtual machines** page appears.
2. In the **Virtual machines** page, click the name of the virtual machine of interest.
The **Virtual machines [VM name]** page appears.
3. In **Virtual machines [VM name]** page, navigate to the **Host devices** section and perform any of the following actions:

Add host device

Perform these steps:

- a. Click **Add host device**.
- b. In the **Add** dialog, configure the following properties.
 - **Type:** Select **USB** or **PCI** to populate the Device product list.
 - **Device:** Select the check box for the host device you want to add.
- c. Click **Add**. The host device appears in the Host devices section.
- d. Restart the VM instance to detect the newly added host device.
For details, see [Start, Shutdown, Remove, or Interrupt a VM Instance](#).
Upon restarting the VM, the newly added device appears in the Host devices section of the VM instance page.

Troubleshooting: If an operational error appears when restarting the VM, you can fix the error by removing the host device from the virtual machine, or correcting the issues identified in the error message.

Remove

Optional Requirement:

- Before detaching a host device, consider using the `virsh dumpxml` command to create a backup of the VM XML configuration file:

```
virsh dumpxml VM_name >
VM_name.xml
cat VM_name.xml
<domain type='kvm'
xmlns:qemu='http://libvirt.org/
schemas/domain/qemu/1.0'>
  <name>VM_name</name>
  <uuid>ede29304-fe0c-4ca4-abcd-
d246481acd18</uuid>
  [...]
```

Perform these steps:

- Click **Remove**.
- In the confirmation dialog, click **Remove** to detach the host device from the VM instance.

Troubleshooting:

- If removing a host device causes the VM instance to become unbootable, use the `virsh define` utility to restore the backed up XML configuration.

```
virsh define VM_name.xml
```

- The removal of an attached USB host device might fail because of an incorrect correlation between the device and bus numbers of the USB device. As a workaround, edit VM's XML configuration file to remove the USB device entry [hostdev]. For example:

```
virsh edit VM_name_xml_file
```

Create, Delete, or Revert a Snapshot of VM Instance

Using the **VM [name]** page, Cockpit administrators can create, delete, or revert a snapshot of the VM instance.

A *snapshot* is a copy of a virtual machine's OS and applications at a particular moment in time.

What Do You Need?

- The Cockpit web console must be installed and accessible. For details, see these topics: [Install and Enable Cockpit](#) and [Log in to the Cockpit Web Console](#).
- The Cockpit-machines add-on application must be installed in the web console and the Oracle Linux virtualization packages must be installed and enabled on the host system. For more information, see [Install Cockpit-Machines and Enable Virtualization](#).

- One or more virtual machine instances must exist on managed the host system and appear in the Virtual Machines page.

Steps

Using the Cockpit web console, follow these steps to manage snapshots associated with a VM instance.

1. In the Cockpit navigation pane, click **Virtual machines**.
The **Virtual machines** page appears.
2. In the Virtual machines page, click the name of the virtual machine of interest.
The **Virtual machines [VM name]** page appears.
3. In **Virtual machines [VM name]** page, navigate to the **Snapshot** section and perform any of the following actions:

Create

Perform these steps:

- a. Click **Create** to take a snapshot of the VM instance current state.
- b. In the dialog, specify the following:
 - **Name**. Optionally accept the name provided or change it.
 - **Description**. Optionally provide text about the snapshot.
- c. Click **Create**.
The name of the newly created snapshot appears in the Snapshot section of the VM instance page.

Revert

Caution:

Restoring a snapshot affects the virtual hard drives that are connected to the VM instance. This means that all files changes that occurred after creating the snapshot are lost during the restore process. To avoid losing files and other changes, create a snapshot of the revised VM instance before restoring the VM to an earlier state.

Perform these steps:

- a. Click **Revert**.
- b. In the dialog, confirm the restore operation by clicking **Revert**.

Remove

Requirement: VM instance must be running to remove Snapshot.
Perform these steps:

- Click **Remove**.
- In the confirmation dialog, click **Remove** to delete the selected snapshot.

Share a Host Directory with VM Instance

Using the **VM [name]** page, Cockpit Administrators can share files on the managed host system with a VM instance.

Note:

To share files between the host and a VM instance, Cockpit uses the `virtiofs` virtualization feature. *Virtiofs* is a shared file system feature that lets virtual machines access a directory tree on the managed host system.

What Do You Need?

- The Cockpit web console must be installed and accessible.
For details, see these topics: [Install and Enable Cockpit](#) and [Log in to the Cockpit Web Console](#).
- The Cockpit-machines add-on application must be installed in the web console and the Oracle Linux virtualization packages must be installed and enabled on the host system.

Note:

Virtiofs is available in most Linux distributions. It works with the virtualization packages (libvirt).

For more information, see [Install Cockpit-Machines and Enable Virtualization](#).

- The VM instance that you want to share host files with must have a Linux guest OS.
- The VM instance must be in a Shut off state to add shared directories.
- The host directory that you want to share must already exist.
- The shared host directory must be manually mounted inside the VM. For example:

```
mount -t virtiofs hostshare [mount point]
```

Where:

- mount point: The mount point inside the VM instance.

For more details, click the question mark icon [?] next to the Shared directories title on the VM [name] page.

Steps

Using the Cockpit web console, follow these steps to add a shared host directory to a VM instance.

1. In the Cockpit navigation pane, click **Virtual machines**.
The **Virtual machines** page appears.
2. In the **Virtual machines** page, click the name of the virtual machine of interest.
The **Virtual machines [VM name]** page appears.
3. In **Virtual machines [VM name]** page, navigate to the **Shared directory** section and click **Add shared directory**.
4. In the **Add shared directory** dialog, specify the following properties and then click **Share**.

Source path

In the **Source path** drop-down list box, specify the full host file path that you want to share with the VM instance.

Mount tag

In the **Mount tag** list box, specify a tag name that the VM instance uses to mount the source path.

Additional Options: **Extended attributes**

Select the **Extended attributes** checkbox to enable (`xattr`) on the shared directory and files.

or

Clear the **Extended attributes** checkbox to disable this option.

5. The newly created shared directory appears in the Shared directory section on the **VM [name]** page. You can now open the stored files on the shared host directory from the VM instance.

 **Note:**

You can remove an existing shared directory from the VM instance by clicking **Remove** on the VM page.

6

Podman Management Tasks

Using the Podman page in the web console, Cockpit administrators can monitor and manage containers, pods, and images on a host system. The Podman page provides configurable options that enable administrators to create containers and pods, and download Podman images. Options are also available on the Podman page for filtering container and pod view by owner (for example, user, system-wide, or all).

For further details about how to use the Podman management functionality from the Cockpit web console, see the following topics:

- [Install and Configure Cockpit-Podman](#)
- [Podman Image Management](#)
- [Podman Container Management](#)
- [Podman Pod Management](#)

Install and Configure Cockpit-Podman

Before Cockpit administrators can access and use the Podman functionality in the web console, the following tasks must be completed:

- Install the Cockpit- Podman add-on application in the web console.
- Verify that the Podman API socket service is enabled. If this service is disabled, start the Podman API socket service.
- Verify that the Podman proxy server settings on the host system are configured for use with the Cockpit web console service.

What Do You Need?

- The Cockpit web console must be installed and accessible.
For details, see these topics: [Install and Enable Cockpit](#) and [Log in to the Cockpit Web Console](#).
- The Podman container tools must be installed on the host system. For information on how to install Podman, see *Installing Podman and Related Utilities* in [Oracle Linux: Podman User's Guide](#).

Steps

Follow these steps to install the Cockpit-Podman add-on application, start the Podman API service if not already started, and to verify that the Podman proxy server settings on the host are configured for use with Cockpit.

1. In the Cockpit web console, click **Terminal**. The **Terminal shell** page appears.
2. Type the following to install the Cockpit-Podman add-on application in the web console.

```
sudo dnf install -y cockpit-podman
```

For more information about adding applications to the Cockpit web console, see [Install and Manage Add-on Applications](#)

3. In the Cockpit web console, click **Podman** and verify that Podman API service is running. For example, when the:
 - **Podman API service is enabled and running** – The Podman page displays the containers, pods, and images that are available on the local system.
 - **Podman API service is disabled** – The Podman page displays a warning indicating that the Podman service is disabled. In the warning dialog, perform the following steps to start the Podman API service.
 - a. Select **Automatically Start**.
When this option is checked, the Podman API service automatically starts upon each system power up.
 - b. Click **Start Podman**.
4. See the following documentation to ensure that the Podman proxy server settings on the host are configured for use with the Cockpit web console service. *Configuring Proxy Server Settings* in [Oracle Linux: Podman User's Guide](#).

Podman Image Management

Cockpit administrators, as needed, can easily view, inspect, and manage Podman images on the Podman page in the web console. For more details, see the following topics:

- [Search and Download New Images](#)
- [View and Inspect Available Images](#)
- [Remove Images](#)

Search and Download New Images

Using the **Podman** page, Cockpit administrators can search and download images from configured Podman registries on the host system.

What Do You Need?

- The Cockpit web console must be installed and accessible. For details, see these topics: [Install and Enable Cockpit](#) and [Log in to the Cockpit Web Console](#).
- The Podman container tools must be installed on the host system. For information on how to install Podman, see *Installing Podman and Related Utilities* in [Oracle Linux: Podman User's Guide](#).
- The Cockpit-Podman add-on application must be installed. For further information on how to configure the Cockpit web console to interact with Podman, see [Install and Configure Cockpit-Podman](#).

Steps

Using the Cockpit web console, follow these steps to search and download a copy of an image from a Podman configured registry on the host system.

1. In the Cockpit navigation pane, click **Podman**.
The **Podman** page appears.
2. In the **Podman** page, navigate to the upper right corner of the Images section, click the actions menu [⋮], and then select **Download new image**.
The **Search for an image** dialog appears.
3. In the **Search for an image** dialog, specify the following properties and then click **Download** to download a copy of the image to the local container image store.

Owner

Select one of the following:

- **System:** Download images as a system root user. Images are by default stored in the `/var/lib/containers` directory.
- **Username (logged in user):** Download images as a standard user. Images are typically stored in `$HOME/.local/share/containers/storage/` directory.

Search for

In the **Search for** section, perform the following:

- a. Type the name of an image in the first field, and then select the name of the target registry to conduct the search. A list of images appear below the search fields.
Example: Type Oracle Linux in the first field and select All registries in the second field. A list of all the Oracle Linux images found in all configured registries appear.
- b. In the list of images, select the image you want to download.

Tag

The *Tag* identifies the image version. If a tag name isn't specified, the default tag `latest` is appended to the image filename (*image:tag*).

 **WARNING:**

Many tools default to using the `latest` tag if no tag is specified but this can lead to errors and is now considered bad practice.

Recommended action:

- In the **Tag** field, enter a unique name that represents the version of the image. Example: If the image name is Oracle Linux and you know that the software version is 8, in the tag field you would enter 8 (`OracleLinux:8`).

For more information about tagging images, see *Oracle Linux Container Image Tagging Conventions* in [Oracle Linux: Podman User's Guide](#).

The **Search for image** dialog closes and the name of the downloaded image appears in the Images table on the **Podman** page.

View and Inspect Available Images

Using the **Podman** page, Cockpit users can easily view a listing of all Podman images stored on the host system. Clicking an image provides more details about an image such as image name, owner, creation date, size, and so on.

What Do You Need?

- The Cockpit web console must be installed and accessible. For details, see these topics: [Install and Enable Cockpit](#) and [Log in to the Cockpit Web Console](#).
- The Podman container tools must be installed on the host system. For information on how to install Podman, see *Installing Podman and Related Utilities* in [Oracle Linux: Podman User's Guide](#).
- The Cockpit-Podman add-on application must be installed. For further information on how to configure the Cockpit web console to interact with Podman, see [Install and Configure Cockpit-Podman](#).
- A copy of one or more images must appear in the Images section on the Podman page. For details on how to search and download images, see [Search and Download New Images](#).

Steps

Using the Cockpit web console, follow these steps to view and inspect available Podman images stored on the host system.

1. In the Cockpit navigation pane, click **Podman**.
The **Podman** page appears.
2. In the **Podman** page, use the following read-only properties to view and inspect Podman images appearing in the **Images** section.

Image totals

At the top of the **Podman** page the following Image total properties appear.

- **Total number of images:** Identifies the total number of images that are available for use on the local host.
- **Total number of unused images:** Identifies the total number of images that aren't use by a container.
- **Total number of used images:** Identifies the total number of images in use by a container.

Show and Hide Images (toggle switch)

Under the **Image Totals** property, a toggle link appears enabling users to show or hide the current list of images.

 **Tip:**

To find an image of interest, click Show image, and then use the Owner filtering options at the top of the page.

Listing of Images

The Listing of images appear in a table format. Clicking an image in the table displays the following properties about the image:

- **Image directory path:** Identifies the location of the image file relative to the registry host and the repository directory. For example:
`registry.host/repository/
imagename:tag`
- **Owner:** Identifies the owner of the image file.
- **Created:** Identifies the image file creation date.
- **ID:** Identifies the unique system ID assigned to the image file.
- **Disk space:** Identifies the size of the image file on the disk.
- **Used by:** Identifies whether the image file is or isn't in use by a container.

In addition, clicking the arrow next to the Image directory path displays the following information:

- **Details:** Identifies the image configuration details such as entry point, runtime command, and the exposed ports when available.
- **History:** Identifies the image file timestamp history.

Remove Images

Using the **Podman** page, Cockpit administrators can remove a single unused image or remove all unused images on a host system. Typically, images are removed to free up disk space or to download newer versions of an image.

Note:

Images can't be removed when associated with a Podman container. All container image dependencies must be removed before removing an image. For more information about deleting container images, see *Working With Container Images* in the [Install and Configure Cockpit-Podman](#).

What Do You Need?

- The Cockpit web console must be installed and accessible. For details, see these topics: [Install and Enable Cockpit](#) and [Log in to the Cockpit Web Console](#).
- The Podman container tools must be installed on the host system. For information on how to install Podman, see *Installing Podman and Related Utilities* in [Oracle Linux: Podman User's Guide](#).

- The Cockpit-Podman add-on application must be installed. For further information on how to configure the Cockpit web console to interact with Podman, see [Install and Configure Cockpit-Podman](#).
- A copy of one or more images must appear in the Images section on the Podman page. For details on how to search and download images, see [Search and Download New Images](#).

Steps

Using the Cockpit web console, follow these steps to remove a single unused image or all unused images on a host system.

1. In the Cockpit navigation pane, click **Podman**.
The **Podman** page appears.
2. In the **Podman** page, navigate to the Image section and expand the image listing (if not already expanded) by clicking Show images.
3. To remove unused images, perform the one of the following:
 - **Remove all unused images:**
 - a. Navigate to the upper top corner of the Podman page, click the actions menu [⋮] and then select **Prune unused images**.
The **Prune unused images** dialog appears.
 - b. In the **Prune unused images** dialog, click **Prune**.
The unused images are removed, and a refreshed image listing appears.
 - **Remove a single unused image:**
 - a. In the Images table, click the row that contains the image you want to delete, and then find the actions [⋮] menu in that same row and select **Delete**.
The **Delete [image file]** dialog appears.
 - b. In the **Delete [image file]** dialog, click **Delete tagged image**.
The unused image is removed, and a refreshed image listing appears.

Podman Container Management

Cockpit users can use the **Podman** page in the web console to monitor and manage all things to do with containers. For example, the Podman page provides update-to-date container performance details, container CLI interaction ability, and options to create, run, and change container instances. For more information about using Cockpit to perform container Podman management tasks, see the following topics:

- [Create and Run Container](#)
- [Special Considerations for Non Administrator Containers](#)
- [Inspect Container and Access Container Logs and CLI](#)
- [Rename, Pause, Stop, or Restart Container](#)
- [Commit Container Changes to Create New Image](#)
- [Checkpoint and Restore a System Container](#)
- [Remove Container or Pod Group](#)

Create and Run Container

Using the **Podman web console** page, Cockpit administrators can create and run containers with registry images. The Podman page provides different options to create a container. For example, administrators can create a container from either the image table, container table, or inside an existing pod group.

What Do You Need?

- The Cockpit web console must be installed and accessible. For details, see these topics: [Install and Enable Cockpit](#) and [Log in to the Cockpit Web Console](#).
- The Podman container tools must be installed on the host system. For information on how to install Podman, see *Installing Podman and Related Utilities* in [Oracle Linux: Podman User's Guide](#)
- The Cockpit-Podman add-on application must be installed. For further information on how to configure the Cockpit web console to interact with Podman, see [Install and Configure Cockpit-Podman](#).
- One or more registry images must exist in the Image table to create a container from an Image. For details on how to search and download registry images, see [Search and Download New Images](#).
- One or more pod groups must exist in the Container table to create a container inside a pod group. For instructions on how to create a pod group, see [Create a Pod Group](#).

Steps

Using the Cockpit web console, follow these steps to create or create and run a container.

1. In the Cockpit navigation pane, click **Podman**.
The **Podman** page appears.
2. In the **Podman** page, perform one of the following:
 - **Create a container from an image in the Image table:** Navigate to the Image table, find the row with the image that you want to use to create a container, and then in that table row, click **Create Container**.
The **Create Container** dialog appears.
 - OR-
 - **Create a container from the Container table:** Navigate to the **Container** table and click **Create Container**.
The **Create container** dialog appears.
 - OR-
 - **Create a container in a pod group.** Navigate to the **Container** table, find the row with the pod group that you want to add a container, and then in that table row click **Create a container in pod**.

 **Note:**

The options to Create a container in a pod group only appear when one or more pod groups exist. For information on how to create a pod group, see [Create a Pod Group](#)

The **Create container in [pod name group]** appears.

3. In the **Create container** dialog, perform the following
 - a. Specify the applicable properties:

Name

By default, a system generated container name appears in the **Name** text box. Choose to keep this name or remove it and specify a new name.

Details tab property: **Owner**

The following Owner options appear only for users with administrator or root privileges.

- **System:** Select to create a system ownership pod group.
- **Username** (logged in user): Select to create a local user ownership pod group.

 **Note:**

The local user ownership pod group is created by default for Cockpit users with limited access privileges. For more information about running pods or containers as a non-root user, see [Special Considerations for Non Administrator Containers](#).

Details tab property: **Image**

Use the **Image** list box to specify a registry image for the container.

For example:

If the **Create container** dialog is created from an image in the Image table, the name of the image automatically appears in the Image list.

If an image isn't already specified, perform the following to specify a registry image.

- Click the **Image** list box and select an image saved to cache.
-or-
- Type a search string in the **Image** drop down list box and then select one of the following search criteria: All, Local, Oracle Linux, or Docker. In the search results select the appropriate registry image.

Details tab property: **Command**

Use the **Command** text box to specify the applicable command to run the container image.
By default, the run command appears. If required, you can change the command.
Select the option **With Terminal** to run the container in a terminal.

Details tab: **Memory limit**

Use the **Memory limit** controls to specify the minimum memory allocated to run the container.
Optional:
Select the **Memory limit** checkbox and then using the controls specify a minimum memory allocation value.

Details tab: **CPU Shares**

 **Note:**

The **CPU Shares** property applies only to System container configurations.

CPU shares decide the priority for running containers by the amount of CPU shares allocated to the container. Default value: 1024

Optional:

Select the **CPU Shares** checkbox and then using the controls specify a CPU shares allocation value.

Details tab: Restart Policy

 **Note:**

The Restart Policy property applies only to System container configurations.

Select one of the following:

- **No** (default value): No action.
- **On Failure**: Restarts a container on failure.
- **Always**: Restarts container when exits or after system boot.

Integration tab: Port Mapping

Use the Port Mapping properties to set port mappings between the container and host system. Specifying port mapping exposes services running inside a host container.

To set port mappings, do the following:

- i. Click **Add Port Mapping**. The **Add Port Mapping** dialog appears.
- ii. Enter an IP address, host name, and container port.
- iii. Select a Protocol from the list.

For more information about configuring Port Mappings, see *Configuring Port Mappings for Containers* in [Oracle Linux: Podman User's Guide](#).

Integration tab property: Volumes

Use the Volume properties to share file system space on host system with container.

To configure the storage volume properties, do the following:

- i. Click **Add Volumes**. The **Add Volumes** dialog appears.
- ii. Enter a host path and container path.
- iii. (Optional) Select the **Writable** checkbox to create a writable volume.
- iv. In the **SELinux** drop down list, select one of the following options: No Label, Shared, or Private.

For more information about volumes, see *Using Volumes with Containers* in [Use storage with Podman containers](#).

For more information about configuring SELinux with Podman, see *Setting SELinux permission in Containers* in [Oracle Linux: Podman User's Guide](#).

Integration tab: Environment variables

Use the **Environment variable** properties for when you want to start a process inside the container.

To add variables, click **Add Variables**, and then enter a key and value.

For more information about the use of environment variables for container processes, see the Environment variables section in the Podman man-page.

Health tab:

Use the **Health** properties to help monitor the health or readiness of a process running in a container. To set container health check properties in the Health tab, do the following:

- i. In the **Command** text box, enter the `healthcheck run command`.
- ii. Set the following health check properties: Interval (30 second default), Timeout (30 second default), Start period, and Retries (3 default).

For more information about setting container health properties, see the `podman-healthcheck-run(1)` manual page.

- b. Click one of the following options:
 - **Create and Run.** The **Create and Run** option creates the container, starts the container image, and lists the active container in the **Container** table as Running.
 - -or-
 - **Create.** The *Create* option creates the container and lists the container in the **Container** table as Created.

 **Note:**

You can later run a created container from the **Container** table by selecting **Start** from the actions [] menu.

Special Considerations for Non Administrator Containers

Review the following special considerations when you're running containers as a non administrator:

- The storage path for the host container is different for root users (`/var/lib/containers/storage`) and non administrator users (`HOME/.local/share/containers/storage`).
- non administrators running containers are provided special permission to run as a range of user and group IDs on the host system. However, they have no root privileges to the host OS.
- In cases where a non administrator needs to change the `/etc/subuid` or `/etc/subgid` manually, the changes take effect only after issuing the `podman system migrate` command.
- Some system features are uneditible by non administrators. For example, non administrators are unable to change the system clock by setting a `SYS_TIME` capability inside a container and running the network time service (`ntpd`).
- A non administrator container is unable to access a port numbered less than 1024.

Inspect Container and Access Container Logs and CLI

Using the **Podman** page, Cockpit users can inspect container configuration details, log files, and interact with the container CLI as needed.

What Do You Need?

- The Cockpit web console must be installed and accessible. For details, see these topics: [Install and Enable Cockpit](#) and [Log in to the Cockpit Web Console](#).
- The Podman container tools must be installed on the host system. For information on how to install Podman, see *Installing Podman and Related Utilities* in [Oracle Linux: Podman User's Guide](#)
- The Cockpit-Podman add-on application must be installed. For further information on how to configure the Cockpit web console to interact with Podman, see [Install and Configure Cockpit-Podman](#).
- One or more containers or pod groups must already exist in the Container table. For information on how to create a container or a pod group, see [Create and Run Container](#) or [Create a Pod Group](#).
- Administrators or root users can access and change all containers and pods groups. Users with limited access privileges can access and change only the containers and pod groups that they created.

Steps

Using the Cockpit web console, follow these steps to view container details, generated log files, and gain access to the container CLI.

1. In the Cockpit navigation pane, click **Podman**. The **Podman** page appears.
2. In the **Podman** page, navigate to the **Container** table.
3. In the **Container** table, find the container that want you to view and then perform the following:
 - a. In the container row, view the following properties:
 - Container name.
 - Container owner (system or user)
 - Container usage properties for CPU and memory.
 - Container state (created, running, stopped, and so on)
 - b. Click the arrow icon next to container name to view further details about the container. The row expands with the following properties:

Details tab

Use the **Details** tab to view the container ID, creation date, image file path, runtime command, and the container process state. In addition, the following properties are viewable for system containers: IP address, MAC address, and Gateway address.

Integration tab	User the Integration tab to view the environment variables, port mappings, and configured volumes associated with the container.
Logs tab	Use the Logs tab to view the log files associated with the container.
Console tab	Use Console tab to display and interact with the container CLI.

Rename, Pause, Stop, or Restart Container


Use the **Podman** page in the Cockpit web console to rename a container or change its operating state.

What Do You Need?

- The Cockpit web console must be installed and accessible. For details, see these topics: [Install and Enable Cockpit](#) and [Log in to the Cockpit Web Console](#).
- The Podman container tools must be installed on the host system. For information on how to install Podman, see *Installing Podman and Related Utilities* in [Oracle Linux: Podman User's Guide](#)
- The Cockpit-Podman add-on application must be installed. For further information on how to configure the Cockpit web console to interact with Podman, see [Install and Configure Cockpit-Podman](#).
- One or more containers or pod groups must already exist in the Container table. For information on how to create a container or a pod group, see [Create and Run Container](#) or [Create a Pod Group](#).
- Administrators and root users can access and change all containers. Users with limited access privileges can access and change only the containers they created.

Steps

Using the Cockpit web console, follow these steps to rename a container or change the state of a container.

1. In the Cockpit navigation pane, click **Podman**.
The **Podman** page appears.
2. In the **Podman** page, navigate to the **Container** table.
3. In the **Container** table, find the container row of interest and then click the actions  menu in that same row and select one of the following:
 - **Rename**. The **Rename** dialog appears enabling you to rename the container.
 - OR-
 - **Start, Stop, Restart, or Pause**. Select one of the applicable actions to change the current operating state of the container.

Commit Container Changes to Create New Image

Use the **Podman** page in the Cockpit web console to commit container changes to a new image.

What Do You Need?

- The Cockpit web console must be installed and accessible. For details, see these topics: [Install and Enable Cockpit](#) and [Log in to the Cockpit Web Console](#).
- The Podman container tools must be installed on the host system. For information on how to install Podman, see *Installing Podman and Related Utilities* in [Oracle Linux: Podman User's Guide](#)
- The Cockpit-Podman add-on application must be installed. For further information on how to configure the Cockpit web console to interact with Podman, see [Install and Configure Cockpit-Podman](#).
- One or more containers or pod groups must already exist in the Container table. For information on how to create a container or a pod group, see [Create and Run Container](#) or [Create a Pod Group](#).
- Administrator or root users can access and change all containers. Users with limited access privileges can access and change only the containers they created.

Steps

Using the Cockpit web console, follow these steps to commit a new image based on the container changes.

1. In the Cockpit navigation pane, click **Podman**.
The **Podman** page appears.
2. In the **Podman** page, navigate to the **Container** table.
3. In the **Container** table, find the container row of interest and then in that same row select **Commit** from the actions [⋮] menu.
The **Commit** dialog appears.
4. In the **Commit** dialog, specify the following properties and then click **Commit**.

New Image Name	Enter a name for the new image.
Tag (Optional)	Enter information to describe the image.
Author (Optional)	Enter the author's name that submitted the changes.
Command (Optional)	Keep or change the runtime command
Options (Optional)	Select any of the following options that apply: <ul style="list-style-type: none"> • Pause container when creating image: When selected, the container, and its processes are paused while the image is committed. • Use legacy Docker format: When selected, the Docker image format is used. Otherwise, the OCI format is used.

The newly created image appears in the **Image** table.

Checkpoint and Restore a System Container

Use the **Podman** page in the Cockpit web console to set a checkpoint on a system container and save its state to disk. After creating a checkpoint and rebooting the system, you can then as needed, restore the state of the system container to an earlier checkpoint in time.

What Do You Need?

- The Cockpit web console must be installed and accessible. For details, see these topics: [Install and Enable Cockpit](#) and [Log in to the Cockpit Web Console](#).
- The Podman container tools must be installed on the host system. For information on how to install Podman, see *Installing Podman and Related Utilities* in [Oracle Linux: Podman User's Guide](#)
- The Cockpit-Podman add-on application must be installed. For further information on how to configure the Cockpit web console to interact with Podman, see [Install and Configure Cockpit-Podman](#).
- One or more running containers must already exist in the Container table. For information on how to create a container or a pod group, see [Create and Run Container](#).
- Users must have administrator or root privileges to checkpoint and restore a system container.

Steps

Using the **Cockpit** web console, follow these steps to create and restore a checkpoint on a system container.

1. In the Cockpit navigation pane, click **Podman**.
The **Podman** page appears.
2. In the **Podman** page, navigate to the **Container** table.
3. In the **Container** table, find the row with the running system container of interest and then in that same row select **Checkpoint** from the actions [⋮] menu.
The **Checkpoint** dialog appears.
4. In the **Checkpoint** dialog, select any of the following properties that apply and then click **Checkpoint**.

Keep all temporary checkpoint files. (Optional)	When selected, this option keeps all CRIO created logs and statistical data.
Leave running after writing checkpoint to disk. (Optional)	When selected, this option leaves the container running after the checkpoint process completes.
Support preserving established TCP connections. (Optional)	When selected, this option preserves the current container TCP property connections.

The checkpoint state for the running container is saved to disk.

 **Note:**

The system container can be restored to the saved checkpoint state after a reboot.

5. To restore the checkpoint container, perform the following steps:
 - a. If the system didn't reboot after creating the checkpoint, do the following:
 - i. Manually reboot the system.
 - ii. Access the Cockpit web console and open the **Podman** page.
 - b. In the **Podman** page of the Cockpit web console, navigate to the row with the checkpoint container and select **Restore** from the actions [⋮] menu. The **Restore** dialog appears.
 - c. In the **Restore** dialog, specify any of the following properties that apply and then click Restore.

Keep all temporary checkpoint files.
(Optional)

When selected, this option saves all CRIU temporary logs and properties created during checkpoint process.

Support preserving established TCP connections. (Optional)

When selected, this option reserves the current container TCP property connections.

Ignore IP address if set statically.
(Optional)

This option applies only when port mappings were configured. When selected, this option tries to use the same IP address that was used earlier to start the container.

Ignore MAC address if set statically.
(Optional)

When selected, this option tries to use the same MAC address that was used earlier to run the container.

After the system container is restored to the checkpoint state, it appears in the **Container** table as a running container.

 **Note:**

In the case that the checkpoint operation fails, the temporary files, and their properties remain available for debugging purposes.

Remove Container or Pod Group

Using the **Podman** page, Cockpit users can choose to remove a single container or a pod group of containers.

What Do You Need?

- The Cockpit web console must be installed and accessible.

For details, see these topics: [Install and Enable Cockpit](#) and [Log in to the Cockpit Web Console](#).

- The Podman container tools must be installed on the host system. For information on how to install Podman, see *Installing Podman and Related Utilities* in [Oracle Linux: Podman User's Guide](#)
- The Cockpit-Podman add-on application must be installed. For further information on how to configure the Cockpit web console to interact with Podman, see [Install and Configure Cockpit-Podman](#).
- Stop the container to be remove. For more details, see [Rename, Pause, Stop, or Restart Container](#).
- Users with limited access privileges can only remove containers and pod groups they created. For more details, see [Special Considerations for Non Administrator Containers](#)

Steps

Using the Cockpit web console, follow these steps to remove a single container or the containers associated with a pod group.

1. In the Cockpit navigation pane, click **Podman**.
The **Podman** page appears.
2. In the **Podman** page, navigate to the **Container** table.
3. In the **Container** table, perform any of the following:
 - **Remove a single container:** Find the row with the stopped container and select Delete from actions [⋮] menu.
 - a. Find the row with the stopped container.
 - b. In the stopped container row, select **Delete** from the actions [⋮] menu.
A dialog appears confirming that you want to delete the container.
 - c. In the dialog, click **Delete**.
 - **Remove a container in a pod group.**
 - a. Find the pod group row with a stopped container.
 - b. In the stopped container row, select **Delete** from the actions [⋮] menu.
A dialog appears confirming that you want to delete the container.
 - c. In the dialog, click **Delete**.
 - **Remove a pod group and all containers.**
 - a. Find the row with the stopped pod group.
 - b. In the row with the stopped pod group, select **Delete** from actions [⋮] menu.
A dialog appears confirming that you want to delete the pod group and all the containers in that group.
 - c. In the dialog, click **Delete**.

Podman Pod Management

Pods consist of one or more containers that share the same network communication settings, namespace, and service processes. Cockpit users can use the **Podman**

page in the web console to create and maintain pods. For example, the **Podman** page provides configurable options for creating pods, adding containers to pods, stopping pods, or starting pods. For more information about using Cockpit to perform pod management tasks, see the following topics:

- [Create a Pod Group](#)
- [Add a Container to a Pod Group](#)
- [Inspect and Change a Pod Group](#)

Create a Pod Group

Using the Podman web console page, Cockpit users can create an empty pod group. After creating a pod group, users can add and manage containers within the pod group.

What Do You Need?

- The Cockpit web console must be installed and accessible. For details, see these topics: [Install and Enable Cockpit](#) and [Log in to the Cockpit Web Console](#).
- The Podman container tools must be installed on the host system. For information on how to install Podman, see *Installing Podman and Related Utilities* in [Oracle Linux: Podman User's Guide](#)
- The Cockpit-Podman add-on application must be installed. For further information on how to configure the Cockpit web console to interact with Podman, see [Install and Configure Cockpit-Podman](#).

Steps

Using the Cockpit web console, follow these steps to create a pod group.

1. In the Cockpit navigation pane, click **Podman**.
The **Podman** page appears.
2. In the **Podman** page, navigate to the **Container** table and click **Create pod**.
The **Create pod** dialog appears.
3. In the **Create pod** dialog, specify the following properties and then click **Create**.

Name	By default, a system provided pod name appears in the Name text box. Choose to keep this name or remove it and specify a new name.
-------------	---

Owner

The following Owner options appear only for users with administrator or root privileges.

- **System:** Select to create a system ownership pod group.
- **User:** Select to create a local user ownership pod group.

 **Note:**

The local user ownership pod group is created by default for Cockpit users with limited access privileges. For more information about running pods or containers as a non-root user, see [Special Considerations for Non Administrator Containers](#).

Port Mapping (optional)

Use the Port Mapping properties to set port mappings between the pod group containers and host system. Specifying port mapping exposes services running inside a host container.

To set port mappings, do the following:

- a. Click Add Port Mapping. The Add Port Mapping dialog appears.
- b. Enter an IP address, host name, and container port.
- c. Select a Protocol from the drop-down list box.

For more information about configuring Port Mappings, see *Configuring Port Mappings for Containers* in [Oracle Linux: Podman User's Guide](#).

Volumes (optional)

Use the Volume properties to share file system space on host system with pod group containers.

To configure the storage volume properties, do the following:

- a. Click **Add Volumes**. The **Add Volumes** dialog appears.
- b. Enter a host path and container path.
- c. (Optional) Select the **Writable** checkbox to create a writable volume.
- d. In the **SELinux** drop down list, select one of the following options: No Label, Shared or Private.

For more information about volumes, see *Using Volumes with Containers* in [Use storage with Podman containers](#).

For more information about configuring SELinux with Podman, see Setting SELinux permission in Containers in [Oracle Linux: Podman User's Guide](#).

An empty pod group is created and appears in the **Container** table.

Add a Container to a Pod Group

Using the **Podman** web console page, Cockpit users can add a container to a pod group by creating a container within that group.

What Do You Need?

- The Cockpit web console must be installed and accessible. For details, see these topics: [Install and Enable Cockpit](#) and [Log in to the Cockpit Web Console](#).
- The Podman container tools must be installed on the host system. For information on how to install Podman, see *Installing Podman and Related Utilities* in [Oracle Linux: Podman User's Guide](#)
- The Cockpit-Podman add-on application must be installed. For further information on how to configure the Cockpit web console to interact with Podman, see [Install and Configure Cockpit-Podman](#).
- One or more pod groups must already exist in the Container table. For information on how to create a Pod Group, see [Create a Pod Group](#).
- Users with limited access privileges can only add containers to a pod group they created.

Steps

Using the Cockpit web console, follow these steps to add a container to a pod group.

1. In the Cockpit navigation pane, click **Podman**. The **Podman** page appears.
2. In the **Podman** page, navigate to the **Container** table and find the row with the pod group that you want to add a container and click Create container in pod. The **Create container in [pod group name]** dialog appears.

3. In the **Create container [pod group name]** dialog, provide the required properties to create a container. For more details, see Step 3 in [Create and Run Container](#) for instructions.
After creating a container in the pod group, the newly added container appears in pod group row of the **Container** table.

Inspect and Change a Pod Group

Using the **Podman** page in the Cockpit web console, you can inspect, and change existing pod group configurations as needed. For example, you can view pod groups to inspect container configurations. You can also commit changes, rename, or change the operating state of any container within a pod group. Finally, you can checkpoint, and restore a running pod group container, or remove a pod group as needed.

What Do You Need?

- The Cockpit web console must be installed and accessible.
For details, see these topics: [Install and Enable Cockpit](#) and [Log in to the Cockpit Web Console](#).
- The Podman container tools must be installed on the host system. For information on how to install Podman, see *Installing Podman and Related Utilities* in [Oracle Linux: Podman User's Guide](#)
- The Cockpit-Podman add-on application must be installed. For further information on how to configure the Cockpit web console to interact with Podman, see [Install and Configure Cockpit-Podman](#).
- One or more pod groups must already exist in the Container table. For information on how to create a Pod Group, see [Create a Pod Group](#).
- Users with limited access privileges can only access and change the containers and pod groups they created.

Steps

Using the Cockpit web console, follow these steps to inspect and change a container in a pod group.

1. In the Cockpit navigation pane, click **Podman**.
The **Podman** page appears.
2. In the **Podman** page, navigate to the **Container** table and find the row with the pod group of interest and perform any of the following:
 - **Inspect a container and access logs or CLI within pod group:** For instructions, see [Inspect Container and Access Container Logs and CLI](#).
 - **Rename or change a container operating state within a pod group:** For instructions, see [Rename, Pause, Stop, or Restart Container](#).
 - **Commit container changes to a new image:** For instructions, see [Commit Container Changes to Create New Image](#).
 - **Checkpoint and restore a system container within a pod group:** For instructions, see [Checkpoint and Restore a System Container](#).
 - **Remove a pod group:** For instructions, see [Remove Container or Pod Group](#).

7

Debugging Tools

When system issues occur, Cockpit administrators can use the following debugging tools in the web console to help evaluate and resolve system issues as they occur. For more details, see the following topics:

- [Evaluate System Problems Using Diagnostic Reports](#)
- [Capture Crash Dump Details Using Kdump](#)

Evaluate System Problems Using Diagnostic Reports

The Cockpit web console includes a Diagnostic Reports application that gathers various debugging information on the fly that helps to evaluate and resolve system related problems. For more details about using the Cockpit diagnostics reports available in the web console, see these topics:

- [Generate Diagnostic Reports](#)
- [Download or Remove Generated Reports](#)

Generate Diagnostic Reports

Using the **Diagnostics report** page in the web console, Cockpit administrators can collect configuration and diagnostic information from a host system to help with diagnosing system related problems.

What Do You Need?

- The Cockpit web console must be installed and accessible.
For details, see these topics: [Install and Enable Cockpit](#) and [Log in to the Cockpit Web Console](#).
- The `cockpit-sosreport` package is installed. Typically, this package is installed by default.
For package installation details, see [Install and Manage Add-on Applications](#).
- Administrator privileges.

Steps

Using the Cockpit web console, follow these steps to generate diagnostic reports about the hardware and setup of the managed host system.

1. In the Cockpit navigation pane, click **Diagnostic reports**.
The **Systems diagnostics** page appears.
2. In the **System diagnostics** page, click **Run reports**.
The **Run new report** dialog appears.
3. In the **Run new report** dialog, enter the following information and click **Run report**.

Report label	In the Report label text box, enter a name for this report. For example, you might provide your first initial and last name; or, if this report is being generated for a specific support case, you could enter the support case ID number.
Encryption passphrase (Optional)	If you need to secure the contents of the report, you can enter an encryption passphrase in the Encryption passphrase text box. This level of security is helpful to use when transferring a report over a public network to a third-party.
<div style="border-left: 2px solid #f0e68c; padding-left: 10px;"> <p>⚠ WARNING:</p> <p>Ensure that enough disk space is available when creating an encrypted report, as this process temporarily uses double the disk space.</p> </div>	
Options	Optionally, select any of the following options: <ul style="list-style-type: none"> • Obfuscate network addresses, hostnames, and usernames – Select this checkbox to intentionally render this information unclear and harder to understand. • Use verbose logging – Select this checkbox to gather more logging information than the normal logging mode.

Upon clicking **Run report** on the **Run new report** dialog, a progress label appears. This progress label indicates the completion percentage for rendering a completed report.

ⓘ Important:

If the system has a lot of packages installed, the report collection process might take a longer time to complete. To stop the report from generating, click **Stop report**.

Upon rendering a completed report, the **Run new report** dialog automatically closes, and the name of the newly generated report appears on the **System diagnostics** page.

 **Note:**

By default, the Cockpit generated report data is saved to a tar file format in the Cockpit user `home/downloads` directory.

Download or Remove Generated Reports


Using the **Diagnostic report** page in the web console, Cockpit administrators can download and extract generated report data or remove a generated report as needed.

What Do You Need?

- The Cockpit web console must be installed and accessible.
For details, see these topics: [Install and Enable Cockpit](#) and [Log in to the Cockpit Web Console](#).
- The `cockpit-sosreport` package is installed. Typically, this package is installed by default.
For package installation details, see [Install and Manage Add-on Applications](#).
- Administrator privileges.

Steps

Using the Cockpit web console, follow these steps to download or remove a Cockpit generated diagnostic report.

1. In the Cockpit navigation pane, click **Diagnostic reports**.
The **Systems diagnostics** page appears.
2. In the **Reports** panel on the **System diagnostics** page, perform any of the following:
 - **Download the report tar file and optionally extract content:**
 - a. Click **Download**. A dialog appears indicating the report is being downloaded.
 - b. Click the download dialog (or the browser window download arrow) to view the report's downloaded location.
By default, the generated report tar file is downloaded to the Cockpit user `/home/downloads` directory.
 - c. Optionally, extract the generated tar file to the local directory.
Upon extracting the tar file, the extracted report data appears across multiple subdirectories, for example: `sos_commands`, `sos_logs`, and `sos_reports`
 - **Remove a generated report:**
 - a. Click the actions  menu, which appears next to the report name you want to delete, and then select **Delete**.
A prompt appears confirming to you want to permanently delete the report.
 - b. Click **Delete** to confirm the deletion process.

Capture Crash Dump Details Using Kdump

Cockpit administrators can use Kdump in the web console to safeguard their system in the event of a kernel crash. For more information about using the Kdump features that are available in the web console, see the following topics:

- [Overview of Crash Recovery in Kdump](#)
- [Change Kdump Service State](#)
- [Change Crash Dump Target Location](#)
- [Test the Kdump Configuration](#)

Overview of Crash Recovery in Kdump

Kdump is a vital feature in Oracle Linux that helps to preserve the system in the event of a system failure. Kdump, when enabled, permanently reserves part of the system memory to capture and save kernel crash dump information for later analysis. Additionally, in the case of a kernel crash, kdump safeguards the system by automatically booting in to a clean and reliable second Linux kernel, called the captured kernel, which also resides as part of the reserved system memory.



Note:

The Kdump service is initially enabled at the time of installation. In the event that the Kdump package or the service wasn't configured at installation, administrators can install the package and enable the service from the command line. For more information about installing and enabling the Kdump service, see *Installing Kdump* in [Oracle Linux 9: Monitoring and Tuning the System](#)

When Kdump is properly configured on the host system, Cockpit administrators can use the Kdump application in the web console to view and manage Kdump configuration properties and test the Kdump configuration as needed.



Important:

The Kdump reserved memory usage property is configurable only from the command line. For details about the setting the reserved memory usage property from the command line, see *Working with Kernel Dumps* in [Oracle Linux 9: Monitoring and Tuning the System](#).

Change Kdump Service State

As needed, Cockpit administrators can change the Kdump service at boot using the web console.

What Do You Need?

- The Cockpit web console must be installed and accessible. For details, see these topics: [Install and Enable Cockpit](#) and [Log in to the Cockpit Web Console](#).
- The Kdump package (`kexec-tools`) is installed on the host.

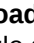
 **Note:**

An option for installing the Kdump package appears during the installation set up of Oracle Linux. If this package wasn't installed, see [Installing Kdump in Oracle Linux 9: Monitoring and Tuning the System](#)

- The `cockpit-kdump` package is installed on the host. Typically, this package is installed by default. For Cockpit add-on application installation details, see [Install and Manage Add-on Applications](#).
- Administrator privileges.

Steps

Using the Cockpit web console, follow these steps to change the kernel dump service state on the host system.

1. In the Cockpit navigation pane, click **Kernel dump**. The **Kernel crash dump** page appears.
2. In the **Kernel crash dump** page, find the **Status** for the service state and click **more details**. The **Service** page for the `kdump.service` appears.
3. In the **Service** page for the `kdump.service`, perform any of the following:
 - **Toggle Kdump service state at boot:** Turn off the toggle switch to deactivate the service state at boot or turn on the toggle switch to activate the service state at boot.
 - **Reload, restart, or stop Kdump service state:** Click the actions  menu next to the toggle switch and select the appropriate service state option, for instance:

Reload	Select Reload to reload the service configuration files but keep the process running.
Restart	Select Restart to shut down the service and then start the service.
Stop	Select Stop to shut down the service.

Change Crash Dump Target Location

Cockpit administrators, as needed, can change the Kdump crash target location by using the **Kernel dump** page in the web console.

What Do You Need?

- The Cockpit web console must be installed and accessible. For details, see these topics: [Install and Enable Cockpit](#) and [Log in to the Cockpit Web Console](#).
- The Kdump package (`kexec-tools`) is installed on the host.

 **Note:**

An option for installing the Kdump package appears during the installation set up of Oracle Linux. If this package wasn't installed, see *Installing Kdump* in [Oracle Linux 9: Monitoring and Tuning the System](#)

- The `cockpit-kdump` package is installed. Typically, this package is installed by default. For Cockpit package installation details, see [Install and Manage Add-on Applications](#).
- Administrator privileges.

Steps

Using the Cockpit web console, follow these steps to change the host target location for the saving the kernel crash dump information.


1. In the Cockpit navigation pane, click **Kernel dump**. The **Kernel crash dump** page appears.
2. In the **Kernel crash dump** page, click the link next to the **Crash dump** location property.

 **Important:**

By default, the Kdump crash files are stored locally to the `var/crash` host system directory.

The **Crash dump location** dialog appears.

3. In the **Crash dump location** dialog, change any of the following properties and then click **Save** to apply the changes.

Location	<p>Click the Location drop-down menu and select one of the following options:</p> <ul style="list-style-type: none"> • Local file system (default): Select this option to save the Kdump crash files to a local directory on the host system. • Remote over SSH: Select this option to transfer the Kdump crash files over a secure shell connection. The following SSH configuration properties are required: <div data-bbox="836 451 1461 714" style="border: 1px solid #ccc; padding: 10px; background-color: #e6f2ff;"> <p> Note:</p> <p>The secure shell option requires the configuration of an SSH client and server. For OpenSSH server configuration details, see Oracle Linux: Connecting to Remote Systems With OpenSSH.</p> </div> <ul style="list-style-type: none"> – Server: To establish a connection to a SSH server, specify the username and server name, for example: <i>username@servername.com</i>. – SSH Key: To identify yourself to the SSH server, enter the path to your public RSA key, for example: <i>root/ssh/kdump_id_rsa</i>. <ul style="list-style-type: none"> • Remote over NFS: Select this option to export the Kdump crash files to a configured shared network location. The following NFS configuration properties are required: <ul style="list-style-type: none"> – Server: Enter the name of the NFS host server, for example: <i>nfserver.example.com</i> – Export: Enter the NFS shared export directory path, for example: <i>/export/cores</i> <p>For more information about Oracle Linux NFS configuration, see Create an NFS server on Oracle Linux.</p>
Directory	<p>In the Directory text box, enter the target directory for where the Kdump crash files are stored. For example, <i>/var/crash</i>.</p>
Compression (Optional)	<p>Select the Compress crash dump to save space checkbox to enable this option.</p>

Test the Kdump Configuration

Cockpit administrators, as needed, can test the Kdump configuration by crashing the kernel on the host system using the **Kernel dump** page in the web console.

What Do You Need?

- The Cockpit web console must be installed and accessible. For details, see these topics: [Install and Enable Cockpit](#) and [Log in to the Cockpit Web Console](#).
- The Kdump package (`kexec-tools`) is installed on the host.

 **Note:**

An option for installing the Kdump package appears during the installation set up of Oracle Linux. If this package was not installed, see *Installing Kdump* in [Oracle Linux 9: Monitoring and Tuning the System](#)

- The `cockpit-kdump` package is installed. Typically, this package is installed by default.
For package installation details, see [Install and Manage Add-on Applications](#).
- Administrator privileges.

Steps

Using the Cockpit web console, follow these steps for testing the Kdump configuration.

1. In the Cockpit navigation pane, click **Kernel dump**.
The **Kernel dump** page appears.
2. In the **Kernel dump** page, click **Test configuration**.
The **Test kdump settings** dialog appears.

 **WARNING:**

Testing the Kdump configuration will result in disrupting the execution of the kernel. Proceeding with this operation (Crash system) results in a system crash and a possible loss of data.

3. In the **Test kdump settings** dialog, click **Crash system** to proceed with testing the Kdump configuration, otherwise, click **Cancel** to cancel the operation.