Red Hat Enterprise Linux 8

Managing systems using the RHEL 8 web console

Server management with a graphical web-based interface

Red Hat Enterprise Linux 8 Managing systems using the RHEL 8 web console

Server management with a graphical web-based interface

Legal Notice

Abstract

The RHEL web console is a web-based graphical interface, which is based on the upstream Cockpit project. By using it, you can perform system administration tasks, such as inspecting and controlling systemd services, managing storage, configuring networks, analyzing network issues, and inspecting logs.

Table of Contents

MAKING OPEN SOURCE MORE INCLUSIVE

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

PROVIDING FEEDBACK ON RED HAT DOCUMENTATION

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

Submitting feedback through Jira (account required)

1. Log in to the Jira website.

2. Click **Create** in the top navigation bar.

3. Enter a descriptive title in the **Summary** field.

4. Enter your suggestion for improvement in the **Description** field. Include links to the relevant parts of the documentation.

5. Click **Create** at the bottom of the dialogue.

CHAPTER 1. GETTING STARTED USING THE RHEL WEB CONSOLE

Learn how to install the Red Hat Enterprise Linux 8 web console, how to add and manage remote hosts through its convenient graphical interface, and how to monitor the systems managed by the web console.

1.1. WHAT IS THE RHEL WEB CONSOLE

The RHEL web console is a Red Hat Enterprise Linux web-based interface designed for managing and monitoring your local system, as well as Linux servers located in your network environment.

The RHEL web console enables you to perform a wide range of administration tasks, including:

- Managing services

- Managing user accounts

- Managing and monitoring system services

- Configuring network interfaces and firewall

- Reviewing system logs

- Managing virtual machines

- Creating diagnostic reports

- Setting kernel dump configuration

- Configuring SELinux

- Updating software

- Managing system subscriptions

The RHEL web console uses the same system APIs as you would use in a terminal, and actions performed in a terminal are immediately reflected in the RHEL web console.

You can monitor the logs of systems in the network environment, as well as their performance, displayed as graphs. In addition, you can change the settings directly in the web console or through the terminal.

1.2. INSTALLING AND ENABLING THE WEB CONSOLE

To access the RHEL 8 web console, first enable the **cockpit.socket** service.

Red Hat Enterprise Linux 8 includes the RHEL 8 web console installed by default in many installation variants. If this is not the case on your system, install the **cockpit** package before enabling the **cockpit.socket** service.

Procedure

1. If the web console is not installed by default on your installation variant, manually install the **cockpit** package:

```
# yum install cockpit
```

2. Enable and start the **cockpit.socket** service, which runs a web server:

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

3. If the web console was not installed by default on your installation variant and you are using a custom firewall profile, add the **cockpit** service to **firewalld** to open port 9090 in the firewall:

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

Verification steps

- To verify the previous installation and configuration, open the web console.

1.3. LOGGING IN TO THE WEB CONSOLE

When the **cockpit.socket** service is running and the corresponding firewall port is open, you can log in to the web console in your browser for the first time.

Prerequisites

- Use one of the following browsers to open the web console:

    - Mozilla Firefox 52 and later

    - Google Chrome 57 and later

    - Microsoft Edge 16 and later

- System user account credentials
  The RHEL web console uses a specific pluggable authentication modules (PAM) stack at **/etc/pam.d/cockpit**. The default configuration allows logging in with the user name and password of any local account on the system.

- Port 9090 is open in your firewall.

Procedure

1. In your web browser, enter the following address to access the web console:

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

NOTE

This provides a web-console login on your local machine. If you want to log in to the web console of a remote system, see Section 1.6, "Connecting to the web console from a remote machine"

If you use a self-signed certificate, the browser displays a warning. Check the certificate, and accept the security exception to proceed with the login.

The console loads a certificate from the **/etc/cockpit/ws-certs.d** directory and uses the last file with a **.cert** extension in alphabetical order. To avoid having to grant security exceptions, install a certificate signed by a certificate authority (CA).

2. In the login screen, enter your system user name and password.

3. Click **Log In**.

After successful authentication, the RHEL web console interface opens.

> **NOTE**
>
> To switch between limited and administrative access, click **Administrative access** or **Limited access** in the top panel of the web console page. You must provide your user password to gain administrative access.

1.4. CHANGING THE DEFAULT STYLE SETTING FOR THE WEB CONSOLE

By default, the web console adopts its style setting from the setting of your browser. You can override the default style setting from your RHEL 8 web console interface.

Prerequisites

- The web console is installed and accessible. For details, see Installing the web console .

Procedure

1. Log in to the RHEL web console. For details, see Logging in to the web console .

2. In the upper right corner, click the **Session** button.

3. In the section **Style**, choose the preferred setting. The **Default** setting uses the same style setting as your browser.

Verification steps

1. The style setting has changed according to set style.

1.5. DISABLING BASIC AUTHENTICATION IN THE WEB CONSOLE

You can modify the behavior of an authentication scheme by modifying the **cockpit.conf** file. Use the **none** action to disable an authentication scheme and only allow authentication through GSSAPI and forms.

Prerequisites

- The web console is installed and accessible. For details, see Installing the web console .

- You have **root** privileges or permissions to enter administrative commands with **sudo**.

Procedure

1. Open or create the **cockpit.conf** file in the **/etc/cockpit/** directory in a text editor of your preference, for example:

   ```
   # vi cockpit.conf
   ```

2. Add the following text:

   ```
   [basic]
   action = none
   ```

3. Save the file.

4. Restart the web console for changes to take effect.

   ```
   # systemctl try-restart cockpit
   ```

1.6. CONNECTING TO THE WEB CONSOLE FROM A REMOTE MACHINE

You can connect to your web console interface from any client operating system and also from mobile phones or tablets.

Prerequisites

- A device with a supported internet browser, such as:

  - Mozilla Firefox 52 and later

  - Google Chrome 57 and later

  - Microsoft Edge 16 and later

- RHEL 8 server you want to access with an installed and accessible web console. For more information about the installation of the web console see Installing the web console .

Procedure

1. Open your web browser.

2. Type the remote server's address in one of the following formats:

   a. With the server's host name:

      ```
      https://<server.hostname.example.com>:<port-number>
      ```

      For example:

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

   b. With the server's IP address:

      ```
      https://<server.IP_address>:<port-number>
      ```

      For example:

> https://192.0.2.2:9090

3. After the login interface opens, log in with your RHEL system credentials.

1.7. LOGGING IN TO THE WEB CONSOLE USING A ONE-TIME PASSWORD

If your system is part of an Identity Management (IdM) domain with enabled one-time password (OTP) configuration, you can use an OTP to log in to the RHEL web console.

> **IMPORTANT**
>
> It is possible to log in using a one-time password only if your system is part of an Identity Management (IdM) domain with enabled OTP configuration. For more information about OTP in IdM, see One-time password in Identity Management .

Prerequisites

- The RHEL web console has been installed.
  For details, see Installing the web console .

- An Identity Management server with enabled OTP configuration.
  For details, see One-time password in Identity Management .

- A configured hardware or software device generating OTP tokens.

Procedure

1. Open the RHEL web console in your browser:

   - Locally: **https://localhost:PORT_NUMBER**

   - Remotely with the server hostname: **https://example.com:PORT_NUMBER**

   - Remotely with the server IP address:
     https://EXAMPLE.SERVER.IP.ADDR:PORT_NUMBER
     If you use a self-signed certificate, the browser issues a warning. Check the certificate and accept the security exception to proceed with the login.

     The console loads a certificate from the **/etc/cockpit/ws-certs.d** directory and uses the last file with a **.cert** extension in alphabetical order. To avoid having to grant security exceptions, install a certificate signed by a certificate authority (CA).

2. The Login window opens. In the Login window, enter your system user name and password.

3. Generate a one-time password on your device.

4. Enter the one-time password into a new field that appears in the web console interface after you confirm your password.

5. Click **Log in**.

6. Successful login takes you to the **Overview** page of the web console interface.

1.8. REBOOTING THE SYSTEM USING THE WEB CONSOLE

You can use the web console to restart a RHEL system that the web console is attached to.

Prerequisites

- The web console is installed and accessible. For details, see Installing the web console.

Procedure

1. Log into the RHEL web console. For details, see Logging in to the web console .

2. In the **Overview** page, click the **Reboot** button.
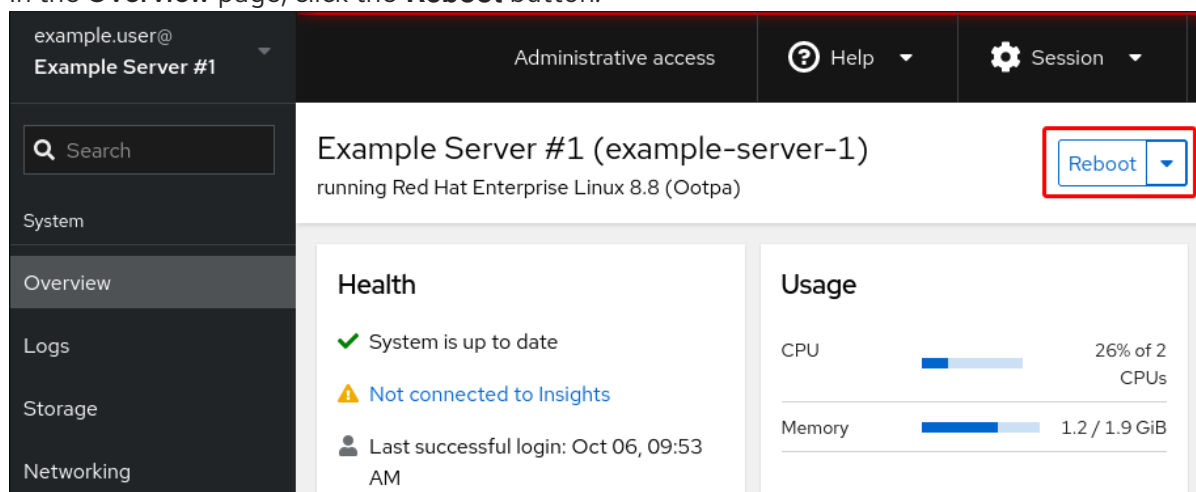
3. If any users are logged in to the system, you can write a message about the restart in the **Reboot** dialog box.

4. Optional: In the **Delay** drop down list, select a time interval for the reboot delay.
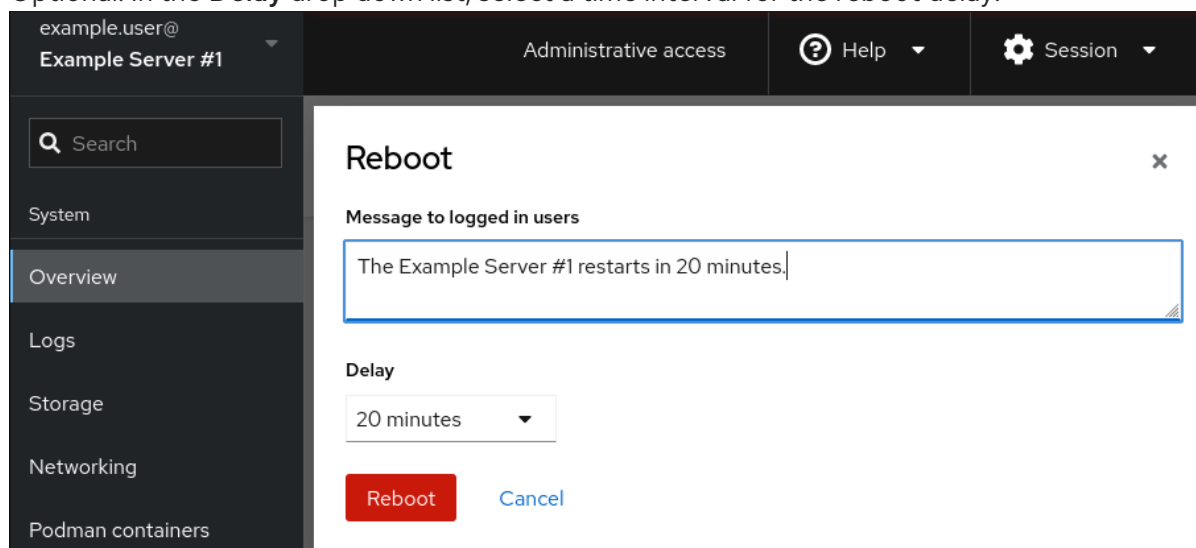
5. Click **Reboot**.

1.9. SHUTTING DOWN THE SYSTEM USING THE WEB CONSOLE

You can use the web console to shut down a RHEL system that the web console is attached to.

Prerequisites

- The web console is installed and accessible.
  For details, see Installing the web console.

Procedure

1. Log into the RHEL web console.
   For details, see Logging in to the web console .

2. Click **Overview**.

3. In the **Restart** drop down list, select **Shut Down**.
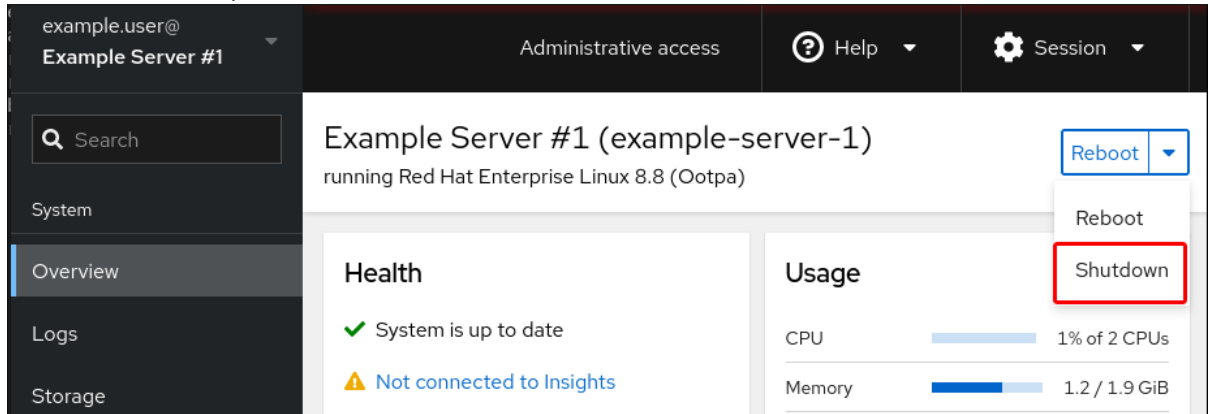
4. If any users are logged in to the system, write a reason for the shutdown in the **Shut Down** dialog box.

5. Optional: In the **Delay** drop down list, select a time interval.

6. Click **Shut Down**.

1.10. CONFIGURING TIME SETTINGS USING THE WEB CONSOLE

You can set a time zone and synchronize the system time with a Network Time Protocol (NTP) server.

Prerequisites

- The web console is installed and accessible.
  For details, see Installing the web console.

Procedure

1. Log in to the RHEL web console.
   For details, see Logging in to the web console .

2. Click the current system time in **Overview**.
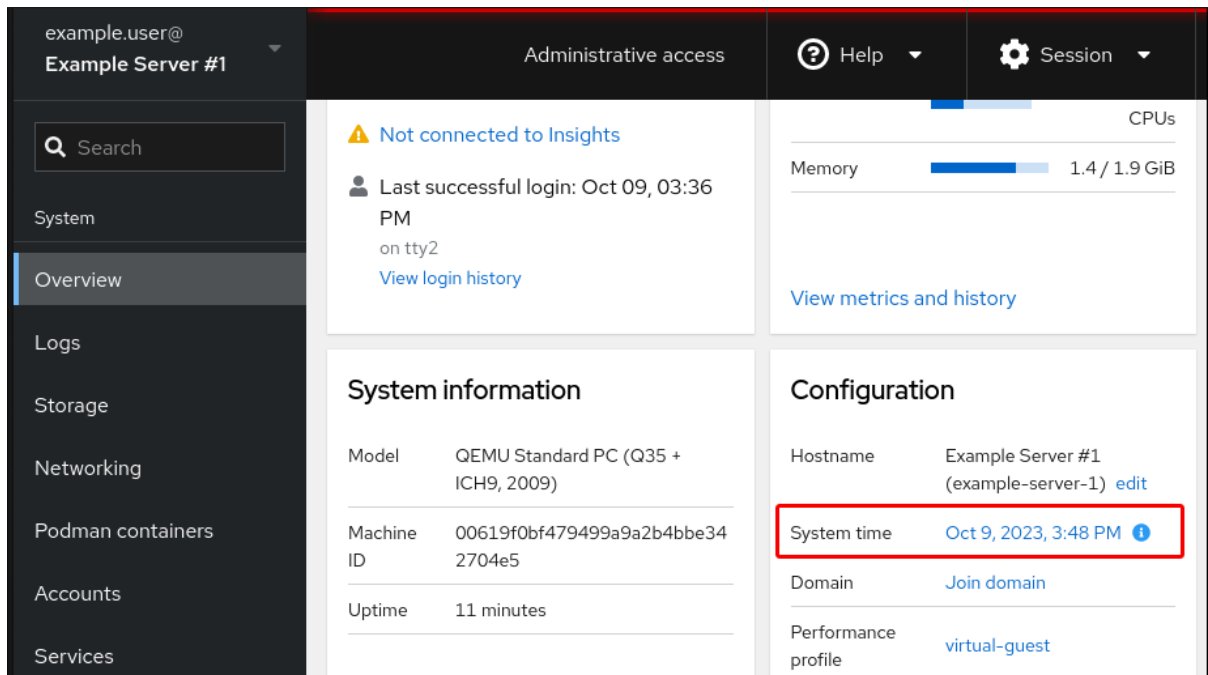
3. Click **System time**.

4. In the **Change System Time** dialog box, change the time zone if necessary.

5. In the **Set Time** drop down menu, select one of the following:

   Manually

   Use this option if you need to set the time manually, without an NTP server.

   Automatically using NTP server

   This is a default option, which synchronizes time automatically with the preset NTP servers.

   Automatically using specific NTP servers

   Use this option only if you need to synchronize the system with a specific NTP server. Specify the DNS name or the IP address of the server.

6. Click **Change**.
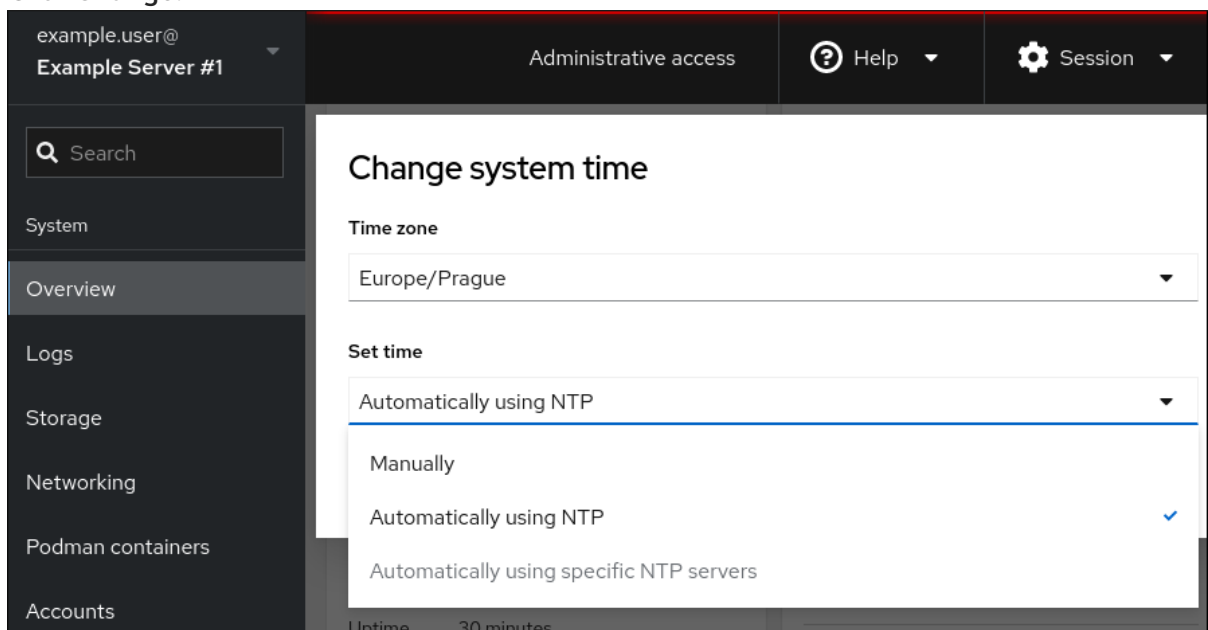
Verification steps

- Check the system time displayed in the **System** tab.

Additional resources

- Using the Chrony suite to configure NTP .

1.11. JOINING A RHEL 8 SYSTEM TO AN IDM DOMAIN USING THE WEB CONSOLE

You can use the web console to join the Red Hat Enterprise Linux 8 system to the Identity Management (IdM) domain.

Prerequisites

- The IdM domain is running and reachable from the client you want to join.

- You have the IdM domain administrator credentials.

Procedure

1. Log into the RHEL web console.
   For details, see Logging in to the web console .

2. In the **Configuration** field of the **Overview** tab click **Join Domain**.

3. In the **Join a Domain** dialog box, enter the host name of the IdM server in the **Domain Address** field.

4. In the **Domain administrator name** field, enter the user name of the IdM administration account.

5. In the **Domain administrator password**, add a password.

6. Click **Join**.

Verification steps

1. If the RHEL 8 web console did not display an error, the system has been joined to the IdM domain and you can see the domain name in the **System** screen.

2. To verify that the user is a member of the domain, click the Terminal page and type the **id** command:

   ```
   $ id
   euid=548800004(example_user) gid=548800004(example_user)
   groups=548800004(example_user) context=unconfined_u:unconfined_r:unconfined_t:s0-
   s0:c0.c1023
   ```

Additional resources

- Planning Identity Management

- Installing Identity Management

- Managing IdM users, groups, hosts, and access control rules

1.12. DISABLING SMT TO PREVENT CPU SECURITY ISSUES USING THE WEB CONSOLE

Disable Simultaneous Multi Threading (SMT) in case of attacks that misuse CPU SMT. Disabling SMT can mitigate security vulnerabilities, such as L1TF or MDS.

IMPORTANT

Disabling SMT might lower the system performance.

Prerequisites

- The web console must be installed and accessible. For details, see Installing the web console.

Procedure

1. Log in to the RHEL web console. For details, see Logging in to the web console .

2. In the **Overview** tab find the **System information** field and click **View hardware details**.

3. On the **CPU Security** line, click **Mitigations**.
   If this link is not present, it means that your system does not support SMT, and therefore is not vulnerable.

4. In the **CPU Security Toggles** table, turn on the **Disable simultaneous multithreading (nosmt)** option.

5. Click the **Save and reboot** button.

After the system restart, the CPU no longer uses SMT.

Additional resources

- L1TF - L1 Terminal Fault Attack - CVE-2018-3620 & CVE-2018-3646

- MDS - Microarchitectural Data Sampling - CVE-2018-12130, CVE-2018-12126, CVE-2018-12127, and CVE-2019-11091

1.13. ADDING A BANNER TO THE LOGIN PAGE

You can set the web console to show a content of a banner file on the login screen.

Prerequisites

- The web console is installed and accessible.
  For details, see Installing the web console.

- You have **root** privileges or permissions to enter administrative commands with **sudo**.

Procedure

1. Open the **/etc/issue.cockpit** file in a text editor of your preference:

```
# vi /etc/issue.cockpit
```

2. Add the content you want to display as the banner to the file, for example:

   This is an example banner for the RHEL web console login page.

   You cannot include any macros in the file, but you can use line breaks and ASCII art.

3. Save the file.

4. Open the **cockpit.conf** file in the **/etc/cockpit/** directory in a text editor of your preference, for example:

   ```
   # vi /etc/cockpit/cockpit.conf
   ```

5. Add the following text to the file:

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

6. Save the file.

7. Restart the web console for changes to take effect.

   ```
   # systemctl try-restart cockpit
   ```

Verification steps

- Open the web console login screen again to verify that the banner is now visible:

1.14. CONFIGURING AUTOMATIC IDLE LOCK IN THE WEB CONSOLE

You can enable the automatic idle lock and set the idle timeout for your system through the web console interface.

Prerequisites

- The web console must be installed and accessible.
  For details, see Installing the web console .

- You have **root** privileges or permissions to enter administrative commands with **sudo**.

Procedure

1. Open the **cockpit.conf** file in the **/etc/cockpit/** directory in a text editor of your preference, for example:

   # vi /etc/cockpit/cockpit.conf

2. Add the following text to the file:

   [Session]
   IdleTimeout=<*X*>

   Substitute <*X*> with a number for a time period of your choice in minutes.

3. Save the file.

4. Restart the web console for changes to take effect.

```
# systemctl try-restart cockpit
```

Verification steps

- Check if the session logs you out after a set period of time.

CHAPTER 2. CONFIGURING THE HOST NAME IN THE WEB CONSOLE

Learn how to use the Red Hat Enterprise Linux web console to configure different forms of the host name on the system that the web console is attached to.

2.1. HOST NAME

The host name identifies the system. By default, the host name is set to **localhost**, but you can change it.

A host name consists of two parts:

Host name

It is a unique name which identifies a system.

Domain

Add the domain as a suffix behind the host name when using a system in a network and when using names instead of just IP addresses.

A host name with an attached domain name is called a fully qualified domain name (FQDN). For example: **mymachine.example.com**.

Host names are stored in the **/etc/hostname** file.

2.2. PRETTY HOST NAME IN THE WEB CONSOLE

You can configure a pretty host name in the RHEL web console. The pretty host name is a host name with capital letters, spaces, and so on.

The pretty host name displays in the web console, but it does not have to correspond with the host name.

> **Example 2.1. Host name formats in the web console**
>
> **Pretty host name**
>   **My Machine**
> **Host name**
>   **mymachine**
> **Real host name – fully qualified domain name (FQDN)**
>   **mymachine.idm.company.com**

2.3. SETTING THE HOST NAME USING THE WEB CONSOLE

This procedure sets the real host name or the pretty host name in the web console.

Prerequisites

- The web console is installed and accessible.
  For details, see [Installing the web console](#).

Procedure

1. Log into the web console.
   For details, see Logging in to the web console .

2. Click **Overview**.

3. Click **edit** next to the current host name.
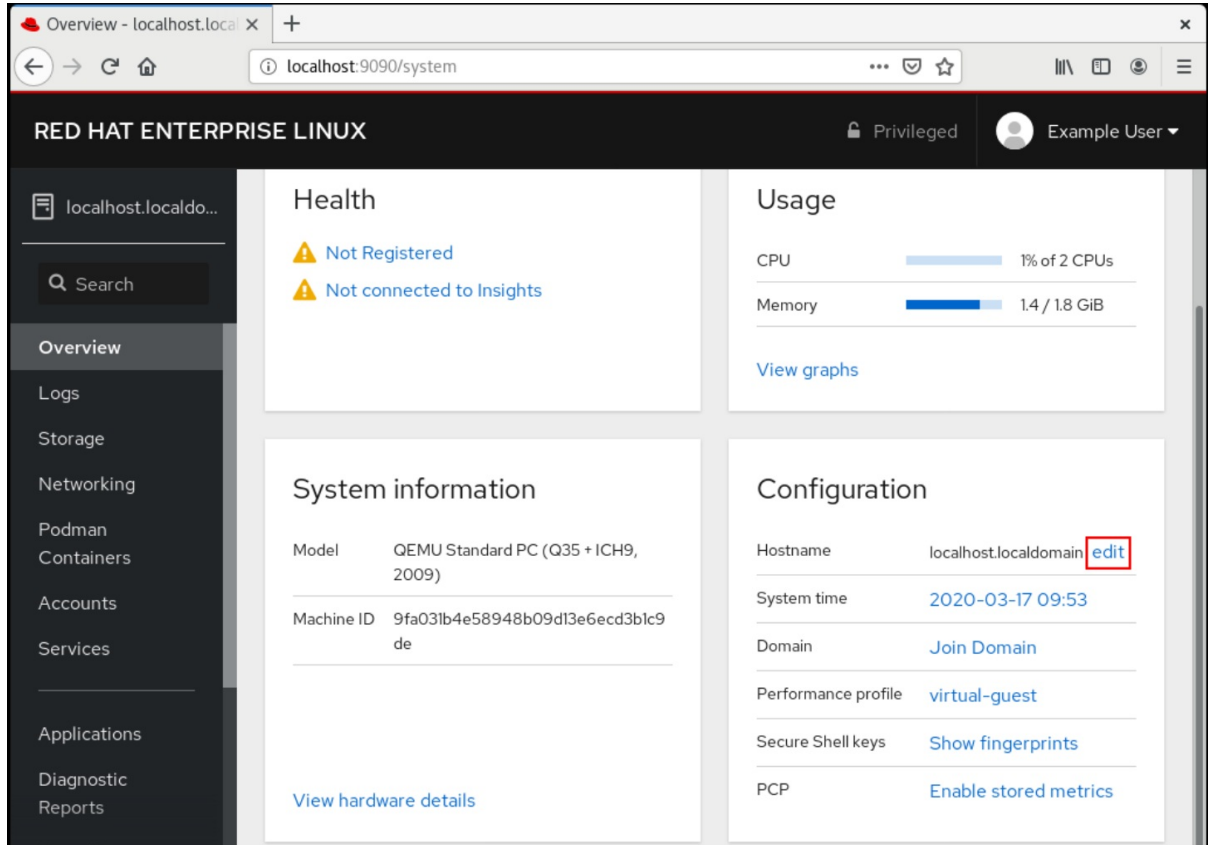
4. In the **Change Host Name** dialog box, enter the host name in the **Pretty Host Name** field.

5. The **Real Host Name** field attaches a domain name to the pretty name.
   You can change the real host name manually if it does not correspond with the pretty host name.

6. Click **Change**.
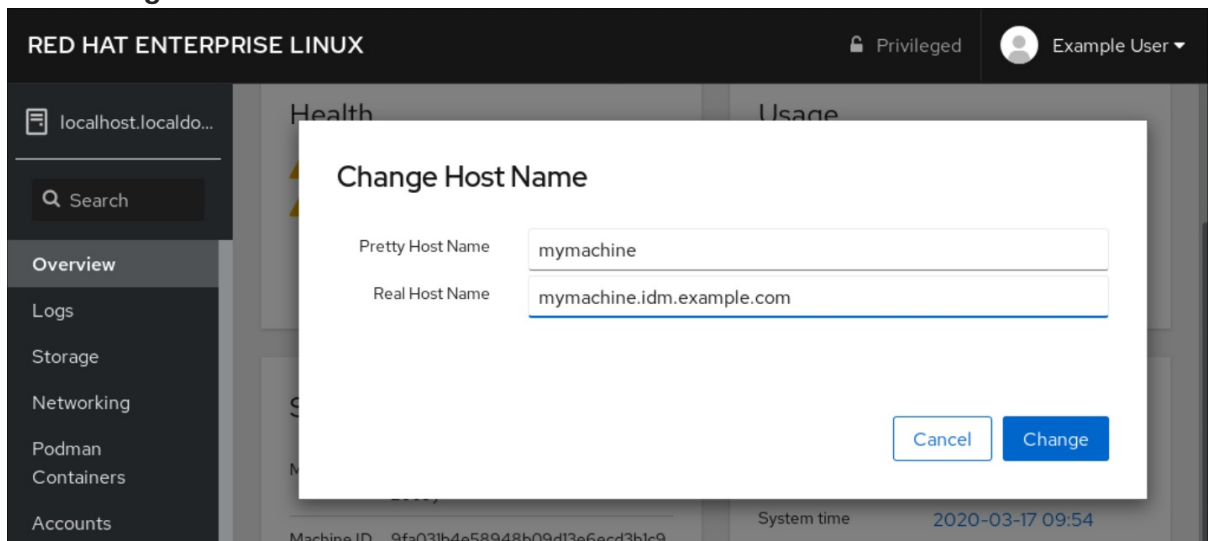
Verification steps

1. Log out from the web console.

2. Reopen the web console by entering an address with the new host name in the address bar of your browser.
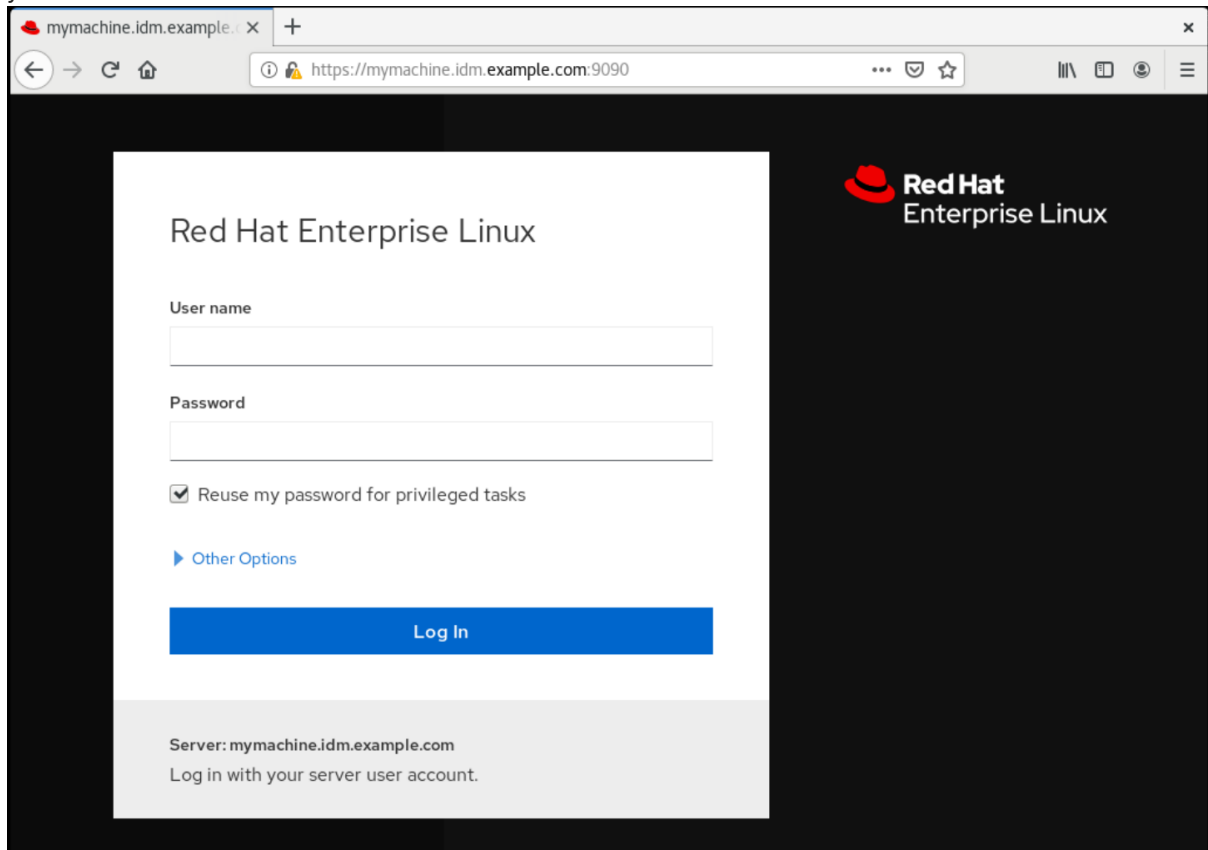
CHAPTER 3. INSTALLING WEB CONSOLE ADD-ONS AND CREATING CUSTOM PAGES

Depending on how you want to use your Red Hat Enterprise Linux system, you can add additional **available** applications to the web console or create custom pages based on your use case.

3.1. ADD-ONS FOR THE RHEL WEB CONSOLE

While the **cockpit** package is a part of Red Hat Enterprise Linux by default, you can install add-on applications on demand using the following command:

> # yum install *<add-on>*

In the previous command, replace *<add-on>* by a package name from the list of available add-on applications for the RHEL web console.

Feature name	Package name	Usage
Composer	**cockpit-composer**	Building custom OS images
Machines	**cockpit-machines**	Managing **libvirt** virtual machines
PackageKit	**cockpit-packagekit**	Software updates and application installation (usually installed by default)
PCP	**cockpit-pcp**	Persistent and more fine-grained performance data (installed on demand from the UI)
Podman	**cockpit-podman**	Managing containers and managing container images (available from RHEL 8.1)
Session Recording	**cockpit-session-recording**	Recording and managing user sessions
Storage	**cockpit-storaged**	Managing storage through **udisks**

3.2. CREATING NEW PAGES IN THE WEB CONSOLE

If you want to add customized functions to your Red Hat Enterprise Linux web console, you must add the package directory that contains the HTML and JavaScript files for the page that runs the required function.

For detailed information about adding custom pages, see Creating Plugins for the Cockpit User Interface on the Cockpit Project website.

Additional resources

- Cockpit Packages section in the Cockpit Project Developer Guide

CHAPTER 4. OPTIMIZING THE SYSTEM PERFORMANCE USING THE WEB CONSOLE

Learn how to set a performance profile in the RHEL web console to optimize the performance of the system for a selected task.

4.1. PERFORMANCE TUNING OPTIONS IN THE WEB CONSOLE

Red Hat Enterprise Linux 8 provides several performance profiles that optimize the system for the following tasks:

- Systems using the desktop

- Throughput performance

- Latency performance

- Network performance

- Low power consumption

- Virtual machines

The **TuneD** service optimizes system options to match the selected profile.

In the web console, you can set which performance profile your system uses.

Additional resources

- Getting started with TuneD

4.2. SETTING A PERFORMANCE PROFILE IN THE WEB CONSOLE

Depending on the task you want to perform, you can use the web console to optimize system performance by setting a suitable performance profile.

Prerequisites

- Make sure the web console is installed and accessible. For details, see Installing the web console.

Procedure

1. Log into the 8 web console. For details, see Logging in to the web console .

2. Click **Overview**.

3. In the **Configuration** section, click the current performance profile.
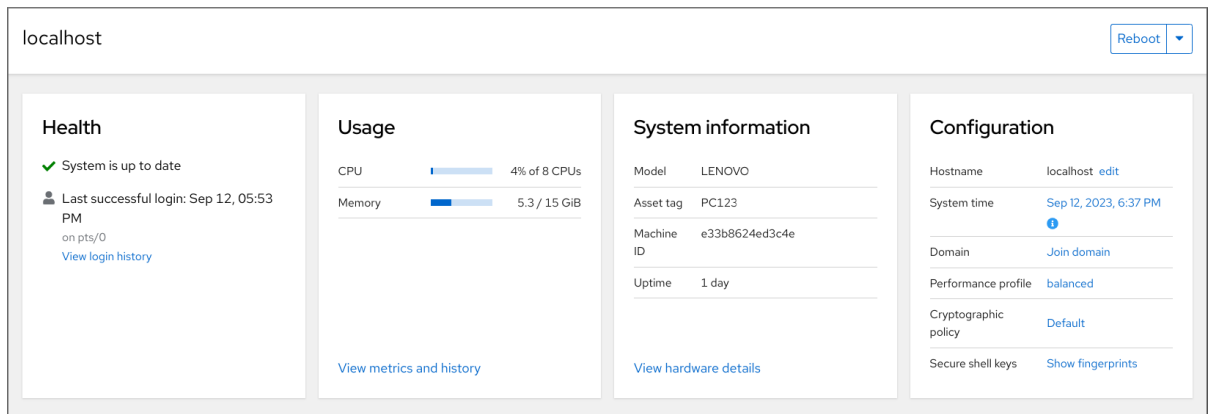
4. In the **Change Performance Profile** dialog box, set the required profile.

5. Click **Change Profile**.

Verification steps

- The **Overview** tab now shows the selected performance profile in the **Configuration** section.

4.3. MONITORING PERFORMANCE ON THE LOCAL SYSTEM USING THE WEB CONSOLE

Red Hat Enterprise Linux web console uses the Utilization Saturation and Errors (USE) Method for troubleshooting. The new performance metrics page has a historical view of your data organized chronologically with the newest data at the top.

In the **Metrics and history** page, you can view events, errors, and graphical representation for resource utilization and saturation.

Prerequisites

- The web console is installed and accessible. For details, see Installing the web console .

- The **cockpit-pcp** package, which enables collecting the performance metrics, is installed:

  a. To install the package from the web console interface:

     i. Log in to the web console with administrative privileges. For details, see Logging in to the web console.

     ii. In the **Overview** page, click **View metrics and history**.

     iii. Click the **Install cockpit-pcp** button.

     iv. In the **Install software** dialog window, click **Install**.

  b. To install the package from the command-line interface, use:

     ```
     # yum install cockpit-pcp
     ```

- The Performance Co-Pilot (PCP) service is enabled:

  ```
  # systemctl enable --now pmlogger.service pmproxy.service
  ```

Procedure

1. Log into the 8 web console. For details, see Logging in to the web console .

2. Click **Overview**.

3. In the **Usage** section, click **View metrics and history**.

localhost			Reboot ▾
Health	**Usage**	**System information**	**Configuration**
✔ System is up to date	CPU 4% of 8 CPUs	Model LENOVO	Hostname localhost edit
👤 Last successful login: Sep 12, 05:53 PM on pts/0	Memory 5.3 / 15 GiB	Asset tag PC123	System time Sep 12, 2023, 6:37 PM ⓘ
View login history		Machine ID e33b8624ed3c4e	Domain Join domain
		Uptime 1 day	Performance profile balanced
			Cryptographic policy Default
	View metrics and history	View hardware details	Secure shell keys Show fingerprints

   The **Metrics and history** section opens:

- The current system configuration and usage:
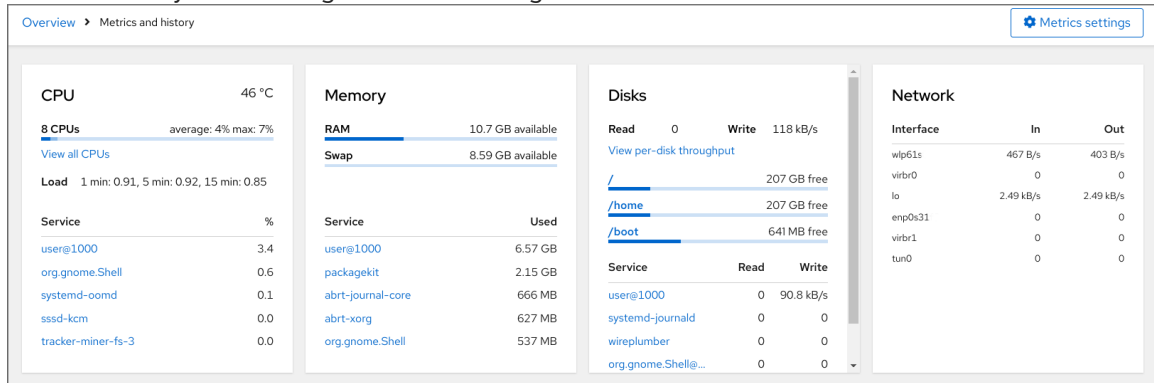
- The performance metrics in a graphical form over a user-specified time interval:
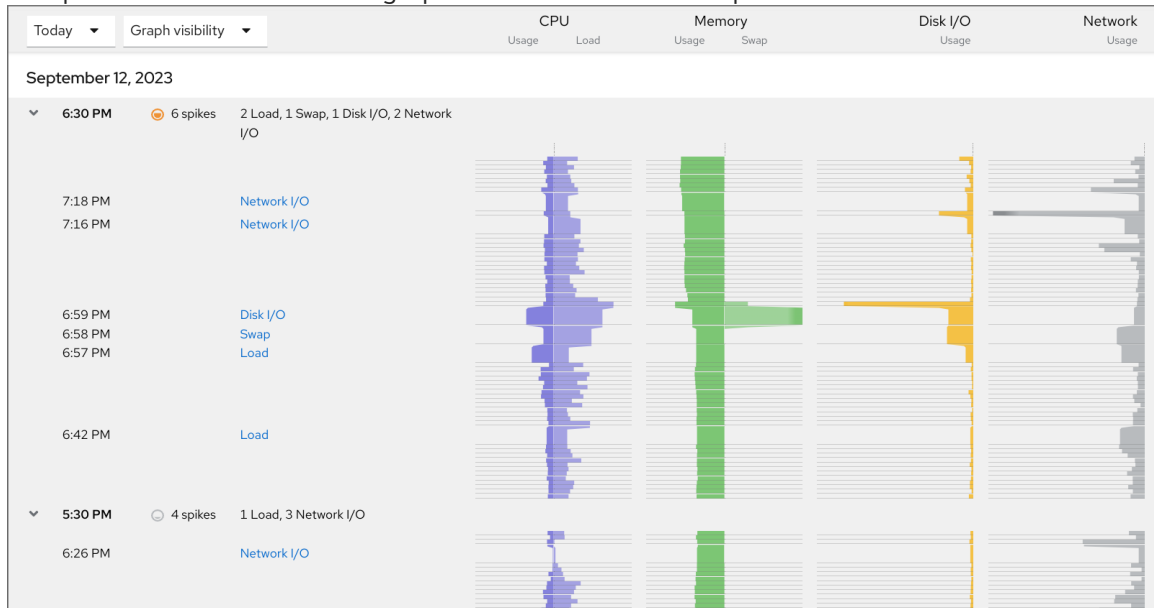
4.4. MONITORING PERFORMANCE ON SEVERAL SYSTEMS USING THE WEB CONSOLE AND GRAFANA

Grafana enables you to collect data from several systems at once and review a graphical representation of their collected Performance Co-Pilot (PCP) metrics. You can set up performance metrics monitoring and export for several systems in the web console interface.

Prerequisites

- The web console must be installed and accessible. For details, see link:Installing the web console.

- Install the **cockpit-pcp** package.

  1. From the web console interface:

     a. Log in to the web console with administrative privileges. For details, see Logging in to the web console.

     b. In the **Overview** page, click **View details and history**.

     c. Click the **Install cockpit-pcp** button.

     d. In the **Install software** dialog window, click **Install**.

e. Log out and in again to see the metrics history.

2. To install the package from the command-line interface, use:

```
# yum install cockpit-pcp
```

- Enable the PCP service:

```
# systemctl enable --now pmlogger.service pmproxy.service
```

- Set up Grafana dashboard. For more information, see Setting up a grafana-server.

- Install the **redis** package.

```
# yum install redis
```

Alternatively, you can install the package from the web console interface later in the procedure.

Procedure

1. In the **Overview** page, click **View metrics and history** in the **Usage** table.

2. Click the **Metrics settings** button.

3. Move the **Export to network** slider to active position.
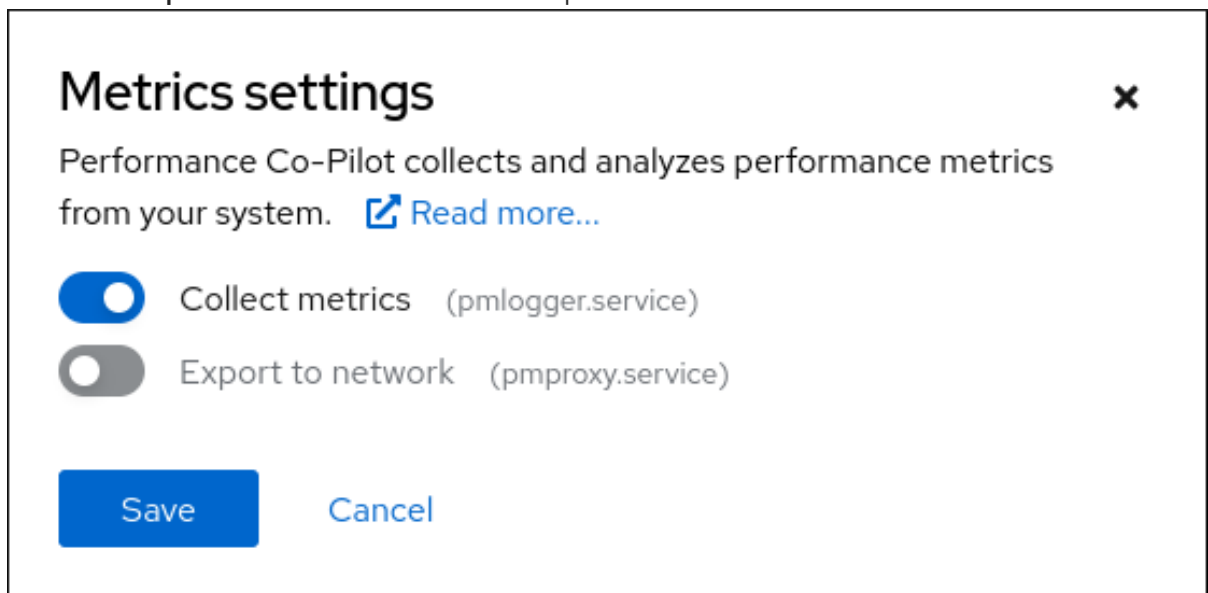
If you do not have the **redis** package installed, the web console prompts you to install it.

4. To open the **pmproxy** service, select a zone from a drop-down list and click the **Add pmproxy** button.

5. Click **Save**.

Verification

1. Click **Networking**.

2. In the **Firewall** table, click the **Edit rules and zones** button.

3. Search for **pmproxy** in your selected zone.

> **IMPORTANT**
>
> Repeat this procedure on all the systems you want to watch.

Additional resources

- Setting up graphical representation of PCP metrics

CHAPTER 5. REVIEWING LOGS IN THE WEB CONSOLE

Learn how to access, review and filter logs in the RHEL 8 web console.

5.1. REVIEWING LOGS IN THE WEB CONSOLE

The RHEL 8 web console Logs section is a UI for the **journalctl** utility. You can access system logs in the web console interface.

Prerequisites

- The RHEL 8 web console has been installed.
  For details, see Installing the web console.

Procedure

1. Log in to the RHEL web console.
   For details, see Logging in to the web console .

2. Click **Logs**.
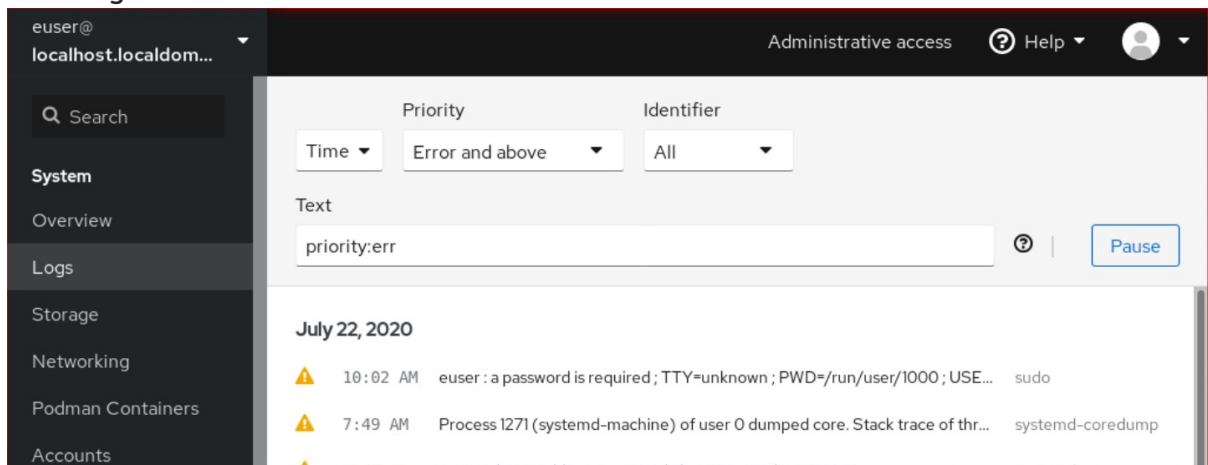
3. Open log entry details by clicking on your selected log entry in the list.

> **NOTE**
>
> You can use the **Pause** button to pause new log entries from appearing. Once you resume new log entries, the web console will load all log entries that were reported after you used the **Pause** button.

You can filter the logs by time, priority or identifier. For more information, see Filtering logs in the web console.

5.2. FILTERING LOGS IN THE WEB CONSOLE

You can filter log entries in the web console.

Prerequisites

- The web console interface must be installed and accessible.
  For details, see Installing the web console.

Procedure

1. Log in to the RHEL 8 web console.
   For details, see Logging in to the web console .

2. Click **Logs**.

3. By default, web console shows the latest log entries. To filter by a specific time range, click the
   Time drop-down menu and choose a preferred option.
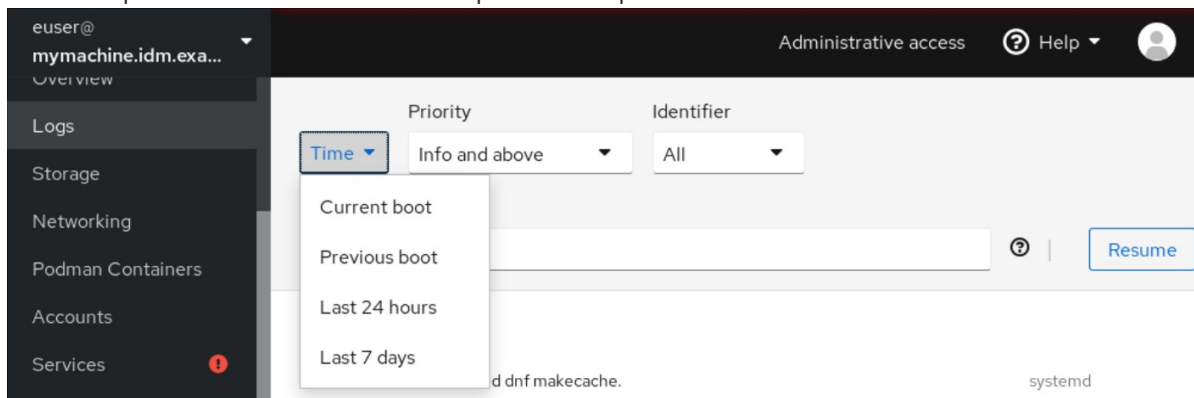
4. **Error and above** severity logs list is shown by default. To filter by different priority, click the
   Error and above drop-down menu and choose a preferred priority.
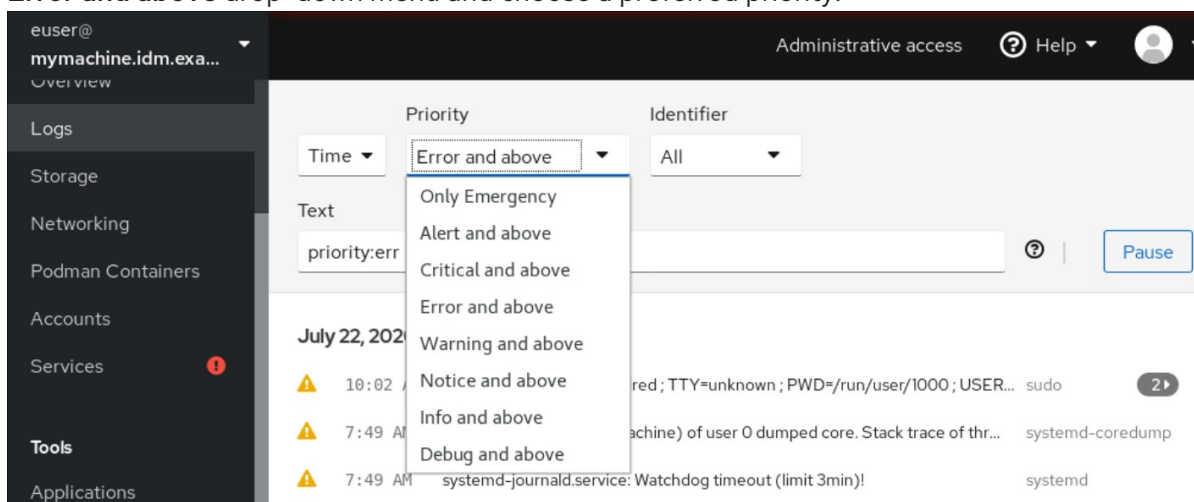
5. By default, web console shows logs for all identifiers. To filter logs for a particular identifier, click
   the **All** drop-down menu and select an identifier.
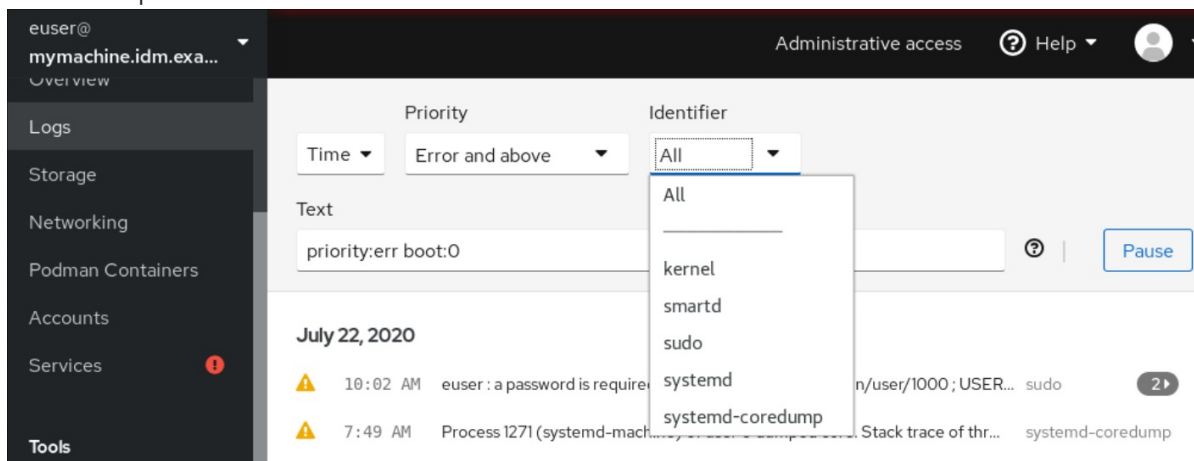
6. To open a log entry, click on a selected log.

5.3. TEXT SEARCH OPTIONS FOR FILTERING LOGS IN THE WEB CONSOLE

The text search option functionality provides a big variety of options for filtering logs. If you decide to filter logs by using the text search, you can use the predefined options that are defined in the three drop-down menus, or you can type the whole search yourself.

Drop-down menus

There are three drop-down menus that you can use to specify the main parameters of your search:

- **Time**: This drop-down menu contains predefined searches for different time ranges of your search.

- **Priority**: This drop-down menu provides options for different priority levels. It corresponds to the **journalctl --priority** option. The default priority value is **Error and above**. It is set every time you do not specify any other priority.

- **Identifier**: In this drop-down menu, you can select an identifier that you want to filter. Corresponds to the **journalctl --identifier** option.

Quantifiers

There are six quantifiers that you can use to specify your search. They are covered in the Options for filtering logs table.

Log fields

If you want to search for a specific log field, it is possible to specify the field together with its content.

Free-form text search in logs messages

You can filter any text string of your choice in the logs messages. The string can also be in the form of a regular expressions.

Advanced logs filtering I

Filter all log messages identified by 'systemd' that happened since October 22, 2020 midnight and journal field 'JOB_TYPE' is either 'start' or 'restart.

1. Type **identifier:systemd since:2020-10-22 JOB_TYPE=start,restart** to search field.

2. Check the results.
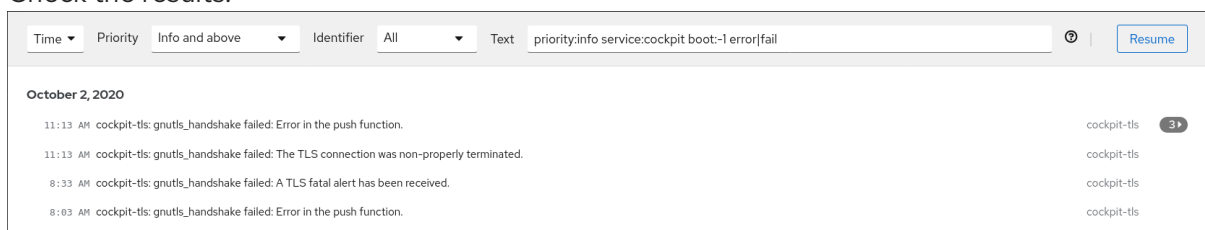
Advanced logs filtering II

Filter all log messages that come from 'cockpit.service' systemd unit that happened in the boot before last and the message body contains either "error" or "fail".

1. Type **service:cockpit boot:-1 error|fail** to the search field.

2. Check the results.

5.4. USING A TEXT SEARCH BOX TO FILTER LOGS IN THE WEB CONSOLE

Using the text search box allows you to filter logs according to different parameters. The search combines usage of the filtering drop-down menus, quantifiers, log fields and free-form string search.
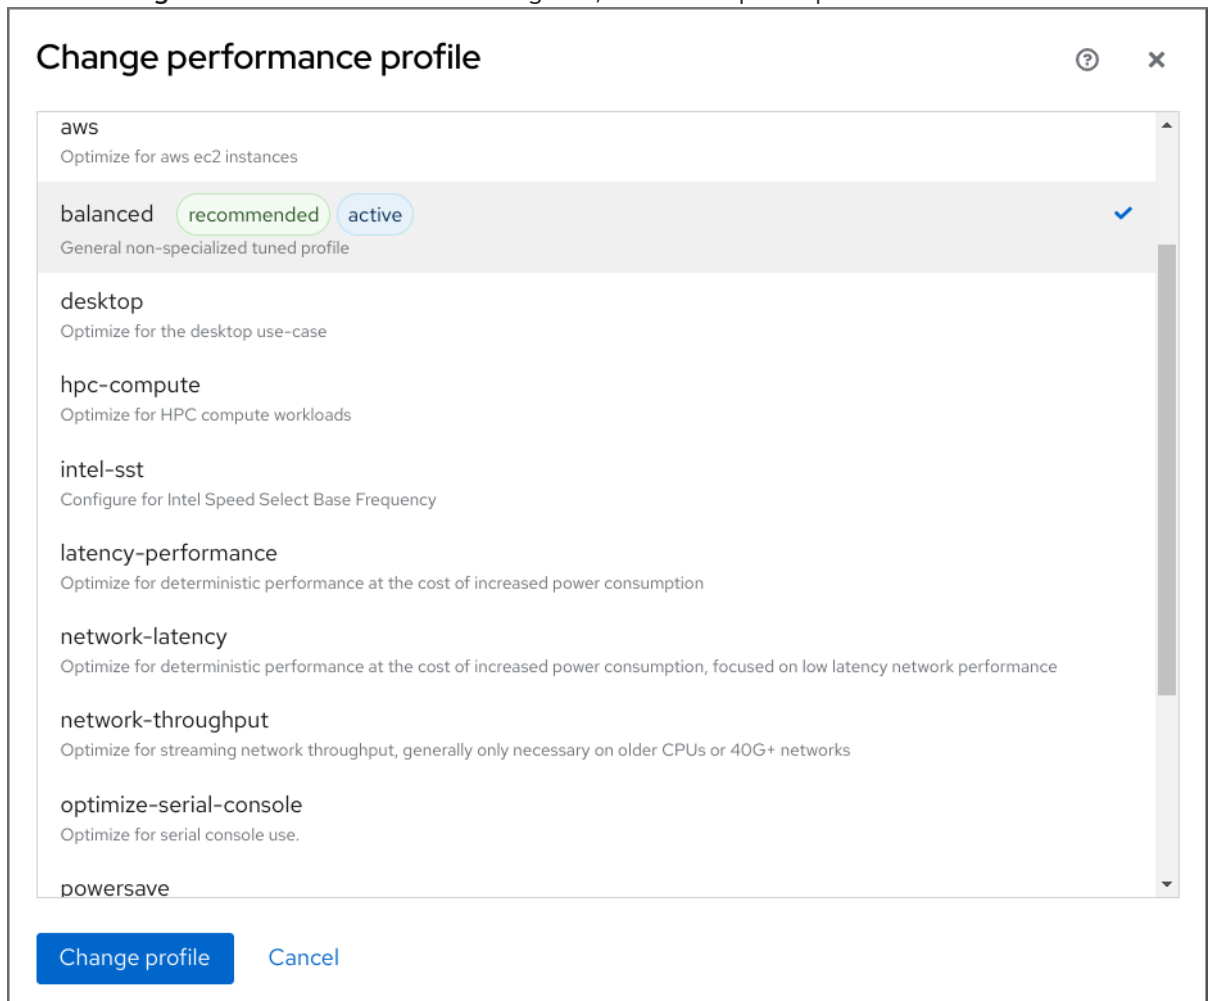
Prerequisites

- The web console interface must be installed and accessible.
  For details, see Installing the web console .

Procedure

1. Log in to the RHEL web console.
   For details, see Logging in to the web console .

2. Click **Logs**.

3. Use the drop-down menus to specify the three main quantifiers - time range, priority, and identifier(s) - you want to filter.
   The **Priority** quantifier always has to have a value. If you do not specify it, it automatically filters the **Error and above** priority. Notice that the options you set reflect in the text search box.

4. Specify the log field you want to filter.
   It is possible to add several log fields.

5. You can use a free-form string to search for anything else. The search box also accepts regular expressions.

5.5. OPTIONS FOR LOGS FILTERING

There are several **journalctl** options, which you can use for filtering logs in the web console, that may be useful. Some of these are already covered as part of the drop-down menus in the web console interface.

Table 5.1. Table

Option name	Usage	Notes

Option name	Usage	Notes
priority	Filter output by message priorities. Takes a single numeric or textual log level. The log levels are the usual syslog log levels. If a single log level is specified, all messages with this log level or a lower (therefore more important) log level are shown.	Covered in the **Priority** drop-down menu.
identifier	Show messages for the specified syslog identifier SYSLOG_IDENTIFIER. Can be specified multiple times.	Covered in the **Identifier** drop-down menu.
follow	Shows only the most recent journal entries, and continuously prints new entries as they are appended to the journal.	Not covered in a drop-down.
service	Show messages for the specified **systemd** unit. Can be specified multiple times.	Is not covered in a drop-down. Corresponds to the **journalctl --unit** parameter.
boot	Show messages from a specific boot. A positive integer will look up the boots starting from the beginning of the journal, and an equal-or-less-than zero integer will look up boots starting from the end of the journal. Therefore, 1 means the first boot found in the journal in chronological order, 2 the second and so on; while -0 is the last boot, -1 the boot before last, and so on.	Covered only as **Current boot** or **Previous boot** in the **Time** drop-down menu. Other options need to be written manually.

Option name	Usage	Notes
since	Start showing entries on or newer than the specified date, or on or older than the specified date, respectively. Date specifications should be of the format "2012-10-30 18:17:16". If the time part is omitted, "00:00:00" is assumed. If only the seconds component is omitted, ":00" is assumed. If the date component is omitted, the current day is assumed. Alternatively the strings "yesterday", "today", "tomorrow" are understood, which refer to 00:00:00 of the day before the current day, the current day, or the day after the current day, respectively. "now" refers to the current time. Finally, relative times may be specified, prefixed with "-" or "+", referring to times before or after the current time, respectively.	Not covered in a drop-down.

CHAPTER 6. MANAGING USER ACCOUNTS IN THE WEB CONSOLE

The RHEL web console offers an interface for adding, editing, and removing system user accounts.

After reading this section, you will know:

- From where the existing accounts come from.

- How to add new accounts.

- How to set password expiration.

- How and when to terminate user sessions.

Prerequisites

- Being logged into the RHEL web console with an account that has administrator permissions assigned. For details, see Logging in to the RHEL web console .

6.1. SYSTEM USER ACCOUNTS MANAGED IN THE WEB CONSOLE

With user accounts displayed in the RHEL web console you can:

- Authenticate users when accessing the system.

- Set the access rights to the system.

The RHEL web console displays all user accounts located in the system. Therefore, you can see at least one user account just after the first login to the web console.

After logging into the RHEL web console, you can perform the following operations:

- Create new users accounts.

- Change their parameters.

- Lock accounts.

- Terminate user sessions.

6.2. ADDING NEW ACCOUNTS USING THE WEB CONSOLE

Use the following steps for adding user accounts to the system and setting administration rights to the accounts through the RHEL web console.

Prerequisites

- The RHEL web console must be installed and accessible. For details, see Installing the web console.

Procedure

1. Log in to the RHEL web console.

2. Click **Accounts**.

3. Click **Create New Account**.

4. In the **Full Name** field, enter the full name of the user.
   The RHEL web console automatically suggests a user name from the full name and fills it in the **User Name** field. If you do not want to use the original naming convention consisting of the first letter of the first name and the whole surname, update the suggestion.

5. In the **Password/Confirm** fields, enter the password and retype it for verification that your password is correct.
   The color bar below the fields shows you the security level of the entered password, which does not allow you to create a user with a weak password.

6. Click **Create** to save the settings and close the dialog box.

7. Select the newly created account.

8. In the **Groups** drop down menu, select the groups that you want to add to the new account.

New User		Terminate session	Delete
Full name	New User		
User name	nuser		
Groups	nuser		▼
Last login	Never		
Options	☐ Disallow interactive password ⑦ Never expire account edit		
Password	Set password Force change Never expire password edit		

Now you can see the new account in the **Accounts** settings and you can use its credentials to connect to the system.

6.3. ENFORCING PASSWORD EXPIRATION IN THE WEB CONSOLE

By default, user accounts have set passwords to never expire. You can set system passwords to expire after a defined number of days. When the password expires, the next login attempt will prompt for a password change.

Procedure

1. Log in to the RHEL 8 web console.

2. Click **Accounts**.

3. Select the user account for which you want to enforce password expiration.

4. Click **edit** on the **Password** line.

Password	Set password Force change Require password change on March 2, 2024 edit

5. In the **Password expiration** dialog box, select **Require password change every ... days**and enter a positive whole number representing the number of days after which the password expires.

6. Click **Change**.
   The web console immediately shows the date of the future password change request on the **Password** line.

6.4. TERMINATING USER SESSIONS IN THE WEB CONSOLE

A user creates user sessions when logging into the system. Terminating user sessions means to log the user out from the system. It can be helpful if you need to perform administrative tasks sensitive to configuration changes, for example, system upgrades.

In each user account in the RHEL 8web console, you can terminate all sessions for the account except for the web console session you are currently using. This prevents you from loosing access to your system.

Procedure

1. Log in to the RHEL 8 web console.

2. Click **Accounts**.

3. Click the user account for which you want to terminate the session.

4. Click **Terminate Session**.
   If the **Terminate Session** button is inactive, the user is not logged in to the system.

   The RHEL web console terminates the sessions.

CHAPTER 7. MANAGING SERVICES IN THE WEB CONSOLE

Learn how to manage system services in the RHEL web console interface. You can activate or deactivate services, restart or reload them, or manage their automatic startup.

7.1. ACTIVATING OR DEACTIVATING SYSTEM SERVICES IN THE WEB CONSOLE

This procedure activates or deactivates system services using the web console interface.

Prerequisites

- The RHEL 8 web console has been installed.
  For details, see Installing the web console .

PROCEDURE

You can filter the services by name or description and also by Enabled, Disabled, or Static automatic startup. The interface shows the current state of the service and its recent logs.

1. Log in to the RHEL web console with administrator privileges.
   For details, see Logging in to the web console .

2. Click **Services** in the web console menu on the left.

3. The default tab for **Services** is **System Services**. If you want to manage targets, sockets, timers, or paths, switch to the respective tab in the menu on top.

4. To open service settings, click on a selected service from the list. You can tell which services are active or inactive by checking the **State** column.

5. Activate or deactivate a service:

   - To activate an inactive service, click the **Start** button.
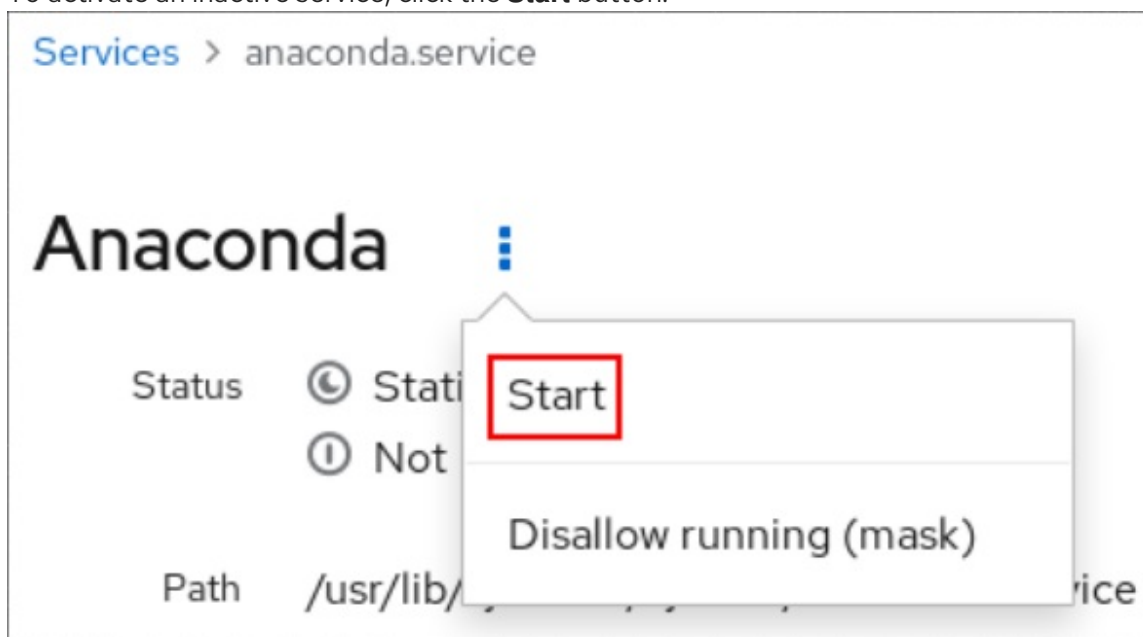
- To deactivate an active service, click the **Stop** button.
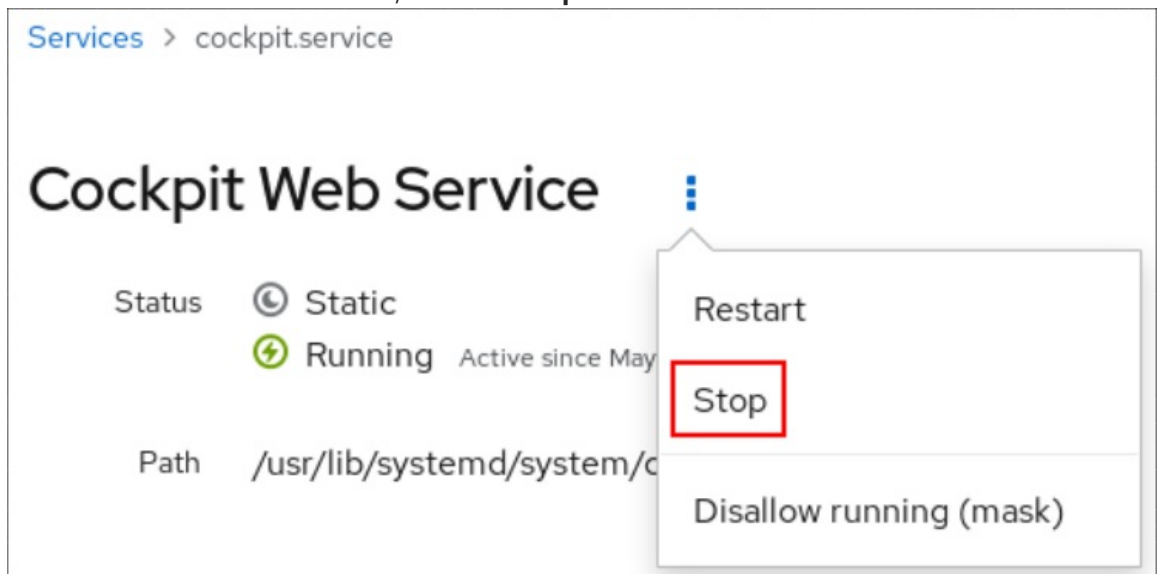
7.2. RESTARTING SYSTEM SERVICES IN THE WEB CONSOLE

This procedure restarts system services using the web console interface.

Prerequisites

- The RHEL 8 web console has been installed.
  For details, see Installing the web console.

PROCEDURE

You can filter the services by name or description and also by Enabled, Disabled, or Static automatic startup. The interface shows the current state of the service and its recent logs.

1. Log in to the RHEL web console with administrator privileges.
   For details, see Logging in to the web console .

2. Click **Services** in the web console menu on the left.

3. The default tab for **Services** is **System Services**. If you want to manage targets, sockets, timers, or paths, switch to the respective tab in the menu on top.

4. To open service settings, click on a selected service from the list.

5. To restart a service, click the **Restart** button.

7.3. OVERRIDING THE MANIFEST SETTINGS IN THE WEB CONSOLE

You can modify the menu of the web console for a particular user and all users of the system. In the **cockpit** project, a package name is a directory name. A package contains the **manifest.json** file along with other files. Default settings are present in the **manifest.json** file. You can override the default **cockpit** menu settings by creating a *<package-name>*.**override.json** file at a specific location for the specified user.

Prerequisites

- The RHEL 8 web console has been installed.
  For details, see Installing the web console.

Procedure

1. Override manifest settings in the ***<systemd>*.override.json** file in a text editor of your choice, for example:

   a. To edit for all users, enter:

      ```
      # vi /etc/cockpit/<systemd>.override.json
      ```

   b. To edit for a single user, enter:

      ```
      # vi ~/.config/cockpit/<systemd>.override.json
      ```

2. Edit the required file with the following details:

   ```
   {
     "menu": {
      "services": null,
      "logs": {
         "order": -1
      }
     }
   }
   ```

   - The **null** value hides the **services** tab

   - The **-1** value moves the **logs** tab to the first place.

3. Restart the **cockpit** service:

   ```
   # systemctl restart cockpit.service
   ```

Additional resources

- **cockpit(1)** man page

- Manifest overrides

CHAPTER 8. CONFIGURING NETWORK BONDS USING THE WEB CONSOLE

Learn how network bonding works and configure network bonds in the RHEL 8 web console.

> **NOTE**
>
> The RHEL 8 web console is build on top of the NetworkManager service.

Prerequisites

- The RHEL 8 web console installed and enabled. For details, see Installing the web console .

8.1. UPSTREAM SWITCH CONFIGURATION DEPENDING ON THE BONDING MODES

Depending on the bonding mode you want to use, you must configure the ports on the switch:

Bonding mode	Configuration on the switch
0 - **balance-rr**	Requires static EtherChannel enabled, not Link Aggregation Control Protocol (LACP)-negotiated.
1 - **active-backup**	No configuration required on the switch.
2 - **balance-xor**	Requires static EtherChannel enabled, not LACP-negotiated.
3 - **broadcast**	Requires static EtherChannel enabled, not LACP-negotiated.
4 - **802.3ad**	Requires LACP-negotiated EtherChannel enabled.
5 - **balance-tlb**	No configuration required on the switch.
6 - **balance-alb**	No configuration required on the switch.

For details how to configure your switch, see the documentation of the switch.

> **IMPORTANT**
>
> Certain network bonding features, such as the fail-over mechanism, do not support direct cable connections without a network switch. For further details, see the Is bonding supported with direct connection using crossover cables? KCS solution.

8.2. BOND MODES

In RHEL 8 there are several mode options. Each mode option is characterize by specific load balancing and fault tolerance. The behavior of the bonded interfaces depends upon the mode. The bonding modes provide fault tolerance, load balancing or both.

Load balancing modes

- **Round Robin**: Sequentially transmit packets from the first available interface to the last one.

Fault tolerance modes

- **Active Backup**: Only when the primary interface fails, one of a backup interfaces replaces it. Only a MAC address used by active interface is visible.

- **Broadcast**: All transmissions are sent on all interfaces.

> **NOTE**
>
> Broadcasting significantly increases network traffic on all the bonded interfaces.

Fault tolerance and load balancing modes

- **XOR**: The destination MAC addresses are distributed equally between interfaces with a modulo hash. Each interface then serves the same group of MAC addresses.

- **802.3ad**: Sets an IEEE 802.3ad dynamic link aggregation policy. Creates aggregation groups that share the same speed and duplex settings. Transmits and receives on all interfaces in the active aggregator.
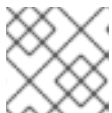
> **NOTE**
>
> This mode requires a switch that is 802.3ad compliant.

- **Adaptive transmit load balancing**: The outgoing traffic is distributed according to the current load on each interface. Incoming traffic is received by the current interface. If the receiving interface fails, another interface takes over the MAC address of the failed one.

- **Adaptive load balancing**: Includes transmit and receive load balancing for IPv4 traffic. Receive load balancing is achieved through Address Resolution Protocol (ARP) negotiation, therefore, it is necessary to set **Link Monitoring** to **ARP** in the bond's configuration.

8.3. CONFIGURING A NETWORK BOND BY USING THE RHEL WEB CONSOLE

Use the RHEL web console to configure a network bond if you prefer to manage network settings using a web browser-based interface.

Prerequisites

- You are logged in to the RHEL web console.

- Two or more physical or virtual network devices are installed on the server.

- To use Ethernet devices as members of the bond, the physical or virtual Ethernet devices must be installed on the server.

- To use team, bridge, or VLAN devices as members of the bond, create them in advance as described in:

- Configuring a network team by using the RHEL web console

- Configuring a network bridge by using the RHEL web console

- Configuring VLAN tagging by using the RHEL web console

Procedure

1. Select the **Networking** tab in the navigation on the left side of the screen.

2. Click **Add bond** in the **Interfaces** section.

3. Enter the name of the bond device you want to create.

4. Select the interfaces that should be members of the bond.

5. Select the mode of the bond.
   If you select **Active backup**, the web console shows the additional field **Primary** in which you can select the preferred active device.

6. Set the link monitoring mode. For example, when you use the **Adaptive load balancing** mode, set it to **ARP**.

7. Optional: Adjust the monitoring interval, link up delay, and link down delay settings. Typically, you only change the defaults for troubleshooting purposes.

Bond settings

Name	bond0
Interfaces	☑ enp7s0
	☑ enp8s0
MAC	▼
Mode	Active backup ▼
Primary	enp7s0 ▼
Link monitoring	MII (recommended) ▼
Monitoring interval	100
Link up delay	0
Link down delay	0

Apply   Cancel

8. Click **Apply**.

9. By default, the bond uses a dynamic IP address. If you want to set a static IP address:

   a. Click the name of the bond in the **Interfaces** section.

   b. Click **Edit** next to the protocol you want to configure.

   c. Select **Manual** next to **Addresses**, and enter the IP address, prefix, and default gateway.

   d. In the **DNS** section, click the **+** button, and enter the IP address of the DNS server. Repeat this step to set multiple DNS servers.

   e. In the **DNS search domains** section, click the **+** button, and enter the search domain.

f. If the interface requires static routes, configure them in the **Routes** section.

IPv4 settings   ✕

| Addresses | | | Manual ▼ | ＋ |

Address	Prefix length or netmask	Gateway	
192.0.2.1	24	192.0.2.254	−

DNS   ⬤ Automatic   ＋

Server

192.0.2.253   −

DNS search domains   ⬤ Automatic   ＋

Search domain

example.com   −

Routes   ⬤ Automatic   ＋

Apply   Cancel

g. Click **Apply**

Verification

1. Select the **Networking** tab in the navigation on the left side of the screen, and check if there is incoming and outgoing traffic on the interface:

Interfaces   Add bond   Add team   Add bridge   Add VLAN

Name	IP address	Sending	Receiving
bond0	192.0.2.1/24	1.11 Mbps	61.2 Mbps

2. Temporarily remove the network cable from the host.
    Note that there is no method to properly test link failure events using software utilities. Tools that deactivate connections, such as the web console, show only the bonding driver's ability to handle member configuration changes and not actual link failure events.

3. Display the status of the bond:

    ```
    # cat /proc/net/bonding/bond0
    ```

8.4. ADDING INTERFACES TO THE BOND USING THE WEB CONSOLE

Network bonds can include multiple interfaces and you can add or remove any of them at any time.

Learn how to add a network interface to an existing bond.

Prerequisites

- Having a bond with multiple interfaces configured as described in Configuring a network bond using the web console

Procedure

1. Log in to the web console.
   For details, see Logging in to the web console .

2. Open **Networking**.

3. In the **Interfaces** table, click on the bond you want to configure.

4. In the bond settings screen, scroll down to the table of members (interfaces).

5. Click the **Add member** drop down icon.

6. Select the interface in the drop down menu and click it.

Verification steps

- Check that the selected interface appeared in the **Interface members** table in the bond settings screen.

8.5. REMOVING OR DISABLING AN INTERFACE FROM THE BOND USING THE WEB CONSOLE

Network bonds can include multiple interfaces. If you need to change a device, you can remove or disable particular interfaces from the bond, which will work with the rest of the active interfaces.

To stop using an interface included in a bond, you can:

- Remove the interface from the bond.

- Disable the interface temporarily. The interface stays a part of the bond, but the bond will not use it until you enable it again.

Prerequisites

- Having a bond with multiple interfaces configured as described in Configuring a network bond using the web console

Procedure

1. Log in to the RHEL web console. For details, see Logging in to the web console .

2. Open **Networking**.

3. Click the bond you want to configure.

4. In the bond settings screen, scroll down to the table of ports (interfaces).

5. Select the interface and remove or disable it:

   - To remove the interface, click the **-** button.

   - To disable or enable the interface, toggle the switch next to the selected interface.

Based on your choice, the web console either removes or disables the interface from the bond and you can see it back in the **Networking** section as a standalone interface.

8.6. REMOVING OR DISABLING A BOND USING THE WEB CONSOLE

Remove or disable a network bond using the web console. If you disable the bond, the interfaces stay in the bond, but the bond will not be used for network traffic.

Prerequisites

- There is an existing bond in the web console.

Procedure

1. Log in to the web console.
   For details, see Logging in to the web console .

2. Open **Networking**.

3. Click the bond you want to remove.

4. In the bond settings screen, you can disable or enable the bond by toggling a switcher or click the **Delete** button to remove the bond permanently.

bond0 Bond 52:54:00:2F:6B:2E	Delete	☑️
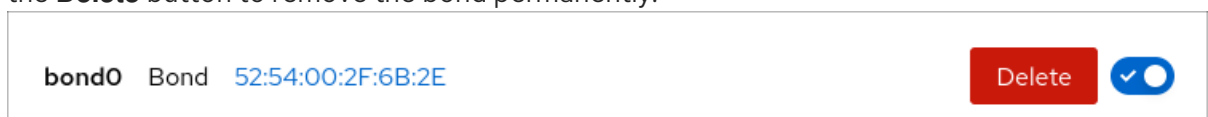
Verification steps

- Go back to **Networking** and verify that all the interfaces from the bond are now standalone interfaces.

CHAPTER 9. CONFIGURING NETWORK TEAMS USING THE WEB CONSOLE

Learn how network bonding works, what are the differences between network teams and network bonds, and what are the possibilities of configuration in the web console.

Additionally you can find guidelines for:

- Adding a new network team

- Adding new interfaces to an existing network team

- Removing interfaces from an existing network team

- Removing a network team

> **IMPORTANT**
>
> Network teaming is deprecated in Red Hat Enterprise Linux 9. If you plan to upgrade your server to a future version of RHEL, consider using the kernel bonding driver as an alternative. For details, see Configuring network bonding.

Prerequisites

- The RHEL web console installed and enabled.
  For details, see Installing the web console.

9.1. CONFIGURING A NETWORK TEAM BY USING THE RHEL WEB CONSOLE

Use the RHEL web console to configure a network team if you prefer to manage network settings using a web browser-based interface.

> **IMPORTANT**
>
> Network teaming is deprecated in Red Hat Enterprise Linux 9. If you plan to upgrade your server to a future version of RHEL, consider using the kernel bonding driver as an alternative. For details, see Configuring network bonding.

Prerequisites

- The **teamd** and **NetworkManager-team** packages are installed.

- Two or more physical or virtual network devices are installed on the server.

- To use Ethernet devices as ports of the team, the physical or virtual Ethernet devices must be installed on the server and connected to a switch.

- To use bond, bridge, or VLAN devices as ports of the team, create them in advance as described in:

  - Configuring a network bond by using the RHEL web console

  - Configuring a network bridge by using the RHEL web console

- Configuring VLAN tagging by using the RHEL web console

Procedure

1. Select the **Networking** tab in the navigation on the left side of the screen.

2. Click **Add team** in the **Interfaces** section.

3. Enter the name of the team device you want to create.

4. Select the interfaces that should be ports of the team.

5. Select the runner of the team.
   If you select **Load balancing** or **802.3ad LACP**, the web console shows the additional field **Balancer**.

6. Set the link watcher:

   - If you select **Ethtool**, additionally, set a link up and link down delay.

   - If you set **ARP ping** or **NSNA ping**, additionally, set a ping interval and ping target.

7. Click **Apply**.

8. By default, the team uses a dynamic IP address. If you want to set a static IP address:

   a. Click the name of the team in the **Interfaces** section.

   b. Click **Edit** next to the protocol you want to configure.

   c. Select **Manual** next to **Addresses**, and enter the IP address, prefix, and default gateway.

   d. In the **DNS** section, click the **+** button, and enter the IP address of the DNS server. Repeat this step to set multiple DNS servers.

   e. In the **DNS search domains** section, click the **+** button, and enter the search domain.

   f. If the interface requires static routes, configure them in the **Routes** section.

   g. Click **Apply**

Verification

1. Select the **Networking** tab in the navigation on the left side of the screen, and check if there is incoming and outgoing traffic on the interface.

Interfaces	Add bond	Add team	Add bridge	Add VLAN
Name	**IP address**	**Sending**	**Receiving**	
team0	192.0.2.1/24	1.11 Mbps	61.2 Mbps	

2. Display the status of the team:

```
# teamdctl team0 state
setup:
  runner: activebackup
ports:
  enp7s0
    link watches:
      link summary: up
      instance[link_watch_0]:
        name: ethtool
        link: up
        down count: 0
  enp8s0
    link watches:
      link summary: up
      instance[link_watch_0]:
        name: ethtool
        link: up
        down count: 0
runner:
  active port: enp7s0
```

In this example, both ports are up.

Additional resources

- [Network team runners](#)

9.2. ADDING NEW INTERFACES TO THE TEAM USING THE WEB CONSOLE

Network teams can include multiple interfaces and it is possible to add or remove any of them at any time. The following section describes how to add a new network interface to an existing team.

Prerequisites

- A network team with is configured.

Procedure

1. Log in to the web console.
   For details, see [Logging in to the web console](#) .

2. Switch to the **Networking** tab.

3. In the **Interfaces** table, click on the team you want to configure.

4. In the team settings window, scroll down to the **Ports** table.

5. Click on the **+** button.

6. Select the interface you want to add from the drop down list.

The RHEL web console adds the interface to the team.

9.3. REMOVING OR DISABLING AN INTERFACE FROM THE TEAM USING THE WEB CONSOLE

Network teams can include multiple interfaces. If you need to change a device, you can remove or disable particular interfaces from the network team, which will work together with the rest of active interfaces.

There are two options how to stop using an interface included in a team:

- Removing the interface from the team

- Temporarily disabling the interface. The interface then stays as part of the team, but the team will not use it until you enable it again.

Prerequisites

- A network team with multiple interfaces exists on the host.

Procedure

1. Log in to the RHEL web console.
   For details, see Logging in to the web console .

2. Switch to the **Networking** tab.

3. Click the team you want to configure.

4. In the team settings window, scroll down to the table of ports (interfaces).

5. Select an interface and remove or disable it.

   a. Switch the **ON/OFF** button to Off to disable the interface.

   b. Click the **-** button to remove the interface.

Ports	Sending	Receiving	+
enp7s0	0 bps	0 bps	
enp8s0	0 bps	0 bps	
enp9s0	0 bps	0 bps	

Based on your choice, the web console either removes or disables the interface. If you remove the interface, it will be available in **Networking** as a standalone interface.

9.4. REMOVING OR DISABLING A TEAM USING THE WEB CONSOLE

Remove or disable a network team using the web console. If you only disable the team, interfaces in the team will stay in it but the team will not be used for network traffic.

Prerequisites

- A network team is configured on the host.

Procedure

1. Log in to the web console.
   For details, see Logging in to the web console .

2. Switch to the **Networking** tab.

3. Click the team you want to remove or disable.

4. Remove or disable the selected team.

   a. You can remove the team by clicking the **Delete** button.

   b. You can disable the team by moving the **ON/OFF** switch to a disabled position.

myteam	Team	52:54:00:25:A9:7D	Delete
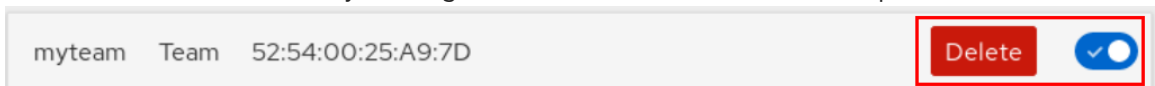
Verification steps

- If you removed the team, go to **Networking**, and verify that all the interfaces from your team are now listed as standalone interfaces.

CHAPTER 10. CONFIGURING NETWORK BRIDGES IN THE WEB CONSOLE

Network bridges are used to connect multiple interfaces to the one subnet with the same range of IP addresses.

Prerequisites

- The RHEL 8 web console installed and enabled.
  For details, see Installing the web console.

10.1. CONFIGURING A NETWORK BRIDGE BY USING THE RHEL WEB CONSOLE

Use the RHEL web console to configure a network bridge if you prefer to manage network settings using a web browser-based interface.

Prerequisites

- Two or more physical or virtual network devices are installed on the server.

- To use Ethernet devices as ports of the bridge, the physical or virtual Ethernet devices must be installed on the server.

- To use team, bond, or VLAN devices as ports of the bridge, you can either create these devices while you create the bridge or you can create them in advance as described in:

    - Configuring a network team using the RHEL web console

    - Configuring a network bond by using the RHEL web console

    - Configuring VLAN tagging by using the RHEL web console

Procedure

1. Select the **Networking** tab in the navigation on the left side of the screen.

2. Click **Add bridge** in the **Interfaces** section.

3. Enter the name of the bridge device you want to create.

4. Select the interfaces that should be ports of the bridge.

5. Optional: Enable the **Spanning tree protocol (STP)** feature to avoid bridge loops and broadcast radiation.

Bridge settings ✕

Name	bridge0

Ports   ☑ enp7s0

☑ enp8s0

Options   ☐ Spanning tree protocol (STP)

Apply   Cancel

6. Click **Apply**.

7. By default, the bridge uses a dynamic IP address. If you want to set a static IP address:

   a. Click the name of the bridge in the **Interfaces** section.

   b. Click **Edit** next to the protocol you want to configure.

   c. Select **Manual** next to **Addresses**, and enter the IP address, prefix, and default gateway.

   d. In the **DNS** section, click the **+** button, and enter the IP address of the DNS server. Repeat this step to set multiple DNS servers.

   e. In the **DNS search domains** section, click the **+** button, and enter the search domain.

   f. If the interface requires static routes, configure them in the **Routes** section.

g. Click **Apply**

Verification

1. Select the **Networking** tab in the navigation on the left side of the screen, and check if there is incoming and outgoing traffic on the interface:

10.2. REMOVING INTERFACES FROM THE BRIDGE USING THE WEB CONSOLE

Network bridges can include multiple interfaces. You can remove them from the bridge. Each removed interface will be automatically changed to the standalone interface.

Learn how to remove a network interface from a software bridge created in the RHEL 8 system.

Prerequisites

- Having a bridge with multiple interfaces in your system.

Procedure

1. Log in to the RHEL web console. For details, see Logging in to the web console .

2. Open **Networking**.

3. Click the bridge you want to configure.

4. In the bridge settings screen, scroll down to the table of ports (interfaces).

5. Select an interface and click the **-** button.

Verification steps

- Go to **Networking** to check that you can see the interface as a standalone interface in the **Interface members** table.

10.3. DELETING BRIDGES IN THE WEB CONSOLE

You can delete a software network bridge in the RHEL web console. All network interfaces included in the bridge will be changed automatically to standalone interfaces.

Prerequisites

- Having a bridge in your system.

Procedure

1. Log in to the RHEL web console.
   For details, see Logging in to the web console .

2. Open the **Networking** section.

3. Click the bridge you want to configure.

4. Click **Delete**.

   | **bridge0** Bridge 52:54:00:2F:6B:2E | Delete |
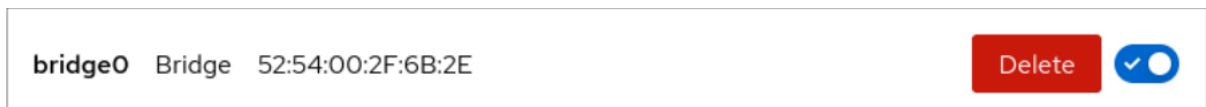
Verification steps

- Go back to **Networking** and verify that all the network interfaces are displayed in the **Interface members** table.

Some interfaces that were previously part of the bridge can become inactive. If necessary, activate them and set network parameters manually.

CHAPTER 11. CONFIGURING VLANS IN THE WEB CONSOLE

This section describes how to configure Virtual Local Area Network (VLAN). A VLAN is a logical network within a physical network. The VLAN interface tags packets with the VLAN ID as they pass through the interface, and removes tags of returning packets.

11.1. CONFIGURING VLAN TAGGING BY USING THE RHEL WEB CONSOLE
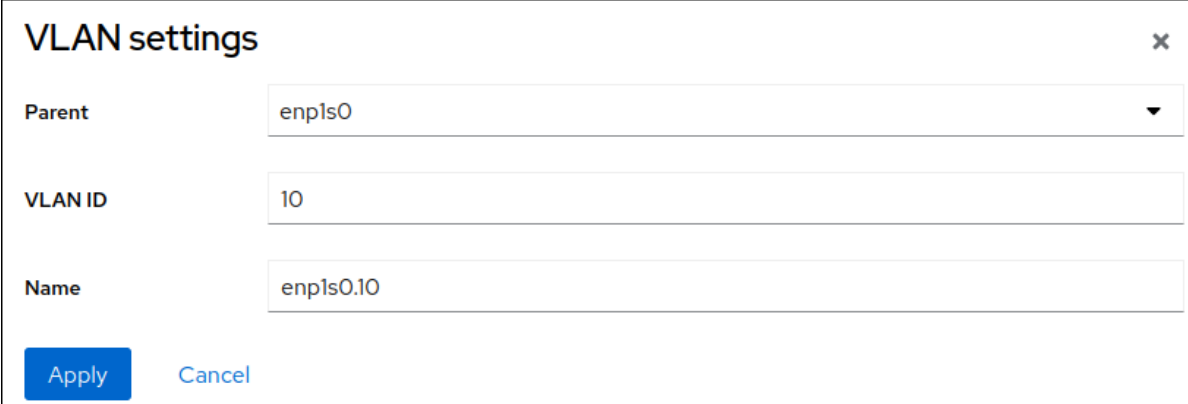
Use the RHEL web console to configure VLAN tagging if you prefer to manage network settings using a web browser-based interface.

Prerequisites

- The interface you plan to use as a parent to the virtual VLAN interface supports VLAN tags.

- If you configure the VLAN on top of a bond interface:

  - The ports of the bond are up.

  - The bond is not configured with the **fail_over_mac=follow** option. A VLAN virtual device cannot change its MAC address to match the parent's new MAC address. In such a case, the traffic would still be sent with the incorrect source MAC address.

  - The bond is usually not expected to get IP addresses from a DHCP server or IPv6 auto-configuration. Ensure it by disabling the IPv4 and IPv6 protocol creating the bond. Otherwise, if DHCP or IPv6 auto-configuration fails after some time, the interface might be brought down.

- The switch, the host is connected to, is configured to support VLAN tags. For details, see the documentation of your switch.

Procedure

1. Select the **Networking** tab in the navigation on the left side of the screen.

2. Click **Add VLAN** in the **Interfaces** section.

3. Select the parent device.

4. Enter the VLAN ID.

5. Enter the name of the VLAN device or keep the automatically-generated name.

6. Click **Apply**.

7. By default, the VLAN device uses a dynamic IP address. If you want to set a static IP address:

   a. Click the name of the VLAN device in the **Interfaces** section.

   b. Click **Edit** next to the protocol you want to configure.

   c. Select **Manual** next to **Addresses**, and enter the IP address, prefix, and default gateway.

   d. In the **DNS** section, click the **+** button, and enter the IP address of the DNS server. Repeat this step to set multiple DNS servers.

   e. In the **DNS search domains** section, click the **+** button, and enter the search domain.

   f. If the interface requires static routes, configure them in the **Routes** section.

   g. Click **Apply**

Verification

- Select the **Networking** tab in the navigation on the left side of the screen, and check if there is incoming and outgoing traffic on the interface:

Interfaces	Add bond	Add team	Add bridge	Add VLAN

Name	IP address	Sending	Receiving
enp1s0.10	192.0.2.1/24	1.11 Mbps	61.2 Mbps

CHAPTER 12. CONFIGURING THE WEB CONSOLE LISTENING PORT

Learn how to allow new ports or change the existing ports using the RHEL 8 web console.

12.1. ALLOWING A NEW PORT ON A SYSTEM WITH ACTIVE SELINUX

Enable the web console to listen on a selected port.

Prerequisites

- The web console must be installed and accessible. For details, see Installing the web console.

Procedure

- For ports that are not defined by any other part of SELinux, run:

  ```
  $ sudo semanage port -a -t websm_port_t -p tcp PORT_NUMBER
  ```

- For ports that already are defined by other part of SELinux, run:

  ```
  $ sudo semanage port -m -t websm_port_t -p tcp PORT_NUMBER
  ```

The changes should take effect immediately.

12.2. ALLOWING A NEW PORT ON A SYSTEM WITH FIREWALLD

Enable the web console to receive connections on a new port.

Prerequisites

- The web console must be installed and accessible. For details, see Installing the web console.

- The **firewalld** service must be running.

Procedure

1. To add a new port number, run the following command:

   ```
   $ sudo firewall-cmd --permanent --service cockpit --add-port=PORT_NUMBER/tcp
   ```

2. To remove the old port number from the **cockpit** service, run:

   ```
   $ sudo firewall-cmd --permanent --service cockpit --remove-port=OLD_PORT_NUMBER/tcp
   ```

> **IMPORTANT**
>
> If you only run the **firewall-cmd --service cockpit --add-port=PORT_NUMBER/tcp** without the **--permanent** option, your change will disappear with the next reload of **firewalld** or a system reboot.

12.3. CHANGING THE WEB CONSOLE PORT

Change default transmission control protocol (TCP) on port **9090** to a different one.

Prerequisites

- The web console must be installed and accessible. For details, see Installing the web console .

- If you have SELinux protecting your system, you need to set it to allow Cockpit to listen on a new port. For more information, see Allowing a new port on a system with active SELinux .

- If you have **firewalld** configured as your firewall, you need to set it to allow Cockpit receive connections on a new port, for more information, see Allowing a new port on a system with **firewalld**.

Procedure

1. Change the listening port with one of the following methods:

   a. Using the **systemctl edit cockpit.socket** command:

      i. Run the following command:

         ```
         $ sudo systemctl edit cockpit.socket
         ```

         This will open the **/etc/systemd/system/cockpit.socket.d/override.conf** file.

      ii. Modify the content of **override.conf** or add a new content in the following format:

         ```
         [Socket]
         ListenStream=
         ListenStream=PORT_NUMBER
         ```

         The **ListenStream** option specifies the desired address and TCP port.

         > **NOTE**
         >
         > The first line with an empty value is intentional. **systemd** allows multiple **ListenStream** directives to be declared in a single socket unit. An empty value in a drop-in file resets the list and disables the default port 9090 from the original unit.

   b. Alternatively, add the above mentioned content to the **/etc/systemd/system/cockpit.socket.d/listen.conf** file.
      Create the **cockpit.socket.d.** directory and the **listen.conf** file if they do not exist yet.

2. Run the following commands for changes to take effect:

   ```
   $ sudo systemctl daemon-reload
   $ sudo systemctl restart cockpit.socket
   ```

   If you used **systemctl edit cockpit.socket** in the previous step, running **systemctl daemon-reload** is not necessary.

Verification steps

Verification steps

- To verify that the change was successful, try to connect to the web console with the new port.

CHAPTER 13. MANAGING FIREWALL USING THE WEB CONSOLE

A firewall is a way to protect machines from any unwanted traffic from outside. It enables users to control incoming network traffic on host machines by defining a set of firewall rules. These rules are used to sort the incoming traffic and either block it or allow through.

Prerequisites

- The RHEL 8 web console configures the **firewalld** service.
  For details about the **firewalld** service, see Getting started with firewalld.

13.1. RUNNING FIREWALL USING THE WEB CONSOLE

The following steps show where and how to run the RHEL 8 system firewall in the web console.

> **NOTE**
>
> The RHEL 8 web console configures the **firewalld** service.

Procedure

1. Log in to the RHEL 8 web console. For details, see Logging in to the web console .

2. Open the **Networking** section.

3. In the **Firewall** section, click the slider to run the firewall.
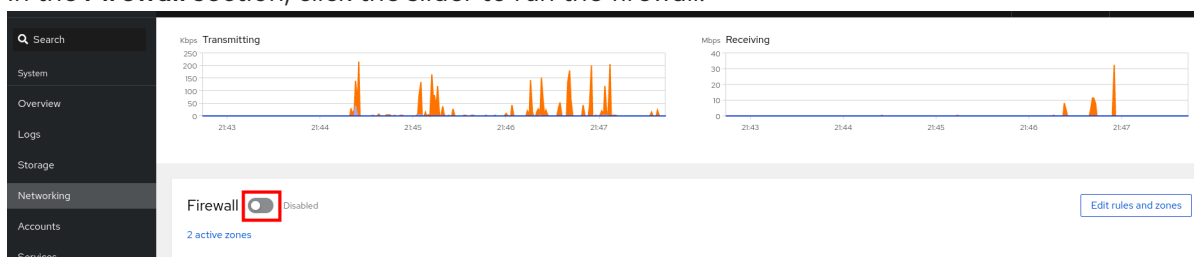
   If you do not see the **Firewall** slider, log in to the web console with the administrative privileges.

At this stage, your firewall is running.

To configure firewall rules, see Enabling services on the firewall using the web console .

13.2. STOPPING FIREWALL USING THE WEB CONSOLE

The following steps show where and how to stop the RHEL 8 system firewall in the web console.
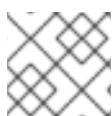
> **NOTE**
>
> The RHEL 8 web console configures the **firewalld** service.

Procedure

1. Log in to the RHEL 8 web console. For details, see Logging in to the web console .

2. Open the **Networking** section.

3. In the **Firewall** section, click the slider to stop the firewall.

If you do not see the **Firewall** slider, log in to the web console with the administrative privileges.

At this stage, the firewall has been stopped and does not secure your system.

13.3. FIREWALL ZONES

You can use the **firewalld** utility to separate networks into different zones according to the level of trust that you have with the interfaces and traffic within that network. A connection can only be part of one zone, but you can use that zone for many network connections.

firewalld follows strict principles in regards to zones:

1. Traffic ingresses only one zone.

2. Traffic egresses only one zone.

3. A zone defines a level of trust.

4. Intrazone traffic (within the same zone) is allowed by default.

5. Interzone traffic (from zone to zone) is denied by default.

Principles 4 and 5 are a consequence of principle 3.

Principle 4 is configurable through the zone option **--remove-forward**. Principle 5 is configurable by adding new policies.

NetworkManager notifies **firewalld** of the zone of an interface. You can assign zones to interfaces with the following utilities:

- **NetworkManager**

- **firewall-config** utility

- **firewall-cmd** utility

- The RHEL web console

The RHEL web console, **firewall-config**, and **firewall-cmd** can only edit the appropriate **NetworkManager** configuration files. If you change the zone of the interface using the web console, **firewall-cmd**, or **firewall-config**, the request is forwarded to **NetworkManager** and is not handled by **firewalld**.

The **/usr/lib/firewalld/zones/** directory stores the predefined zones, and you can instantly apply them to any available network interface. These files are copied to the **/etc/firewalld/zones/** directory only after they are modified. The default settings of the predefined zones are as follows:

block

- Suitable for: Any incoming network connections are rejected with an icmp-host-prohibited message for **IPv4** and icmp6-adm-prohibited for **IPv6**.

- Accepts: Only network connections initiated from within the system.

dmz

- Suitable for: Computers in your DMZ that are publicly-accessible with limited access to your internal network.

- Accepts: Only selected incoming connections.

drop

Suitable for: Any incoming network packets are dropped without any notification.

**Accepts: Only outgoing network connections.

external

- Suitable for: External networks with masquerading enabled, especially for routers. Situations when you do not trust the other computers on the network.

- Accepts: Only selected incoming connections.

home

- Suitable for: Home environment where you mostly trust the other computers on the network.

- Accepts: Only selected incoming connections.

internal

- Suitable for: Internal networks where you mostly trust the other computers on the network.

- Accepts: Only selected incoming connections.

public

- Suitable for: Public areas where you do not trust other computers on the network.

- Accepts: Only selected incoming connections.

trusted

- Accepts: All network connections.

work

Suitable for: Work environment where you mostly trust the other computers on the network.

- Accepts: Only selected incoming connections.

One of these zones is set as the *default* zone. When interface connections are added to **NetworkManager**, they are assigned to the default zone. On installation, the default zone in **firewalld** is the **public** zone. You can change the default zone.

> **NOTE**
>
> Make network zone names self-explanatory to help users understand them quickly.

To avoid any security problems, review the default zone configuration and disable any unnecessary services according to your needs and risk assessments.

Additional resources

- The **firewalld.zone(5)** man page.

13.4. ZONES IN THE WEB CONSOLE

The Red Hat Enterprise Linux web console implements major features of the firewalld service and enables you to:

- Add predefined firewall zones to a particular interface or range of IP addresses

- Configure zones with selecting services into the list of enabled services

- Disable a service by removing this service from the list of enabled service

- Remove a zone from an interface

13.5. ENABLING ZONES USING THE WEB CONSOLE

You can apply predefined and existing firewall zones on a particular interface or a range of IP addresses through the RHEL web console.

Prerequisites

- The RHEL 8 web console has been installed. For details, see Installing the web console.

- The firewall must be enabled. For details, see Running firewall using the web console.

Procedure

1. Log in to the RHEL web console with administrative privileges. For details, see Logging in to the web console.

2. Click **Networking**.

3. Click on the **Edit rules and zones** button.

If you do not see the **Edit rules and zones** button, log in to the web console with the administrator privileges.

4. In the **Firewall** section, click **Add new zone**.

5. In the **Add zone** dialog box, select a zone from the **Trust level** options.
   The web console displays all zones predefined in the **firewalld** service.

6. In the **Interfaces** part, select an interface or interfaces on which the selected zone is applied.

7. In the **Allowed Addresses** part, you can select whether the zone is applied on:

   - the whole subnet

   - or a range of IP addresses in the following format:

     - 192.168.1.0

     - 192.168.1.0/24

     - 192.168.1.0/24, 192.168.1.0

8. Click on the **Add zone** button.

Verification

- Check the configuration in the **Firewall** section:

13.6. ENABLING SERVICES ON THE FIREWALL USING THE WEB CONSOLE

By default, services are added to the default firewall zone. If you use more firewall zones on more network interfaces, you must select a zone first and then add the service with port.

The RHEL 8 web console displays predefined **firewalld** services and you can add them to active firewall zones.

> **IMPORTANT**
>
> The RHEL 8 web console configures the **firewalld** service.
>
> The web console does not allow generic **firewalld** rules which are not listed in the web console.

Prerequisites

- The RHEL 8 web console has been installed. For details, see Installing the web console .

- The firewall must be enabled. For details, see Running firewall using the web console .

Procedure

1. Log in to the RHEL web console with administrator privileges. For details, see Logging in to the web console.

2. Click **Networking**.

3. Click on the **Edit rules and zones** button.

   If you do not see the **Edit rules and zones** button, log in to the web console with the administrator privileges.

4. In the **Firewall** section, select a zone for which you want to add the service and click **Add Services**.
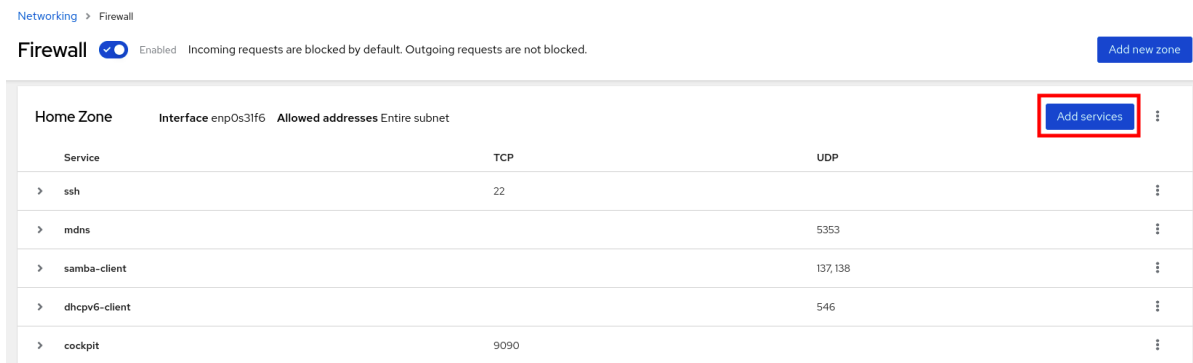
5. In the **Add Services** dialog box, find the service you want to enable on the firewall.

6. Enable desired services.
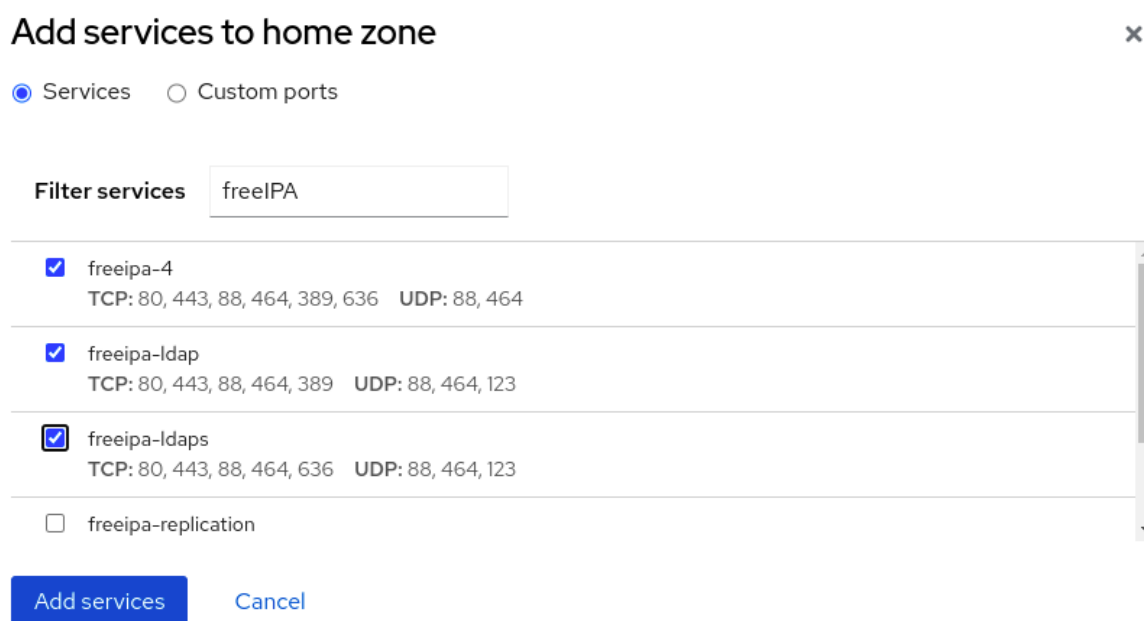
7. Click **Add Services**.

At this point, the RHEL 8 web console displays the service in the zone's list of **Services**.

13.7. CONFIGURING CUSTOM PORTS USING THE WEB CONSOLE

The web console allows you to add:

- Services listening on standard ports: Enabling services on the firewall using the web console

- Services listening on custom ports.

You can add services by configuring custom ports as described.

Prerequisites

- The RHEL 8 web console has been installed. For details, see Installing the web console .

- The firewall must be enabled. For details, see Running firewall using the web console .

Procedure

1. Log in to the RHEL web console with administrator privileges. For details, see Logging in to the web console.

2. Click **Networking**.

3. Click on the **Edit rules and zones** button.

If you do not see the **Edit rules and zones** button, log in to the web console with the administrative privileges.

4. In the **Firewall** section, select a zone for which you want to configure a custom port and click **Add Services**.

5. In the **Add services** dialog box, click on the **Custom Ports** radio button.

6. In the TCP and UDP fields, add ports according to examples. You can add ports in the following formats:

   - Port numbers such as 22

   - Range of port numbers such as 5900-5910

   - Aliases such as nfs, rsync

   > **NOTE**
   >
   > You can add multiple values into each field. Values must be separated with the comma and without the space, for example: 8080,8081,http

7. After adding the port number in the **TCP** filed, the **UDP** filed, or both, verify the service name in the **Name** field.
   The **Name** field displays the name of the service for which is this port reserved. You can rewrite the name if you are sure that this port is free to use and no server needs to communicate on this port.

8. In the **Name** field, add a name for the service including defined ports.

9. Click on the **Add Ports** button.

To verify the settings, go to the **Firewall** page and find the service in the list of zone's **Services**.

13.8. DISABLING ZONES USING THE WEB CONSOLE

You can disable a firewall zone in your firewall configuration using the web console.

Prerequisites

- The RHEL 8 web console has been installed. For details, see Installing the web console.

Procedure

1. Log in to the RHEL web console with administrator privileges. For details, see Logging in to the web console.

2. Click **Networking**.

3. Click on the **Edit rules and zones** button.
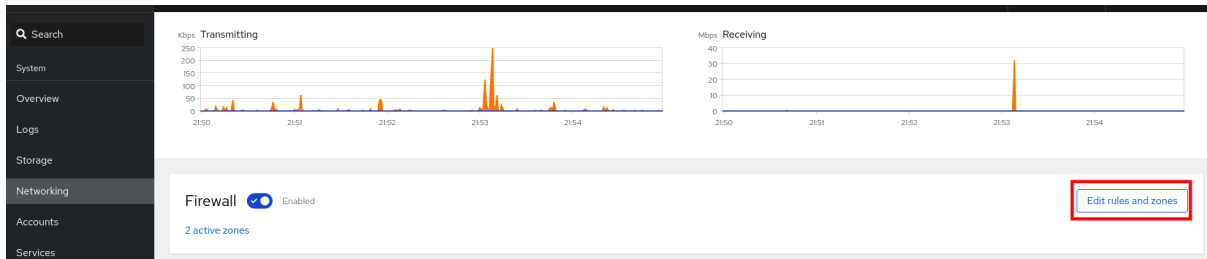
If you do not see the **Edit rules and zones** button, log in to the web console with the administrator privileges.

4. Click on the **Options** icon at the zone you want to remove.
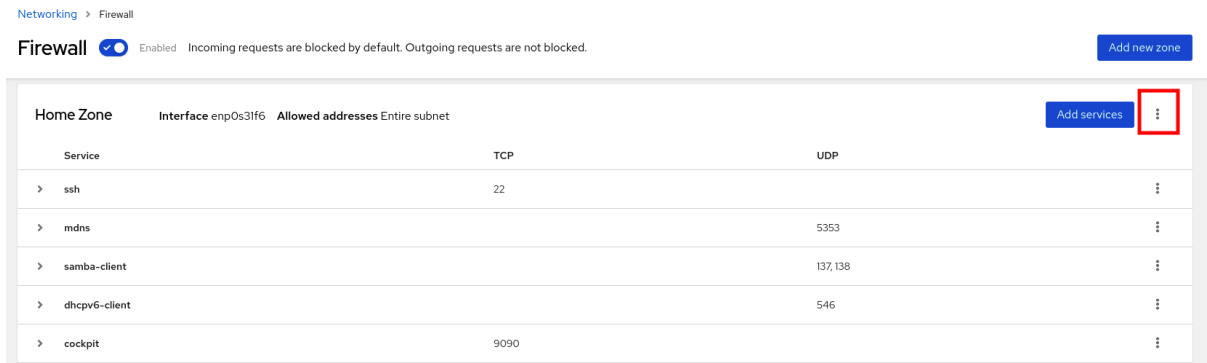
5. Click **Delete**.

The zone is now disabled and the interface does not include opened services and ports which were configured in the zone.

CHAPTER 14. SETTING UP SYSTEM-WIDE CRYPTOGRAPHIC POLICIES IN THE WEB CONSOLE

You can set one of system-wide cryptographic policies and subpolicies directly in the RHEL web console interface. Besides the four predefined system-wide cryptographic policies, you can also apply the following combinations of policies and subpolicies through the graphical interface now:

DEFAULT:SHA1

The **DEFAULT** policy with the **SHA-1** algorithm enabled.

LEGACY:AD-SUPPORT

The **LEGACY** policy with less secure settings that improve interoperability for Active Directory services.

FIPS:OSPP

The **FIPS** policy with further restrictions inspired by the Common Criteria for Information Technology Security Evaluation standard.

Prerequisites

- The RHEL 8 web console has been installed. For details, see Installing and enabling the web console.

- You have **root** privileges or permissions to enter administrative commands with **sudo**.

Procedure

1. Log in to the web console. For more information, see Logging in to the web console .

2. In the **Configuration** card of the **Overview** page, click your current policy value next to **Crypto policy**.
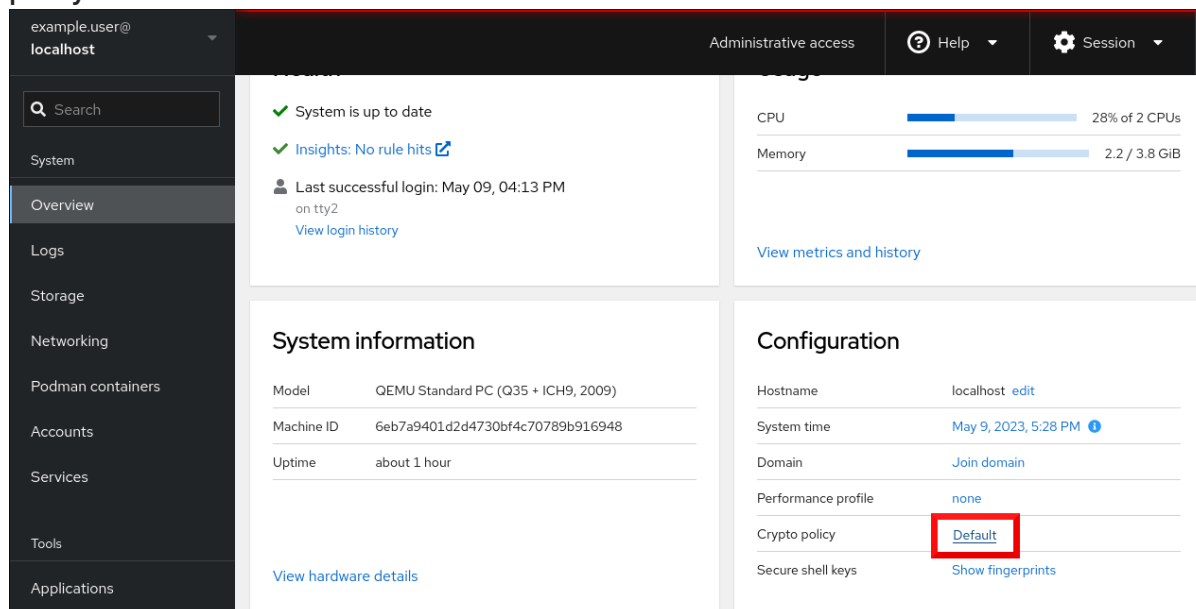
3. In the **Change crypto policy** dialog window, click on the policy you want to start using on your system.
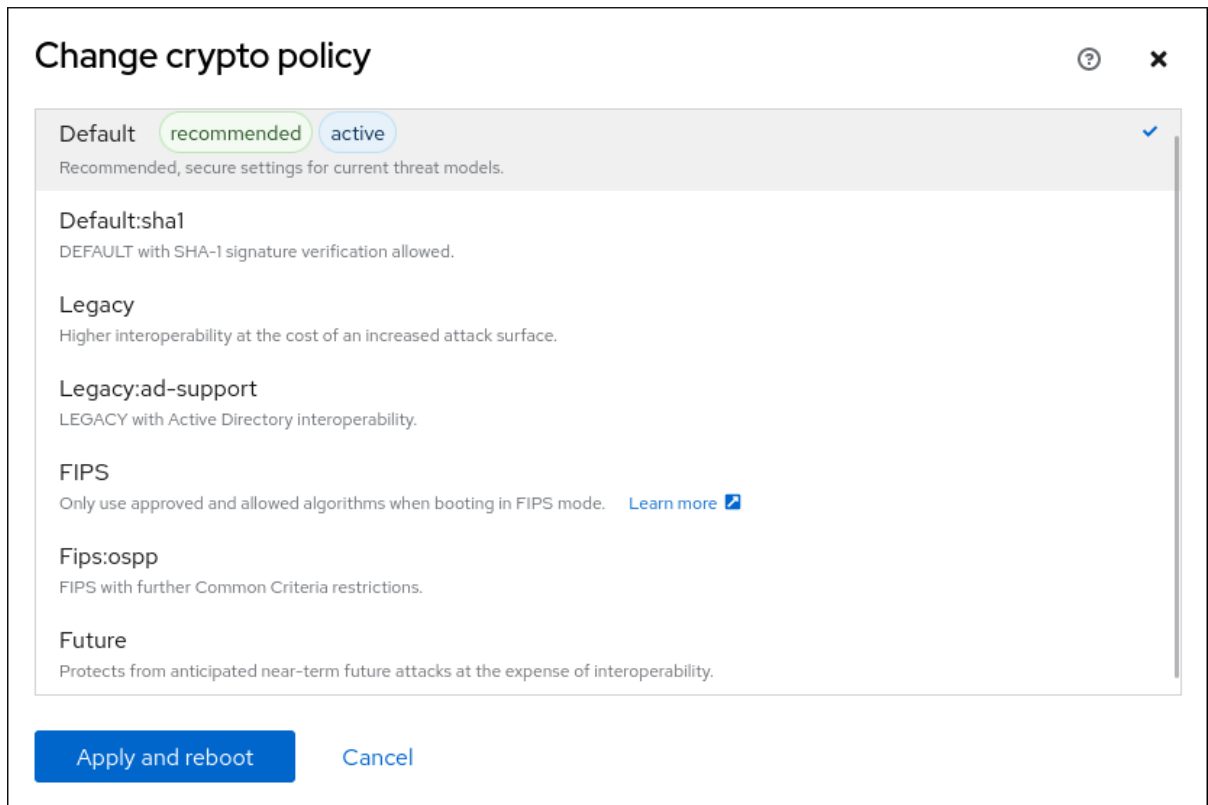
4. Click the **Apply and reboot** button.

Verification

- After the restart, log back in to web console, and check that the **Crypto policy** value corresponds to the one you selected. Alternatively, you can enter the **update-crypto-policies --show** command to display the current system-wide cryptographic policy in your terminal.

Additional resources

- For detailed information about each cryptographic policy, see the System-wide cryptographic policies section in the Security hardening document.

CHAPTER 15. APPLYING A GENERATED ANSIBLE PLAYBOOK

When troubleshooting issues with SELinux, the web console is able to generate a shell script or an Ansible playbook that you can then export and apply for more machines.

Prerequisites

- The web console interface needs to be installed and accessible.
  For details, see Installing the web console .

Procedure

1. Click **SELinux**.

2. Click "View the automation script" on the upper right side.
   A window with the generated script opens. You can navigate between a shell script and an Ansible playbook generation options tab.

   ```
   Automation Script

   Shell Script   Ansible
   - name: Allow virt to sandbox use all caps
     seboolean:
       name: virt_sandbox_use_all_caps
       state: yes
       persistent: yes

   - name: Allow virt to use nfs
     seboolean:
       name: virt_use_nfs
       state: yes
       persistent: yes

   ❷ Create new task file with this content. ⤢Ansible roles documentation

    Copy to clipboard    Close
   ```

3. Click the **Copy to clipboard** button to select the script or playbook and apply it.

As a result, you have an automation script that you can apply to more machines.

Additional resources

- Troubleshootng problems related to SELinux

- Deploying the same SELinux configuration on multiple systems

- For details about the **ansible-playbook** command, see the **ansible-playbook(1)** man page.

CHAPTER 16. MANAGING PARTITIONS USING THE WEB CONSOLE

Learn how to manage file systems on RHEL 8 using the web console.

For details about the available file systems, see the Overview of available file systems.

16.1. DISPLAYING PARTITIONS FORMATTED WITH FILE SYSTEMS IN THE WEB CONSOLE

The **Storage** section in the web console displays all available file systems in the **Filesystems** table.

You can to the list of partitions formatted with file systems displayed in the web console.

Prerequisites

- The **cockpit-storaged** package is installed on your system.

- The web console must be installed and accessible. For details, see Installing the web console.

Procedure

1. Log in to the RHEL 8 web console. For details, see Logging in to the web console .

2. Click on the **Storage** tab.

In the **Filesystems** table, you can see all available partitions formatted with file systems, its name, size and how much space is available on each partition.

Filesystems			
Source ↑	Type ↕	Mount ↕	Size
/dev/nvme0n1p1	ext4	/boot	0.31 / 1.0 GB
/dev/nvme0n1p2	btrfs	/	46 / 250 GB
/dev/sda1	vfat	-	8.09 GB

16.2. CREATING PARTITIONS IN THE WEB CONSOLE

To create a new partition:

- Use an existing partition table

- Create a partition

Prerequisites

- The **cockpit-storaged** package is installed on your system.

- The web console must be installed and accessible. For details, see Installing the web console.

- An unformatted volume connected to the system visible in the **Other Devices** table of the **Storage** tab.

Procedure

1. Log in to the RHEL web console. For details, see Logging in to the web console .

2. Click the **Storage** tab.

3. In the **Other Devices** table, click a volume in which you want to create the partition.

4. In the **Content** section, click the **Create Partition** button.

5. In the **Create partition** dialog box, select the size of the new partition.

6. In the **Erase** drop down menu, select:

   - **Do not overwrite existing data** — the RHEL web console rewrites only the disk header. Advantage of this option is speed of formatting.

   - **Overwrite existing data with zeros** — the RHEL web console rewrites the whole disk with zeros. This option is slower because the program has to go through the whole disk, but it is more secure. Use this option if the disk includes any data and you need to overwrite it.

7. In the **Type** drop down menu, select a file system:

   - **XFS** file system supports large logical volumes, switching physical drives online without outage, and growing an existing file system. Leave this file system selected if you do not have a different strong preference.

   - **ext4** file system supports:

     - Logical volumes

     - Switching physical drives online without outage

     - Growing a file system

     - Shrinking a file system

   Additional option is to enable encryption of partition done by LUKS (Linux Unified Key Setup), which allows you to encrypt the volume with a passphrase.

8. In the **Name** field, enter the logical volume name.

9. In the **Mounting** drop down menu, select **Custom**.
   The **Default** option does not ensure that the file system will be mounted on the next boot.

10. In the **Mount Point** field, add the mount path.

11. Select **Mount at boot**.

12. Click the **Create Partition** button.

Formatting can take several minutes depending on the volume size and which formatting options are selected.

After the formatting has completed successfully, you can see the details of the formatted logical volume on the **Filesystem** tab.

Verification steps

- To verify that the partition has been successfully added, switch to the **Storage** tab and check the **Filesystems** table.

16.3. DELETING PARTITIONS IN THE WEB CONSOLE

The following procedure teaches you how to delete partitions in the web console interface.

Prerequisites

- The **cockpit-storaged** package is installed on your system.

- The web console must be installed and accessible. For details, see Installing the web console.

- Unmount the partition's file system.
  For details about mounting and unmounting partitions, see Mounting and unmounting file systems in the web console.

Procedure

1. Log in to the RHEL web console. For details, see Logging in to the web console .

2. Click on the **Storage** tab.

3. In the **Filesystems** table, select a volume in which you want to delete the partition.

4. In the **Content** section, click on the partition you want to delete.

5. The partition rolls down and you can click on the **Delete** button.
   The partition must not be mounted and used.

Verification steps

- To verify that the partition has been successfully removed, switch to the **Storage** tab and check the **Content** table.

16.4. MOUNTING AND UNMOUNTING FILE SYSTEMS IN THE WEB CONSOLE

To be able to use partitions on RHEL systems, you need to mount a file system on the partition as a device.
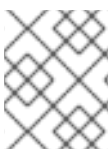
> **NOTE**
>
> You also can unmount a file system and the RHEL system will stop using it. Unmounting the file system enables you to delete, remove, or re-format devices.

Prerequisites

- The **cockpit-storaged** package is installed on your system.

- The web console must be installed and accessible. For details, see Installing the web console .

- If you want to unmount a file system, ensure that the system does not use any file, service, or application stored in the partition.

Procedure

1. Log in to the RHEL web console. For details, see Logging in to the web console .

2. Click on the **Storage** tab.

3. In the **Filesystems** table, select a volume in which you want to delete the partition.

4. In the **Content** section, click on the partition whose file system you want to mount or unmount.

5. Click on the **Mount** or **Unmount** button.
   At this point, the file system has been mounted or unmounted according to your action.

CHAPTER 17. MANAGING NFS MOUNTS IN THE WEB CONSOLE

The RHEL 8 web console enables you to mount remote directories using the Network File System (NFS) protocol.

NFS makes it possible to reach and mount remote directories located on the network and work with the files as if the directory was located on your physical drive.

Prerequisites

- The RHEL 8 web console has been installed.
  For details, see Installing the web console.

- The **cockpit-storaged** package is installed on your system.

- NFS server name or IP address.

- Path to the directory on the remote server.

17.1. CONNECTING NFS MOUNTS IN THE WEB CONSOLE

Connect a remote directory to your file system using NFS.

Prerequisites

- NFS server name or IP address.

- Path to the directory on the remote server.

Procedure

1. Log in to the RHEL 8 web console. For details, see Logging in to the web console .

2. Click **Storage**.

3. Click **+** in the **NFS mounts** section.

NFS mounts	+
No NFS mounts set up	
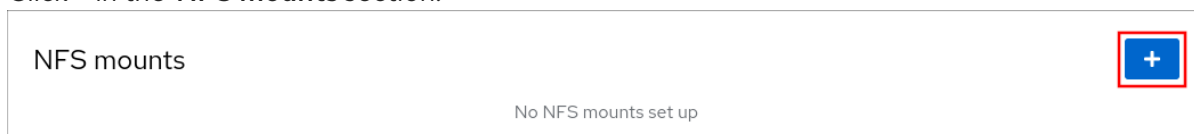
4. In the **New NFS Mount** dialog box, enter the server or IP address of the remote server.

5. In the **Path on Server** field, enter the path to the directory you want to mount.

6. In the **Local Mount Point** field, enter the path where you want to find the directory in your local system.

7. Select **Mount at boot**. This ensures that the directory will be reachable also after the restart of the local system.

8. Optionally, select **Mount read only** if you do not want to change the content.

New NFS mount

Server address	fileserver.example.com
Path on server	/volume1/videotutorials
Local mount point	/mnt/tutorials
Mount options	☑ Mount at boot
	☑ Mount read only
	☐ Custom mount options

Add Cancel

9. Click **Add**.

Verification steps

- Open the mounted directory and verify that the content is accessible.

To troubleshoot the connection, you can adjust it with the Custom Mount Options.

17.2. CUSTOMIZING NFS MOUNT OPTIONS IN THE WEB CONSOLE

Edit an existing NFS mount and add custom mount options.

Custom mount options can help you to troubleshoot the connection or change parameters of the NFS mount such as changing timeout limits or configuring authentication.

Prerequisites

- NFS mount added.

Procedure

1. Log in to the RHEL 8 web console. For details, see Logging in to the web console .

2. Click **Storage**.

3. Click on the NFS mount you want to adjust.

4. If the remote directory is mounted, click **Unmount**.
   The directory must not be mounted during the custom mount options configuration. Otherwise the web console does not save the configuration and this will cause an error.

1. Click **Edit**.

1. In the **NFS Mount** dialog box, select **Custom mount option**.

2. Enter mount options separated by a comma. For example:

   - **nfsvers=4** — the NFS protocol version number

- **soft** — type of recovery after an NFS request times out

- **sec=krb5** — files on the NFS server can be secured by Kerberos authentication. Both the NFS client and server have to support Kerberos authentication.

For a complete list of the NFS mount options, enter **man nfs** in the command line.

1. Click **Apply**.

2. Click **Mount**.

Verification steps

- Open the mounted directory and verify that the content is accessible.

CHAPTER 18. MANAGING REDUNDANT ARRAYS OF INDEPENDENT DISKS IN THE WEB CONSOLE

Redundant Arrays of Independent Disks (RAID) represents a way how to arrange more disks into one storage. RAID protects data stored in the disks against disk failure.

RAID uses the following data distribution strategies:

- Mirroring — data are copied to two different locations. If one disk fails, you have a copy and your data is not lost.

- Striping — data are evenly distributed among disks.

Level of protection depends on the RAID level.

The RHEL web console supports the following RAID levels:

- RAID 0 (Stripe)

- RAID 1 (Mirror)

- RAID 4 (Dedicated parity)

- RAID 5 (Distributed parity)

- RAID 6 (Double Distributed Parity)

- RAID 10 (Stripe of Mirrors)

Before you can use disks in RAID, you need to:

- Create a RAID.

- Format it with file system.

- Mount the RAID to the server.

Prerequisites

- The RHEL 8 web console is installed and accessible. For details, see Installing the web console.

- The **cockpit-storaged** package is installed on your system.

18.1. CREATING RAID IN THE WEB CONSOLE

Configure RAID in the RHEL 8 web console.

Prerequisites

- Physical disks connected to the system. Each RAID level requires different amount of disks.

Procedure

1. Open the RHEL 8 web console.

2. Click **Storage**.

3. Click the menu icon in the **Devices** table.

4. Click **Create RAID device**.

5. In the **Create RAID Device** dialog box, enter a name for a new RAID.

6. In the **RAID Level** drop-down list, select a level of RAID you want to use.

7. In the **Chunk Size** drop-down list, leave the predefined value as it is.
   The **Chunk Size** value specifies how large is each block 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 will be written to the third disk. If you have three disks in your RAID, the fourth 512 KiB will be written to the first disk again.

8. Select disks you want to use for RAID.

9. Click **Create**.

Verification steps

- Go to the **Storage** section and check that you can see the new RAID in the **RAID devices** box and format it.

You have the following options how to format and mount the new RAID in the web console:

- Formatting RAID

- Creating partitions on partition table

- Creating a volume group on top of RAID

18.2. FORMATTING RAID IN THE WEB CONSOLE

Format the new software RAID device created in the RHEL 8 web interface.

Prerequisites

- Physical disks are connected and visible by RHEL 8.

- RAID is created.

- Consider the file system which will be used for the RAID.

- Consider creating of a partitioning table.

Procedure

1. Open the RHEL 8 web console.

2. Click **Storage**.

3. In the **RAID devices** box, choose the RAID you want to format by clicking on it.

4. In the RAID details screen, scroll down to the **Content** part.

5. Click to the newly created RAID.

6. Click the **Format** button.

7. In the **Erase** drop-down list, select:

   - **Don't overwrite existing data**— the RHEL web console rewrites only the disk header. Advantage of this option is speed of formatting.

   - **Overwrite existing data with zeros**— the RHEL web console rewrites the whole disk with zeros. This option is slower because the program has to go through the whole disk. Use this option if the RAID includes any data and you need to rewrite it.

8. In the **Type** drop-down list, select a XFS file system, if you do not have another strong preference.

9. Enter a name of the file system.

10. In the **Mounting** drop down list, select **Custom**.
    The **Default** option does not ensure that the file system will be mounted on the next boot.

11. In the **Mount Point** field, add the mount path.

12. Select **Mount at boot**.

13. Click the **Format** button.
    Formatting can take several minutes depending on the used formatting options and size of RAID.

    After successful finish, you can see the details of the formatted RAID on the **Filesystem** tab.

14. To use the RAID, click **Mount**.

At this point, the system uses the mounted and formatted RAID.

18.3. CREATING A PARTITION TABLE ON RAID USING THE WEB CONSOLE

Format RAID with the partition table on the new software RAID device created in the RHEL 8 interface.

RAID requires formatting as any other storage device. You have two options:

- Format the RAID device without partitions

- Create a partition table with partitions

Prerequisites

- Physical disks are connected and visible by .

- RAID is created.

- Consider the file system used for the RAID.

- Consider creating a partitioning table.

Procedure

Procedure

1. Open the RHEL 8 console.

2. Click **Storage**.

3. In the **RAID devices** box, select the RAID you want to edit.

4. In the RAID details screen, scroll down to the **Content** part.

5. Click to the newly created RAID.

6. Click the **Create partition table** button.

7. In the **Erase** drop-down list, select:

   - **Don't overwrite existing data**— the RHEL web console rewrites only the disk header. Advantage of this option is speed of formatting.

   - **Overwrite existing data with zeros**— the RHEL web console rewrites the whole RAID with zeros. This option is slower because the program has to go through the whole RAID. Use this option if RAID includes any data and you need to rewrite it.

8. In the **Partitioning** drop-down list, select:

   - Compatible with modern system and hard disks > 2TB (GPT) — GUID Partition Table is a modern recommended partitioning system for large RAIDs with more than four partitions.

   - Compatible with all systems and devices (MBR) — Master Boot Record works with disks up to 2 TB in size. MBR also support four primary partitions max.

9. Click **Format**.

At this point, the partitioning table has been created and you can create partitions.

For creating partitions, see Creating partitions on RAID using hte web console .

18.4. CREATING PARTITIONS ON RAID USING THE WEB CONSOLE

Create a partition in the existing partition table.

Prerequisites

   - Partition table is created. For details, see Creating a partition table on RAID using the web console

Procedure

1. Open the RHEL 8 web console.

2. Click **Storage**.

3. In the **RAID devices** box, click to the RAID you want to edit.

4. In the RAID details screen, scroll down to the **Content** part.

5. Click to the newly created RAID.

6. Click **Create Partition**.

7. In the **Create partition** dialog box, set up the size of the first partition.

8. In the **Erase** drop-down list, select:

   - **Don't overwrite existing data**— the RHEL web console rewrites only the disk header. Advantage of this option is speed of formatting.

   - **Overwrite existing data with zeros**— the RHEL web console rewrites the whole RAID with zeros. This option is slower because the program have to go through the whole RAID. Use this option if RAID includes any data and you need to rewrite it.

9. In the **Type** drop-down list, select a XFS file system, if you do not have another strong preference.

10. Enter any name for the file system. Do not use spaces in the name.

11. In the **Mounting** drop down list, select **Custom**.
    The **Default** option does not ensure that the file system will be mounted on the next boot.

12. In the **Mount Point** field, add the mount path.

13. Select **Mount at boot**.

14. Click **Create partition**.

Formatting can take several minutes depending on used formatting options and the size of the RAID.

After a successful finish, you can continue with creating other partitions.

At this point, the system uses mounted and formatted RAID.

18.5. CREATING A VOLUME GROUP ON TOP OF RAID USING THE WEB CONSOLE

Build a volume group from software RAID.

Prerequisites

- RAID device, which is not formatted and mounted.

Procedure

1. Open the RHEL 8 web console.

2. Click **Storage**.

3. Click the **+** button in the **Volume Groups** box.

4. In the **Create Volume Group** dialog box, enter a name for the new volume group.

5. In the **Disks** list, select a RAID device.
   If you do not see the RAID in the list, unmount the RAID from the system. The RAID device must not be used by the RHEL 8 system.

6. Click **Create**.

The new volume group has been created and you can continue with creating a logical volume.

18.6. ADDITIONAL RESOURCES

- To learn more about soft corruption and how you can protect your data when configuring a RAID LV, see Creating a RAID LV with DM integrity .

CHAPTER 19. CONFIGURING LVM LOGICAL VOLUMES USING THE WEB CONSOLE

Red Hat Enterprise Linux 8 supports the LVM logical volume manager. When you install a Red Hat Enterprise Linux 8, it will be installed on LVM automatically created during the installation.
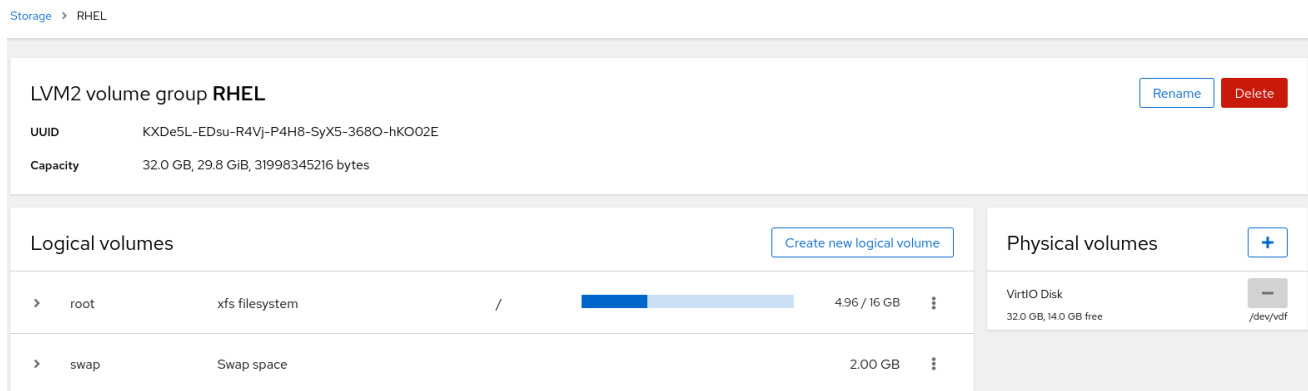
The screenshot shows the web console view of a clean installation of a RHEL 8 system with two logical volumes automatically created during the installation.

To find out more about logical volumes, follow the sections describing:

- What is logical volume manager and when to use it

- What are volume groups and how to create them

- What are logical volumes and how to create them

- How to format logical volumes

- How to resize logical volumes

Prerequisites

- The RHEL 8 web console has been installed.
  For instructions, see Installing and enabling the web console .

- The **cockpit-storaged** package is installed on your system.

- Physical drives, RAID devices, or any other type of block device from which you can create the logical volume.

19.1. LOGICAL VOLUME MANAGER IN THE WEB CONSOLE

The RHEL 8 web console provides a graphical interface to create LVM volume groups and logical volumes.

Volume groups create a layer between physical and logical volumes. It makes you possible to add or remove physical volumes without influencing logical volume itself. Volume groups appear as one drive with capacity consisting of capacities of all physical drives included in the group.

You can join physical drives into volume groups in the web console.

Logical volumes act as a single physical drive and it is built on top of a volume group in your system.

Main advantages of logical volumes are:

- Better flexibility than the partitioning system used on your physical drive.

- Ability to connect more physical drives into one volume.

- Possibility of expanding (growing) or reducing (shrinking) capacity of the volume on-line, without restart.

- Ability to create snapshots.

Additional resources

- Configuring and managing logical volumes

19.2. CREATING VOLUME GROUPS IN THE WEB CONSOLE

Create volume groups from one or more physical drives or other storage devices.

Logical volumes are created from volume groups. Each volume group can include multiple logical volumes.

For details, see Managing LVM volume groups.

Prerequisites

- Physical drives or other types of storage devices from which you want to create volume groups.

Procedure

1. Log in to the RHEL 8 web console.

2. Click **Storage**.

3. In the **Devices** section, select **Create LVM2 volume group** in the drop down menu.
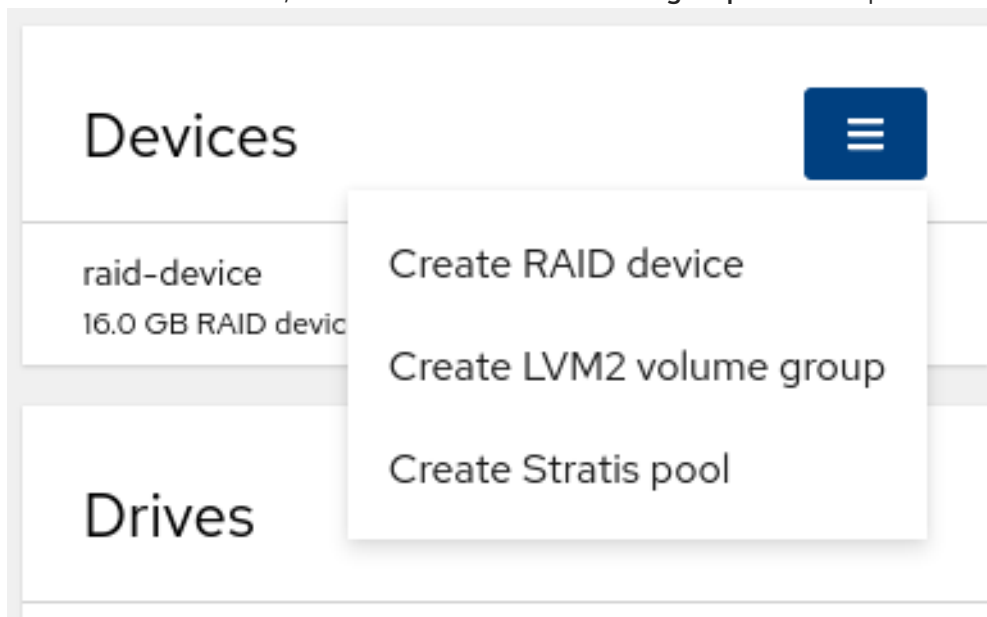
4. In the **Name** field, enter a name of a group without spaces.

5. Select the drives you want to combine to create the volume group.

Create volume group

Name	rhel-volume-group	
Disks	☑ 16.0 GB RAID device raid-device	/dev/md/raid-device
	☑ 16.0 GB VirtIO Disk	/dev/vdb

Create   Cancel

It might happen that you cannot see devices as you expected. The RHEL web console displays only unused block devices. Used devices means, for example:

- Devices formatted with a file system

- Physical volumes in another volume group

- Physical volumes being a member of another software RAID device
  If you do not see the device, format it to be empty and unused.

6. Click **Create**.

The web console adds the volume group in the **Devices** section. After clicking the group, you can create logical volumes that are allocated from that volume group.

Devices

≡

raid-device
16.0 GB RAID device   /dev/md/raid-device

rhel-volume-group
32.0 GB LVM2 volume group   /dev/rhel-volume-group/

19.3. CREATING LOGICAL VOLUMES IN THE WEB CONSOLE

Logical volumes act as physical drives. You can use the RHEL 8 web console to create LVM logical volumes in a volume group.

Prerequisites

- The **cockpit-storaged** package is installed on your system.

- Volume group created. For details, see Creating volume groups in the web console .

Procedure

1. Log in to the RHEL 8 web console.

2. Click **Storage**.

3. In the **Devices** section, click the volume group in which you want to create logical volumes.

4. In the **Logical volumes** section, click **Create new Logical Volume**.

5. In the **Name** field, enter a name for the new logical volume without spaces.

6. In the **Purpose** drop down menu, select **Block device for filesystems**.
   This configuration enables you to create a logical volume with the maximum volume size which is equal to the sum of the capacities of all drives included in the volume group.

7. Define the size of the logical volume. Consider:

   - How much space the system using this logical volume will need.

   - How many logical volumes you want to create.

   You do not have to use the whole space. If necessary, you can grow the logical volume later.

8. Click **Create**.

To verify the settings, click your logical volume and check the details.

Logical volumes   Create new logical volume

⌄   rhel-logical-volume   Unrecognized data   16.0 GB   Format   ⋮

Volume   Unrecognized data

Name   rhel-logical-volume   edit

Size   16.0 GB   Shrink   Grow

At this stage, the logical volume has been created and you need to create and mount a file system with the formatting process.

19.4. FORMATTING LOGICAL VOLUMES IN THE WEB CONSOLE

Logical volumes act as physical drives. To use them, you must format them with a file system.

> **WARNING**
>
> Formatting logical volumes erases all data on the volume.

The file system you select determines the configuration parameters you can use for logical volumes. For example, the XFS file system does not support shrinking volumes. For details, see Resizing logical volumes in the web console.

Prerequisites

- The **cockpit-storaged** package is installed on your system.

- Logical volume created. For details, see Creating logical volumes in the web console .

- You have root access privileges to the system.

Procedure

1. Log in to the RHEL 8 web console.

2. Click **Storage**.

3. In the **Devices** section, click the volume group in which the logical volume is placed.

4. In the **Logical volumes** section, click **Format**.

Logical volumes

Create new logical volume

| ⌄ | rhel-logical-volume | Unrecognized data | 16.0 GB | Format | ⋮ |

Volume | Unrecognized data

| Name | rhel-logical-volume | edit |
| Size | 16.0 GB | Shrink | Grow |

5. In the **Name** field, enter a name for the file system.

6. In the **Mount Point** field, add the mount path.

⚠ **Format /dev/rhel-volume-group/rhel-logical-volume**

Name	rhel-fs
Mount point	/media
Type	XFS (recommended) ▼
Overwrite	☐ Overwrite existing data with zeros (slower)
Encryption	No encryption ▼
At boot	Mount without waiting, ignore failure ▼
	▬ Mounts in parallel with services
	ℹ Boot still succeeds when filesystem does not mount
Mount options	☐ Mount read only
	☐ Custom mount options

Formatting erases all data on a storage device.

[Format and mount] [Format only] Cancel

7. In the **Type** drop-down menu, select a file system:

- **XFS** file system supports large logical volumes, switching physical drives online without outage, and growing an existing file system. Leave this file system selected if you do not have a different strong preference.
  XFS does not support reducing the size of a volume formatted with an XFS file system

- **ext4** file system supports:

  - Logical volumes

  - Switching physical drives online without an outage

  - Growing a file system

  - Shrinking a file system

8. Based on how you want to format the volume, select the **Overwrite** option:

- **Do not overwrite existing data** — the RHEL web console rewrites only the disk header. The advantage of this option is the short time required for the formatting process.

- **Overwrite existing data with zeros** — the RHEL web console rewrites the whole disk with zeros. This option is slower because the program has to go through the whole disk. Use this option if the disk includes any data and you need to overwrite it.

9. In the **Encryption** drop-down menu, select the type of encryption if you want to enable it on the logical volume.
   You can select a version with either the LUKS1 (Linux Unified Key Setup) or LUKS2 encryption, which allows you to encrypt the volume with a passphrase.

10. In the **At boot** drop-down menu, select when you want the logical volume to mount after the system boots.

11. Select the required **Mount options**.

12. Format the logical volume:

    - If you want to format the volume and immediately mount it, click **Format and mount**.

    - If you want to format the volume without mounting it, click **Format only**.
      Formatting can take several minutes depending on the volume size and which formatting options are selected.

Verification

1. After the formatting completes successfully, you can see the details of the formatted logical volume on the **Filesystem** tab.

Logical volumes				Create new logical volume
⌄ rhel-logical-volume	xfs filesystem	/media	16.0 GB	Mount ⋮

Volume	**Filesystem**

Name	rhel-fs edit
Mount point	/media edit
	The filesystem is not mounted.

2. If you had selected the **Format only** option, click **Mount** to use the logical volume.

19.5. RESIZING LOGICAL VOLUMES IN THE WEB CONSOLE

Learn how to extend or reduce logical volumes in the RHEL 8 web console.

Whether you can resize a logical volume depends on which file system you are using. Most file systems enable you to extend (grow) the volume online (without outage).

You can also reduce (shrink) the size of logical volumes, if the logical volume contains a file system which supports shrinking. It should be available, for example, in the ext3/ext4 file systems.

> **⚠ WARNING**
>
> You cannot reduce volumes that contains GFS2 or XFS filesystem.

Prerequisites

- Existing logical volume containing a file system that supports resizing logical volumes.

Procedure

The following steps provide the procedure for growing a logical volume without taking the volume offline:

1. Log in to the RHEL web console.

2. Click **Storage**.

3. In the **Devices** section, click the volume group in which the logical volume is placed.

4. In the **Logical volumes** section, click the logical volume.

5. On the **Volume** tab, click **Grow**.

6. In the **Grow logical volume** dialog box, adjust volume size.

7. Click **Grow**.

LVM grows the logical volume without system outage.

19.6. ADDITIONAL RESOURCES

- [Configuring and managing logical volumes](#)

CHAPTER 20. CONFIGURING THIN LOGICAL VOLUMES USING THE WEB CONSOLE

You can use thin-provisioned logical volumes to allocate more space for designated applications or servers than the actually available physical storage.

For details, see Creating thin-provisioned snapshot volumes .

The following sections describe:

- Creating pools for the thin-provisioned logical volumes.

- Creating thin logical volumes.

- Formatting thin logical volumes.

- Creating snapshots of thin logical volumes.

Prerequisites

- The RHEL 8 web console has been installed.
  For details, see Installing the web console .

- The **cockpit-storaged** package is installed on your system.

- Physical drives or other types of storage devices, which you want to use to create volume groups, are attached to your system.

20.1. CREATING POOLS FOR THIN LOGICAL VOLUMES IN THE WEB CONSOLE

Create a pool for thin-provisioned volumes.

Prerequisites

- Volume group created .

Procedure

1. Log in to the RHEL 8 web console.

2. Click **Storage**.

3. Click the volume group in which you want to create thin volumes.

4. Click **Create new Logical Volume**

5. In the **Name** field, enter a name for the new pool of thin volumes without spaces.

6. In the **Purpose** drop down menu, select **Pool for thin-provisioned volumes** This configuration enables you to create the thin volume.

7. Define the size of the pool of thin volumes. Consider:

   - How many thin volumes you will need in this pool?

- What is the expected size of each thin volume?

  You do not have to use the whole space. If necessary, you can grow the pool later.

8. Click **Create**.
   The pool for thin volumes has been created and you can add thin volumes.

20.2. CREATING THIN LOGICAL VOLUMES IN THE WEB CONSOLE

Create a thin logical volume in the pool. The pool can include multiple thin volumes and each thin volume can be as large as the pool for thin volumes itself.

> **IMPORTANT**
>
> Using thin volumes requires regular checkup of actual free physical space of the logical volume.

Prerequisites

- Pool for thin volumes created.
  For details, see Creating pools for thin logical volumes in the web console .

Procedure

1. Log in to the RHEL 8 web console.

2. Click **Storage**.

3. Click the volume group in which you want to create thin volumes.

4. Click the desired pool.

5. Click **Create Thin Volume**.

6. In the **Create Thin Volume** dialog box, enter a name for the thin volume without spaces.

7. Define the size of the thin volume.

8. Click **Create**.

At this stage, the thin logical volume has been created and you need to format it.

20.3. FORMATTING LOGICAL VOLUMES IN THE WEB CONSOLE

Logical volumes act as physical drives. To use them, you must format them with a file system.

> **WARNING**
>
> Formatting logical volumes erases all data on the volume.

The file system you select determines the configuration parameters you can use for logical volumes. For example, the XFS file system does not support shrinking volumes. For details, see Resizing logical volumes in the web console.

Prerequisites

- The **cockpit-storaged** package is installed on your system.

- Logical volume created. For details, see Creating logical volumes in the web console .

- You have root access privileges to the system.

Procedure

1. Log in to the RHEL 8 web console.

2. Click **Storage**.

3. In the **Devices** section, click the volume group in which the logical volume is placed.

4. In the **Logical volumes** section, click **Format**.

5. In the **Name** field, enter a name for the file system.

6. In the **Mount Point** field, add the mount path.

7. In the **Type** drop-down menu, select a file system:

- **XFS** file system supports large logical volumes, switching physical drives online without outage, and growing an existing file system. Leave this file system selected if you do not have a different strong preference.
  XFS does not support reducing the size of a volume formatted with an XFS file system

- **ext4** file system supports:

  - Logical volumes

  - Switching physical drives online without an outage

  - Growing a file system

  - Shrinking a file system

8. Based on how you want to format the volume, select the **Overwrite** option:

- **Do not overwrite existing data**— the RHEL web console rewrites only the disk header. The advantage of this option is the short time required for the formatting process.

- **Overwrite existing data with zeros**— the RHEL web console rewrites the whole disk with zeros. This option is slower because the program has to go through the whole disk. Use this option if the disk includes any data and you need to overwrite it.

9. In the **Encryption** drop-down menu, select the type of encryption if you want to enable it on the logical volume.
   You can select a version with either the LUKS1 (Linux Unified Key Setup) or LUKS2 encryption, which allows you to encrypt the volume with a passphrase.

10. In the **At boot** drop-down menu, select when you want the logical volume to mount after the system boots.

11. Select the required **Mount options**.

12. Format the logical volume:

    - If you want to format the volume and immediately mount it, click **Format and mount**.

    - If you want to format the volume without mounting it, click **Format only**.
      Formatting can take several minutes depending on the volume size and which formatting options are selected.

Verification

1. After the formatting completes successfully, you can see the details of the formatted logical volume on the **Filesystem** tab.

2. If you had selected the **Format only** option, click **Mount** to use the logical volume.

20.4. USING THE WEB CONSOLE TO CREATE THIN-PROVISIONED SNAPSHOT VOLUMES

You can create snapshots of thin logical volumes in the RHEL web console to backup changes recorded on the disk from the last snapshot.

Prerequisites

- The web console is installed and accessible. For more information, see Installing and enabling the web console.

- The **cockpit-storaged** package is installed on your system.

- A thin-provisioned volume is created. For more information see Configuring thin logical volumes using the web console.

Procedure

1. Log in to the RHEL web console. For more information, see Logging in to the web console .

2. In the RHEL web console, click **Storage**.

3. In the **Devices** box, click the volume group that contains the required logical volume.

4. Find the volume in the **Logical volumes** table and click the menu button ⋮ on its line.

5. Click **Create snapshot**.

6. In the **Name** field, enter a snapshot name.

7. Click **Create**.

8. Find the new snapshot in the **Logical volumes** table and click **Activate** to activate the volume.

CHAPTER 21. CHANGING PHYSICAL DRIVES IN VOLUME GROUPS USING THE WEB CONSOLE

Change the drive in a volume group using the RHEL 8 web console.

The change of physical drives consists of the following procedures:

- Adding physical drives from logical volumes.

- Removing physical drives from logical volumes.

Prerequisites

- The RHEL 8 web console has been installed.
  For details, see Installing the web console.

- The **cockpit-storaged** package is installed on your system.

- A new physical drive for replacing the old or broken one.

- The configuration expects that physical drives are organized in a volume group.

21.1. ADDING PHYSICAL DRIVES TO VOLUME GROUPS IN THE WEB CONSOLE

The RHEL 8 web console enables you to add a new physical drive or other type of volume to the existing logical volume.

Prerequisites

- A volume group must be created.

- A new drive connected to the machine.

Procedure

1. Log in to the RHEL 8 console.

2. Click **Storage**.

3. In the **Volume Groups** box, click the volume group in which you want to add a physical volume.

4. In the **Physical Volumes** box, click the **+** button.

5. In the **Add Disks** dialog box, select the preferred drive and click **Add**.

As a result, the RHEL 8 web console adds the physical volume.

Verification steps

- Check the **Physical Volumes** for section, and the logical volume can immediately start to write on the drive.

21.2. REMOVING PHYSICAL DRIVES FROM VOLUME GROUPS IN THE WEB CONSOLE

If a logical volume includes multiple physical drives, you can remove one of the physical drives online.

The system moves automatically all data from the drive to be removed to other drives during the removal process. Notice that it can take some time.

The web console also verifies, if there is enough space for removing the physical drive.

Prerequisites

- A volume group with more than one physical drive connected.

Procedure

The following steps describe how to remove a drive from the volume group without causing outage in the RHEL 8 web console.

1. Log in to the RHEL 8 web console.

2. Click **Storage**.

3. Click the volume group in which you have the logical volume.

4. In the **Physical Volumes** section, locate the preferred volume.

5. Click the **-** button.
   The RHEL 8 web console verifies whether the logical volume has enough free space for removing the disk. If not, you cannot remove the disk and it is necessary to add another disk first. For details, see Adding physical drives to logical volumes in the web console .

As results, the RHEL 8 web console removes the physical volume from the created logical volume without causing an outage.

CHAPTER 22. MANAGING VIRTUAL DATA OPTIMIZER VOLUMES USING THE WEB CONSOLE

Configure the Virtual Data Optimizer (VDO) using the RHEL 8 web console.

You will learn how to:

- Create VDO volumes

- Format VDO volumes

- Extend VDO volumes

Prerequisites

- The RHEL 8 web console is installed and accessible. For details, see Installing the web console.

- The **cockpit-storaged** package is installed on your system.

22.1. VDO VOLUMES IN THE WEB CONSOLE

Red Hat Enterprise Linux 8 supports Virtual Data Optimizer (VDO).

VDO is a block virtualization technology that combines:

Compression

For details, see Enabling or disabling compression in VDO.

Deduplication

For details, see Enabling or disabling compression in VDO.

Thin provisioning

For details, see Creating and managing thin provisioned volumes (thin volumes).

Using these technologies, VDO:

- Saves storage space inline

- Compresses files

- Eliminates duplications

- Enables you to allocate more virtual space than how much the physical or logical storage provides

- Enables you to extend the virtual storage by growing

VDO can be created on top of many types of storage. In the RHEL 8 web console, you can configure VDO on top of:

- LVM

> **NOTE**
>
> It is not possible to configure VDO on top of thinly-provisioned volumes.

- Physical volume

- Software RAID

For details about placement of VDO in the Storage Stack, see System Requirements.

Additional resources

- For details about VDO, see Deduplicating and compressing storage.

22.2. CREATING VDO VOLUMES IN THE WEB CONSOLE

Create a VDO volume in the RHEL web console.

Prerequisites

- Physical drives, LVMs, or RAID from which you want to create VDO.

Procedure

1. Log in to the RHEL 8 web console.
   For details, see Logging in to the web console .

2. Click **Storage**.

3. Click the **+** button in the **VDO Devices** box.

4. In the **Name** field, enter a name of a VDO volume without spaces.

5. Select the drive that you want to use.

6. In the **Logical Size** bar, set up the size of the VDO volume. You can extend it more than ten times, but consider for what purpose you are creating the VDO volume:

   - For active VMs or container storage, use logical size that is ten times the physical size of the volume.

   - For object storage, use logical size that is three times the physical size of the volume.

   For details, see Deploying VDO.

7. In the **Index Memory** bar, allocate memory for the VDO volume.
   For details about VDO system requirements, see System Requirements.

8. Select the **Compression** option. This option can efficiently reduce various file formats.
   For details, see Enabling or disabling compression in VDO .

9. Select the **Deduplication** option.
   This option reduces the consumption of storage resources by eliminating multiple copies of duplicate blocks. For details, see Enabling or disabling compression in VDO .

10. [Optional] If you want to use the VDO volume with applications that need a 512 bytes block size, select **Use 512 Byte emulation**. This reduces the performance of the VDO volume, but should be very rarely needed. If in doubt, leave it off.

11. Click **Create**.

Verification steps

- Check that you can see the new VDO volume in the **Storage** section. Then you can format it with a file system.

22.3. FORMATTING VDO VOLUMES IN THE WEB CONSOLE

VDO volumes act as physical drives. To use them, you need to format them with a file system.

> ⚠️ **WARNING**
>
> Formatting VDO will erase all data on the volume.

The following steps describe the procedure to format VDO volumes.

Prerequisites

- A VDO volume is created. For details, see Creating VDO volumes in the web console .

Procedure

1. Log in to the RHEL 8 web console. For details, see Logging in to the web console .

2. Click **Storage**.

3. Click the VDO volume.

4. Click on the **Unrecognized Data** tab.

5. Click **Format**.

6. In the **Erase** drop down menu, select:

   Do not overwrite existing data

   The RHEL web console rewrites only the disk header. The advantage of this option is the speed of formatting.

   Overwrite existing data with zeros

   The RHEL web console rewrites the whole disk with zeros. This option is slower because the program has to go through the whole disk. Use this option if the disk includes any data and you need to rewrite them.

7. In the **Type** drop down menu, select a filesystem:

   - The **XFS** file system supports large logical volumes, switching physical drives online without outage, and growing. Leave this file system selected if you do not have a different strong preference.
     XFS does not support shrinking volumes. Therefore, you will not be able to reduce volume formatted with XFS.

- The **ext4** file system supports logical volumes, switching physical drives online without outage, growing, and shrinking.

  You can also select a version with the LUKS (Linux Unified Key Setup) encryption, which allows you to encrypt the volume with a passphrase.

8. In the **Name** field, enter the logical volume name.

9. In the **Mounting** drop down menu, select **Custom**.
   The **Default** option does not ensure that the file system will be mounted on the next boot.

10. In the **Mount Point** field, add the mount path.

11. Select **Mount at boot**.

12. Click **Format**.
    Formatting can take several minutes depending on the used formatting options and the volume size.

    After a successful finish, you can see the details of the formatted VDO volume on the **Filesystem** tab.

13. To use the VDO volume, click **Mount**.

At this point, the system uses the mounted and formatted VDO volume.

22.4. EXTENDING VDO VOLUMES IN THE WEB CONSOLE

Extend VDO volumes in the RHEL 8 web console.

Prerequisites

- The **cockpit-storaged** package is installed on your system.

- The VDO volume created.

Procedure

1. Log in to the RHEL 8 web console.
   For details, see Logging in to the web console .

2. Click **Storage**.

3. Click your VDO volume in the **VDO Devices** box.

4. In the VDO volume details, click the **Grow** button.

5. In the **Grow logical size of VDO** dialog box, extend the logical size of the VDO volume.

1. Click **Grow**.

Verification steps

- Check the VDO volume details for the new size to verify that your changes have been successful.

CHAPTER 23. LOCKING DATA WITH LUKS PASSWORD IN THE RHEL WEB CONSOLE

In the web console's **Storage** tab, you can now create, lock, unlock, resize, and otherwise configure encrypted devices using the LUKS (Linux Unified Key Setup) version 2 format.

This new version of LUKS offers:

- More flexible unlocking policies

- Stronger cryptography

- Better compatibility with future changes

Prerequisites

- The RHEL 8 web console has been installed. For details, see Installing the web console.

- The **cockpit-storaged** package is installed on your system.

23.1. LUKS DISK ENCRYPTION

Linux Unified Key Setup-on-disk-format (LUKS) provides a set of tools that simplifies managing the encrypted devices. With LUKS, you can encrypt block devices and enable multiple user keys to decrypt a master key. For bulk encryption of the partition, use this master key.

Red Hat Enterprise Linux uses LUKS to perform block device encryption. By default, the option to encrypt the block device is unchecked during the installation. If you select the option to encrypt your disk, the system prompts you for a passphrase every time you boot the computer. This passphrase unlocks the bulk encryption key that decrypts your partition. If you want to modify the default partition table, you can select the partitions that you want to encrypt. This is set in the partition table settings.

Ciphers

The default cipher used for LUKS is **aes-xts-plain64**. The default key size for LUKS is 512 bits. The default key size for LUKS with **Anaconda** XTS mode is 512 bits. The following are the available ciphers:

- Advanced Encryption Standard (AES)

- Twofish

- Serpent

Operations performed by LUKS

- LUKS encrypts entire block devices and is therefore well-suited for protecting contents of mobile devices such as removable storage media or laptop disk drives.

- The underlying contents of the encrypted block device are arbitrary, which makes it useful for encrypting swap devices. This can also be useful with certain databases that use specially formatted block devices for data storage.

- LUKS uses the existing device mapper kernel subsystem.

- LUKS provides passphrase strengthening, which protects against dictionary attacks.

- LUKS devices contain multiple key slots, which means you can add backup keys or passphrases.

IMPORTANT

LUKS is not recommended for the following scenarios:

- Disk-encryption solutions such as LUKS protect the data only when your system is off. After the system is on and LUKS has decrypted the disk, the files on that disk are available to anyone who have access to them.

- Scenarios that require multiple users to have distinct access keys to the same device. The LUKS1 format provides eight key slots and LUKS2 provides up to 32 key slots.

- Applications that require file-level encryption.

Additional resources

- LUKS Project Home Page

- LUKS On-Disk Format Specification

- FIPS 197: Advanced Encryption Standard (AES)

23.2. CONFIGURING THE LUKS PASSPHRASE IN THE WEB CONSOLE

If you want to add encryption to an existing logical volume on your system, you can only do so through formatting the volume.

Prerequisites

- The web console must be installed and accessible. For details, see Installing the web console.

- The **cockpit-storaged** package is installed on your system.

- Available existing logical volume without encryption.

Procedure

1. Log in to the RHEL 8 web console.
   For details, see Logging in to the web console .

2. Click **Storage**.

3. Select the storage device you want to format.

4. Click the menu icon and select **Format** option.

5. Select the **Encrypt data** box to activate encryption on your storage device.

6. Set and confirm your new passphrase.

7. [Optional] Modify further encryption options.

8. Finalize formatting settings.

9. Click **Format**.

23.3. CHANGING THE LUKS PASSPHRASE IN THE WEB CONSOLE

Change a LUKS passphrase on an encrypted disk or partition in the web console.

Prerequisites

- The web console must be installed and accessible. For details, see Installing the web console.

- The **cockpit-storaged** package is installed on your system.

Procedure

1. Log in to the web console. For details, see Logging in to the web console .

2. Click **Storage**

3. In the Drives table, select the disk with encrypted data.

4. In **Content**, select the encrypted partition.

5. Click **Encryption**.

6. In the **Keys** table, click the pen icon.

7. In the **Change passphrase** dialog window:

   a. Enter your current passphrase.

   b. Enter your new passphrase.

   c. Confirm your new passphrase.

8. Click **Save**

CHAPTER 24. CONFIGURING AUTOMATED UNLOCKING BY USING A TANG KEY IN THE WEB CONSOLE

You can configure automated unlocking of a LUKS-encrypted storage device using a key provided by a Tang server.

Prerequisites

- The RHEL 8 web console has been installed. See Installing the web console for details.

- The **cockpit-storaged** and **clevis-luks** packages are installed on your system.

- The **cockpit.socket** service is running at port 9090.

- A Tang server is available. See Deploying a Tang server with SELinux in enforcing mode for details.

Procedure

1. Open the RHEL web console by entering the following address in a web browser:

   > https://*<localhost>*:9090

   Replace the *<localhost>* part by the remote server's host name or IP address when you connect to a remote system.

2. Provide your credentials and click **Storage**. In the **Filesystems** section, click the disk that contains an encrypted volume you plan to add to unlock automatically.

3. In the following window listing partitions and drive details of the selected disk, click **>** next to the encrypted file system to expand details of the encrypted volume you want to unlock using the Tang server, and click **Encryption**.

4. Click **+** in the **Keys** section to add a Tang key:
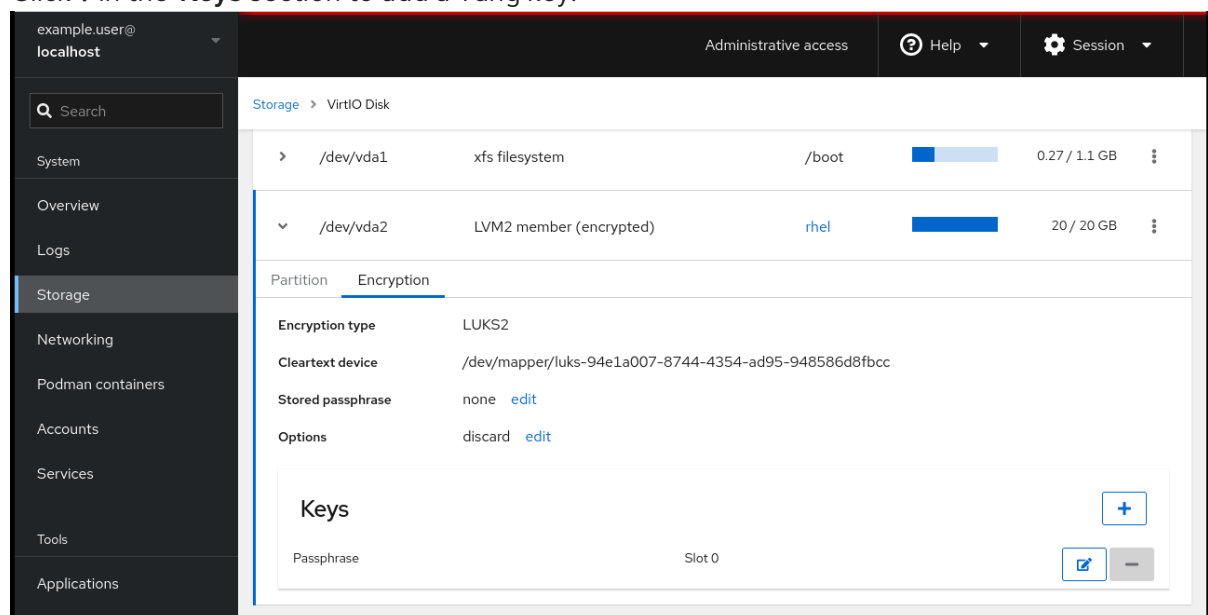
5. Select **Tang keyserver** as **Key source**, provide the address of your Tang server, and a password that unlocks the LUKS-encrypted device. Click **Add** to confirm:

The following dialog window provides a command to verify that the key hash matches.

6. In a terminal on the Tang server, use the **tang-show-keys** command to display the key hash for comparison. In this example, the Tang server is running on the port *7500*:

```
# tang-show-keys 7500
fM-EwYeiTxS66X3s1UAywsGKGnxnpll8ig0KOQmr9CM
```

7. Click **Trust key** when the key hashes in the web console and in the output of previously listed commands are the same:

8. In RHEL 8.8 and later, after you select an encrypted root file system and a Tang server, you can skip adding the **rd.neednet=1** parameter to the kernel command line, installing the **clevis-dracut** package, and regenerating an initial RAM disk (**initrd**). For non-root file systems, the web console now enables the **remote-cryptsetup.target** and **clevis-luks-akspass.path systemd** units, installs the **clevis-systemd** package, and adds the **_netdev** parameter to the **fstab** and **crypttab** configuration files.

Verification

1. Check that the newly added Tang key is now listed in the **Keys** section with the **Keyserver** type:

2. Verify that the bindings are available for the early boot, for example:

```
# lsinitrd | grep clevis
clevis
clevis-pin-null
clevis-pin-sss
clevis-pin-tang
clevis-pin-tpm2
lrwxrwxrwx  1 root    root       48 Feb 14 17:45
etc/systemd/system/cryptsetup.target.wants/clevis-luks-askpass.path…
…
```

Additional resources

- [Configuring automated unlocking of encrypted volumes using policy–based decryption](#)

CHAPTER 25. MANAGING SOFTWARE UPDATES IN THE WEB CONSOLE

Learn how to manage software updates in the RHEL 8 web console and ways to automate them.

The Software Updates module in the web console is based on the **yum** utility. For more information about updating software with **yum**, see the Updating software packages section.

25.1. MANAGING MANUAL SOFTWARE UPDATES IN THE WEB CONSOLE

You can manually update your software using the web console.

Prerequisites

- The web console must be installed and accessible. For details, see Installing the web console .

Procedure

1. Log in to the RHEL 8 web console.
   For details, see Logging in to the web console .

2. Click **Software Updates**.
   The list of available updates refreshes automatically if the last check happened more than 24 hours ago. To trigger a refresh, click the **Check for Updates** button.

3. Apply updates. You can watch the update log while the update is running.

   a. To install all available updates, click the **Install all updates** button.

   b. If you have security updates available, you can install them separately by clicking the **Install Security Updates** button.

   c. If you have kpatch updates available, you can install them separately by clicking the **Install kpatch updates** button.

4. Optional: You can turn on the **Reboot after completion** switch for an automatic restart of your system.
   If you perform this step, you can skip the remaining steps of this procedure.

5. After the system applies updates, you get a recommendation to restart your system.
   We recommend this especially if the update included a new kernel or system services that you do not want to restart individually.

6. Click **Ignore** to cancel the restart, or **Restart Now** to proceed with restarting your system.
   After the system restart, log in to the web console and go to the **Software Updates** page to verify that the update has been successful.

25.2. MANAGING AUTOMATIC SOFTWARE UPDATES IN THE WEB CONSOLE

In the web console, you can choose to apply all updates, or security updates and also manage periodicity and time of your automatic updates.

Prerequisites

- The web console must be installed and accessible. For details, see Installing the web console .

Procedure

1. Log in to RHEL 8 web console. For details, see Logging in to the web console .

2. Click **Software Updates**.

3. In the **Settings** table, click the **Edit** button.

4. Pick one of the types of automatic updates. You can select from **Security updates only**, or **All updates**.

5. To modify the day of the automatic update, click on the **every day** drop-down menu and select a specific day.

6. To modify the time of the automatic update, click into the **6:00** field and select or type a specific time.

7. If you want to disable automatic software updates, select the **No updates** type.

25.3. MANAGING ON-DEMAND RESTARTING AFTER APPLYING SOFTWARE UPDATES IN THE WEB CONSOLE

The intelligent restarting feature informs the users whether it is necessary to reboot the whole system after you apply a software update or if it is sufficient to only restart certain services.

Prerequisites

- The web console must be installed and accessible. For details, see Installing the web console .

Procedure

1. Log in to the RHEL 8 web console. For details, see Logging in to the web console .

2. Click **Software Updates**.

3. Apply an update of your system.

4. After a successful update, click **Reboot system...**, **Restart services...**, or **Ignore**

5. If you decide to ignore, you can return to the restart or reboot menu by doing one of the following:

    a. Rebooting:

        i. Click the **Reboot system** button in the **Status** field of the **Software Updates** page.

        ii. (Optional) Write a message to the logged in users.

        iii. Select a delay from the **Delay** drop down menu.

        iv. Click **Reboot**.

b. Restarting services:

   i. Click the **Restart services…** button in the **Status** field of the **Software Updates** page. You will see a list of all the services that require a restart.

   ii. Click **Restart services**.
   Depending on your choice, the system will reboot or your services will restart.

25.4. APPLYING PATCHES WITH KERNEL LIVE PATCHING IN THE WEB CONSOLE

The web console allows users to apply kernel security patches without forcing reboots by using the **kpatch** framework. The following procedure shows how to set up the preferred type of patching.

Prerequisites

- The web console must be installed and accessible. For details, see Installing the web console.

Procedure

1. Log in to the web console with administrative privileges. For details, see Logging in to the web console.

2. Click **Software Updates**.

3. Check the status of your kernel patching settings.

   a. If the patching is not installed, click **Install**.
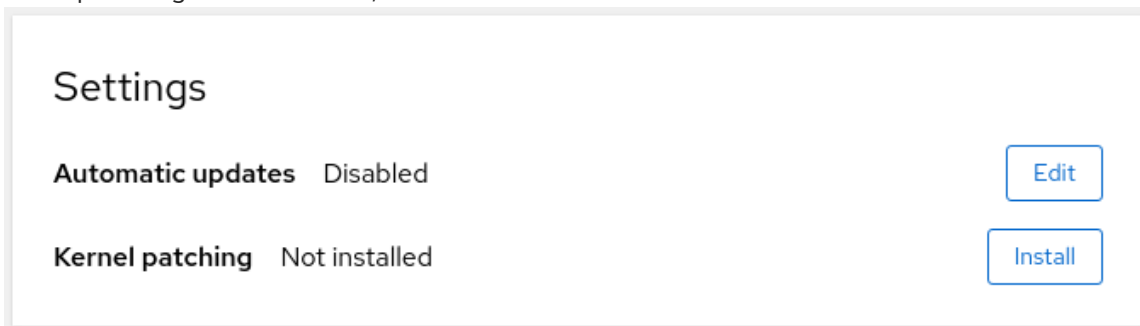
   b. To enable kernel patching, click **Enable**.
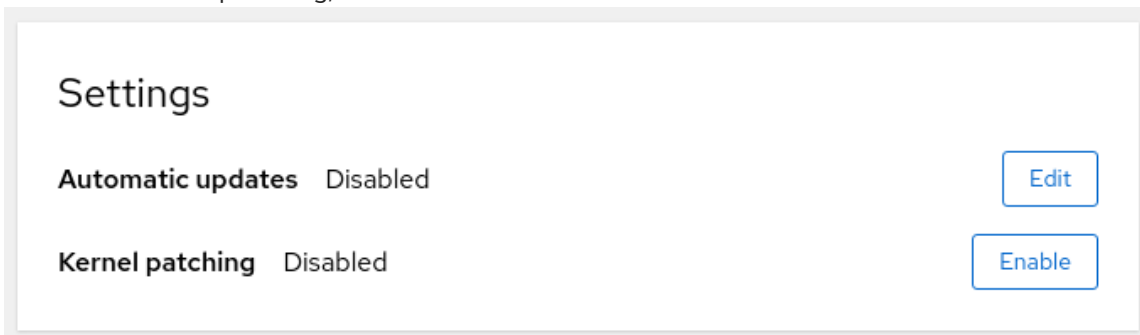
   c. Check the check box for applying kernel patches.

   d. Select whether you want to apply patches for current and future kernels, or for the current kernel only. If you choose to subscribe to applying patches for future kernels, the system will apply patches also for the upcoming kernel releases.
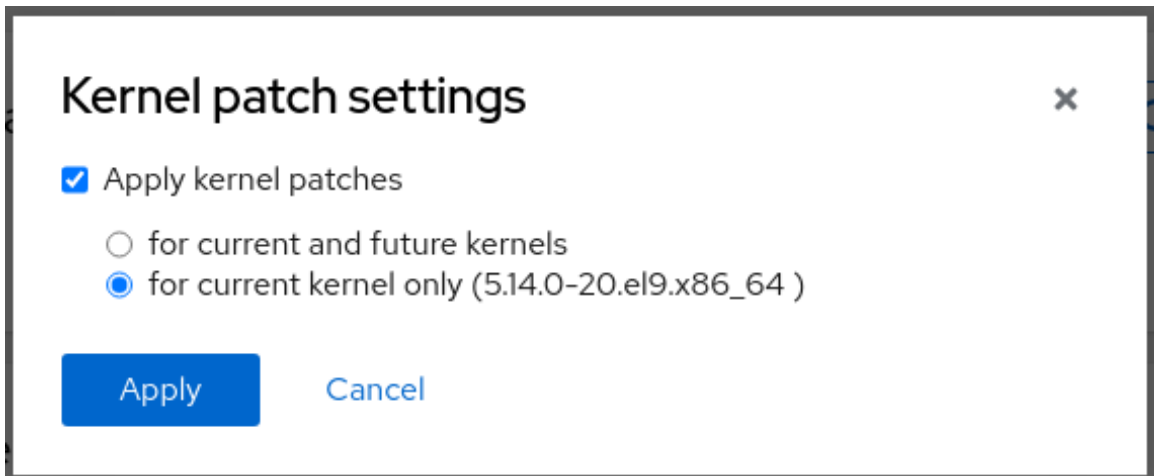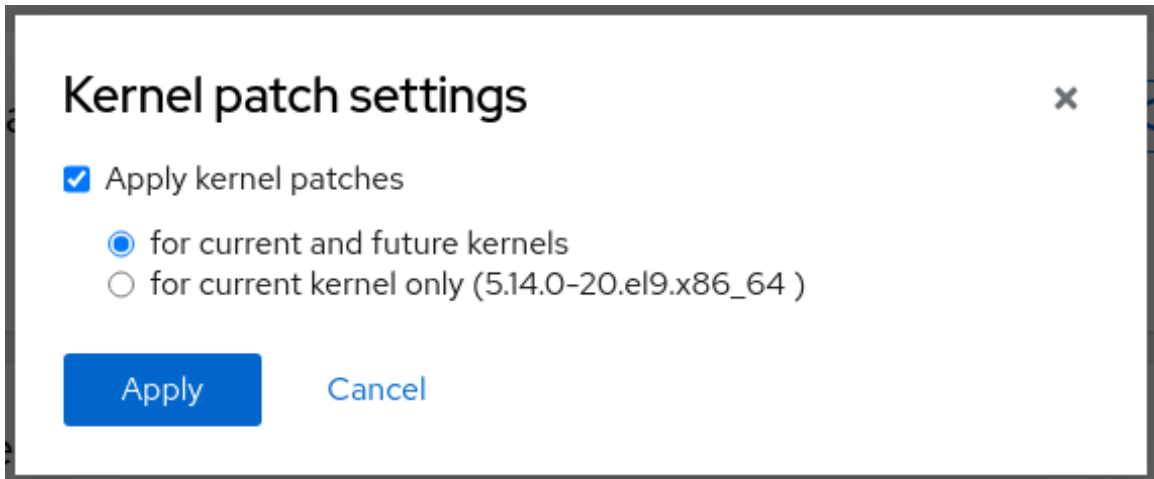
e. Click **Apply**.

Verification

- Check that the kernel patching is now **Enabled** in the **Settings** table of the **Software updates** section.
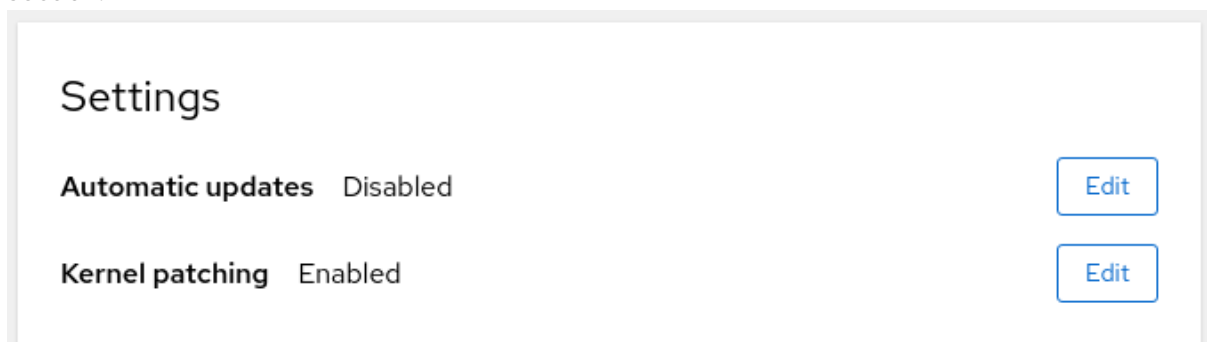
Additional resources

- Applying patches with kernel live patching

CHAPTER 26. MANAGING SUBSCRIPTIONS IN THE WEB CONSOLE

Manage your subscription for Red Hat Enterprise Linux 8 from the web console.

To get a subscription for your Red Hat Enterprise Linux, you need to have an account in the Red Hat Customer Portal or an activation key.

This chapter covers:

- Subscription management in the RHEL 8 web console.

- Registering subscriptions for your system in the web console with the Red Hat user name and password.

- Registering subscriptions with the activation key.

Prerequisites

- Purchased subscriptions.

- The system subjected to subscription has to be connected to the internet because the web console needs to communicate with the Red Hat Customer Portal.

26.1. SUBSCRIPTION MANAGEMENT IN THE WEB CONSOLE

The RHEL 8 web console provides an interface for using Red Hat Subscription Manager installed on your local system.

The Subscription Manager connects to the Red Hat Customer Portal and verifies all available:

- Active subscriptions

- Expired subscriptions

- Renewed subscriptions

If you want to renew the subscription or get a different one in Red Hat Customer Portal, you do not have to update the Subscription Manager data manually. The Subscription Manager synchronizes data with Red Hat Customer Portal automatically.

26.2. REGISTERING SUBSCRIPTIONS WITH CREDENTIALS IN THE WEB CONSOLE

Use the following steps to register a newly installed Red Hat Enterprise Linux with account credentials using the RHEL web console.

Prerequisites

- A valid user account on the Red Hat Customer Portal.
  See the Create a Red Hat Login page.

- Active subscription for your RHEL system.

Procedure
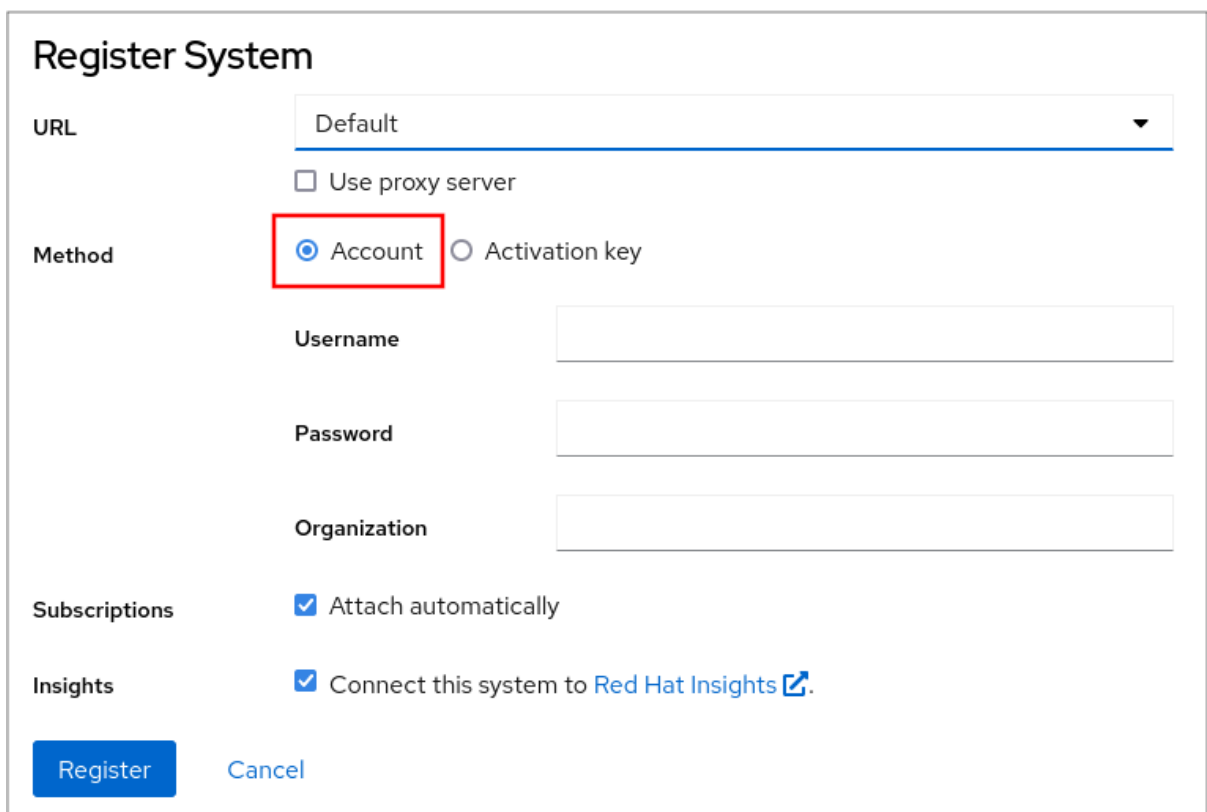
1. Log in to the RHEL web console. For details, see Logging in to the web console .

2. In the **Health** filed in the **Overview** page, click the **Not registered** warning, or click **Subscriptions** in the main menu to move to page with your subscription information.

   .
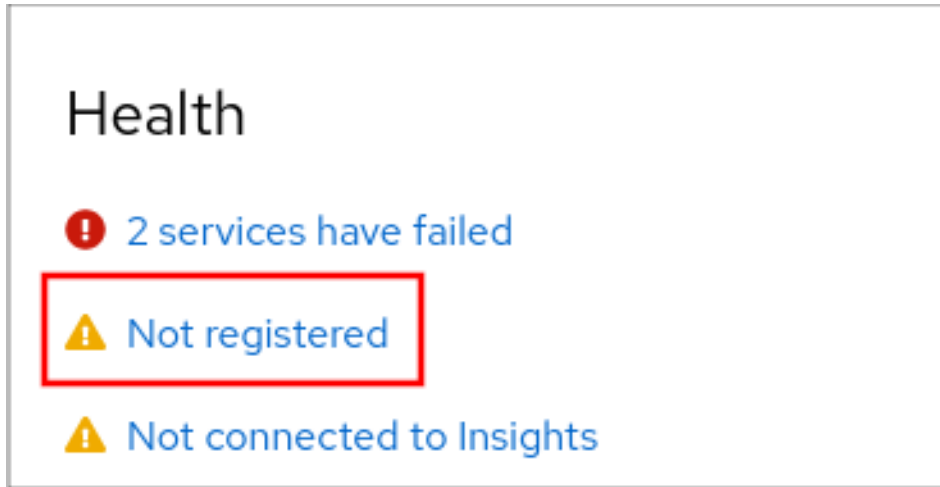
3. In the **Overview** filed, click **Register**.
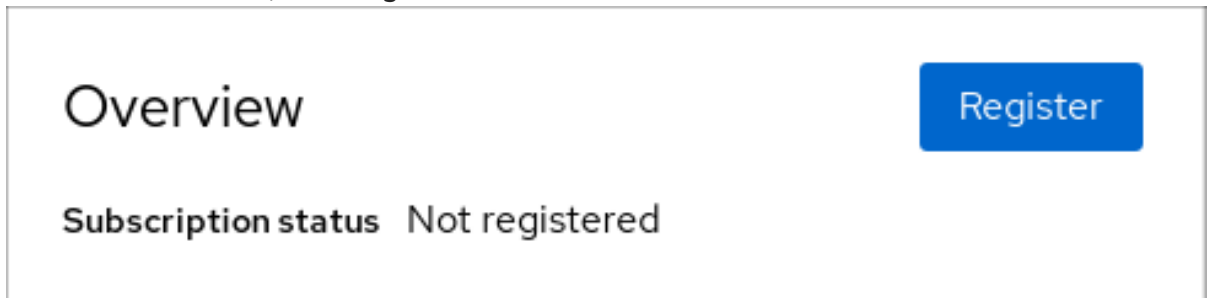
4. In the **Register system** dialog box, select that you want to register using your account credentials.

5. Enter your username.

6. Enter your password.

7. Optionally, enter your organization's name or ID.
   If your account belongs to more than one organization on the Red Hat Customer Portal, you have to add the organization name or organization ID. To get the org ID, go to your Red Hat contact point.

   - If you do not want to connect your system to Red Hat Insights, clear the **Insights** check box.

8. Click the **Register** button.

At this point, your Red Hat Enterprise Linux Enterprise Linux system has been successfully registered.

26.3. REGISTERING SUBSCRIPTIONS WITH ACTIVATION KEYS IN THE WEB CONSOLE

Use the following steps to register a newly installed Red Hat Enterprise Linux with an activation key using the RHEL web console.

Prerequisites

- If you do not have a user account in the portal, your vendor provides you with the activation key.

Procedure

1. Log in to the RHEL web console. For details, see Logging in to the web console .

2. In the **Health** filed in the **Overview** page, click the **Not registered** warning, or click **Subscriptions** in the main menu to move to page with your subscription information.
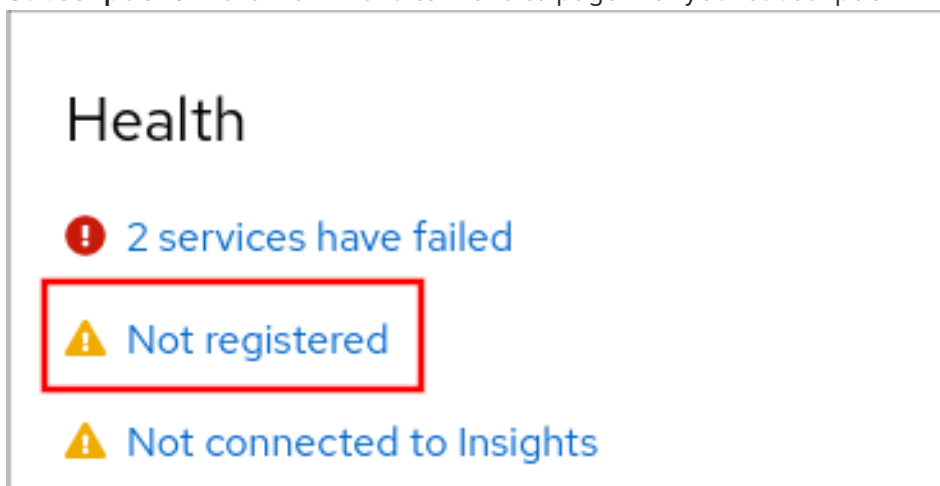
3. In the **Overview** filed, click **Register**.
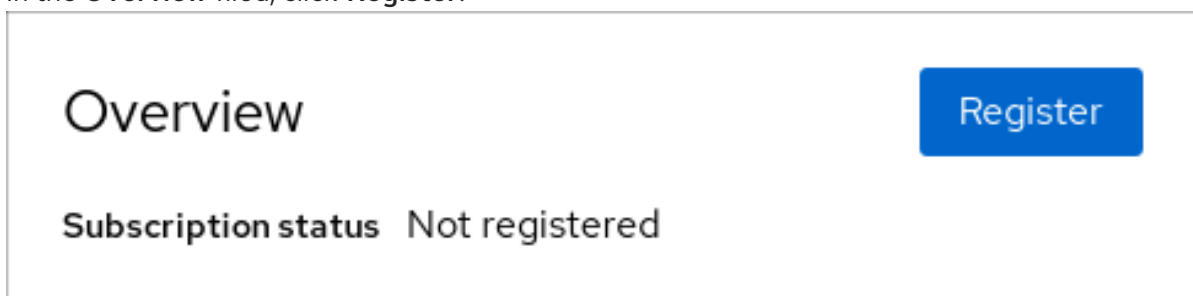
4. In the **Register system** dialog box, select that you want to register using an activation key.

Register System

URL	Default ▾
	☐ Use proxy server
Method	○ Account ● Activation key
Activation Key	key_one,key_two
Organization	
Subscriptions	☑ Attach automatically
Insights	☑ Connect this system to Red Hat Insights ⬏.

Register Cancel

5. Enter your key or keys.

6. Enter your organization's name or ID.
   To get the organization ID, go to your Red Hat contact point.

   - If you do not want to connect your system to Red Hat Insights, clear the **Insights** check box.

7. Click the **Register** button.

At this point, your Red Hat Enterprise Linux system has been successfully registered.

CHAPTER 27. CONFIGURING KDUMP IN THE WEB CONSOLE

You can setup and test the **kdump** configuration by using the RHEL 8 web console. The web console is part of a default installation of RHEL 8 and enables or disables the **kdump** service at boot time. Further, the web console enables you to configure the reserved memory for **kdump**; or to select the *vmcore* saving location in an uncompressed or compressed format.

27.1. CONFIGURING KDUMP MEMORY USAGE AND TARGET LOCATION IN WEB CONSOLE

You can configure the memory reserve for the **kdump** kernel and also specify the target location to capture the **vmcore** dump file with the RHEL web console interface.

Procedure

1. In the web console, open the **Kernel Dump** tab and start the **kdump** service by setting the **Kernel crash dump** switch to on.

2. Configure the **kdump** memory usage in the command line.

3. In the **Kernel Dump** tab, go to **Crash dump location** and click the link with the path to the dump location.
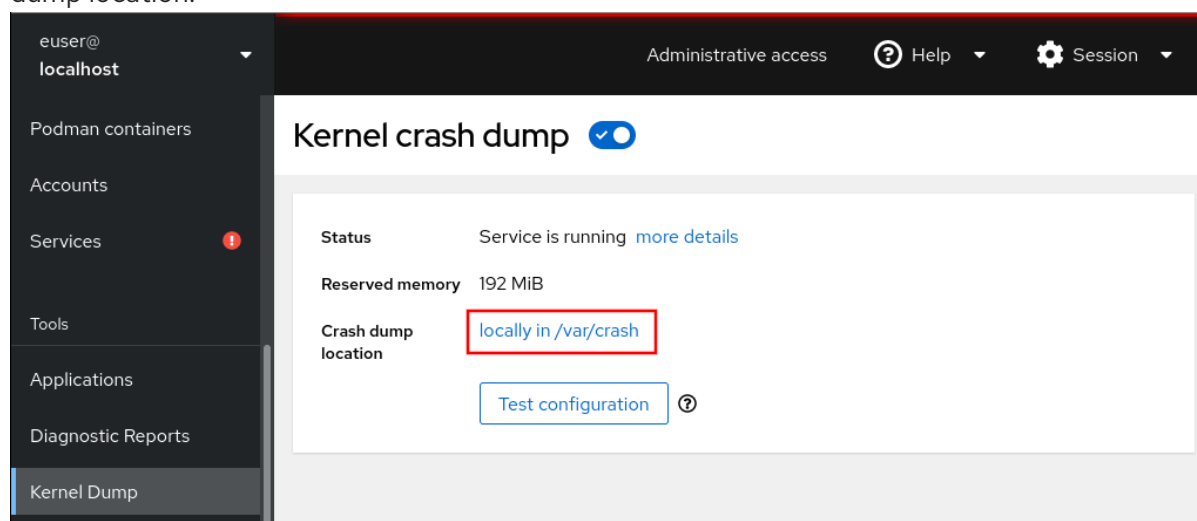
4. Specify the target directory for saving the **vmcore** dump file:

   - For a local filesystem, select **Local Filesystem** from the drop-down menu.

- For a remote system by using the SSH protocol, select **Remote over SSH** from the drop-down menu and specify the following fields:

  - In the **Server** field, enter the remote server address

  - In the **ssh key** field, enter the ssh key location

  - In the **Directory** field, enter the target directory

- For a remote system by using the NFS protocol, select **Remote over NFS** from the drop-down menu and specify the following fields:

  - In the **Server** field, enter the remote server address

  - In the **Export** field, enter the location of the shared folder of an NFS server

  - In the **Directory** field, enter the target directory

> **NOTE**
>
> You can reduce the size of the **vmcore** file by selecting the **Compression** check box.

Verification

1. Click the **Test configuration**.

Status	Service is running more details
Reserved memory	192 MiB
Crash dump location	locally in /var/crash
	Test configuration ⓘ

2. Click **Crash system** under **Test kdump settings**.

> **WARNING**
>
> When you initiate the system crash, the kernel's operation stops and results in a system crash with data loss.

Additional resources

- Supported kdump targets

- Using secure communications between two systems with OpenSSH

- Getting started using the RHEL web console

CHAPTER 28. MANAGING VIRTUAL MACHINES IN THE WEB CONSOLE

To manage virtual machines in a graphical interface on a RHEL 8 host, you can use the **Virtual Machines** pane in the RHEL 8 web console.
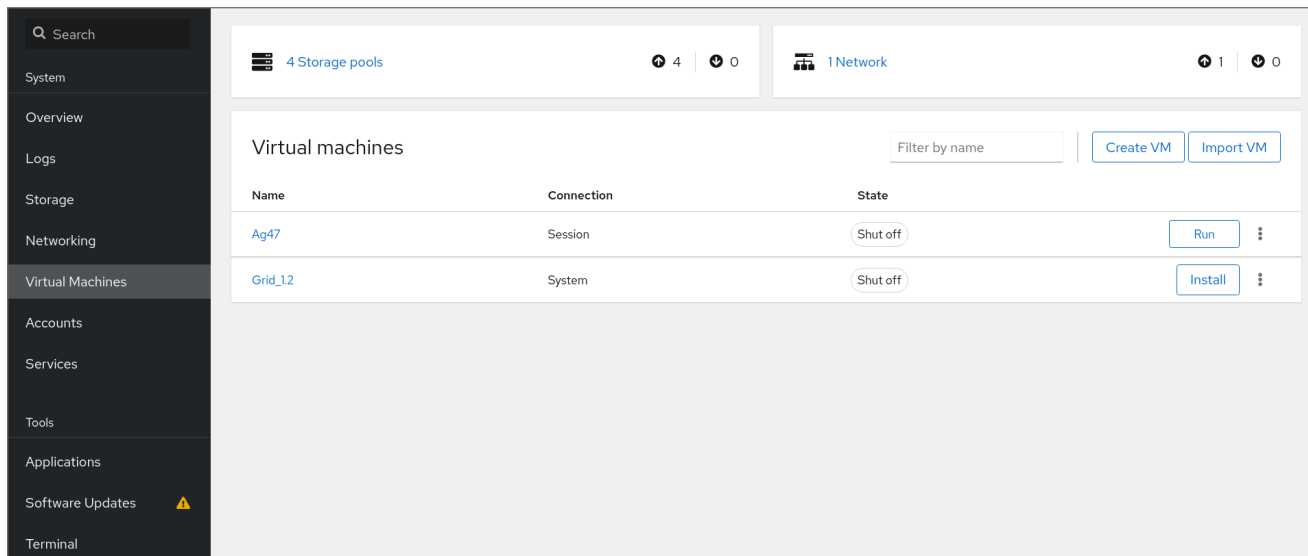
28.1. OVERVIEW OF VIRTUAL MACHINE MANAGEMENT BY USING THE WEB CONSOLE

The RHEL 8 web console is a web-based interface for system administration. As one of its features, the web console provides a graphical view of virtual machines (VMs) on the host system, and makes it possible to create, access, and configure these VMs.

Note that to use the web console to manage your VMs on RHEL 8, you must first install a web console plug-in for virtualization.

Next steps

- For instructions on enabling VMs management in your web console, see Setting up the web console to manage virtual machines.

- For a comprehensive list of VM management actions that the web console provides, see Virtual machine management features available in the web console.

- For a list of features that are currently not available in the web console but can be used in the **virt-manager** application, see Differences between virtualization features in Virtual Machine Manager and the web console.

28.2. SETTING UP THE WEB CONSOLE TO MANAGE VIRTUAL MACHINES

Before using the RHEL 8 web console to manage virtual machines (VMs), you must install the web console virtual machine plug-in on the host.

Prerequisites

- Ensure that the web console is installed and enabled on your machine.

  > # **systemctl status cockpit.socket**
  > cockpit.socket - Cockpit Web Service Socket
  > Loaded: loaded (/usr/lib/systemd/system/cockpit.socket
  > [...]

  If this command returns **Unit cockpit.socket could not be found**, follow the Installing the web console document to enable the web console.

Procedure

- Install the **cockpit-machines** plug-in.

  > # **yum install cockpit-machines**

Verification

1. Access the web console, for example by entering the **https://localhost:9090** address in your browser.

2. Log in.

3. If the installation was successful, **Virtual Machines** appears in the web console side menu.
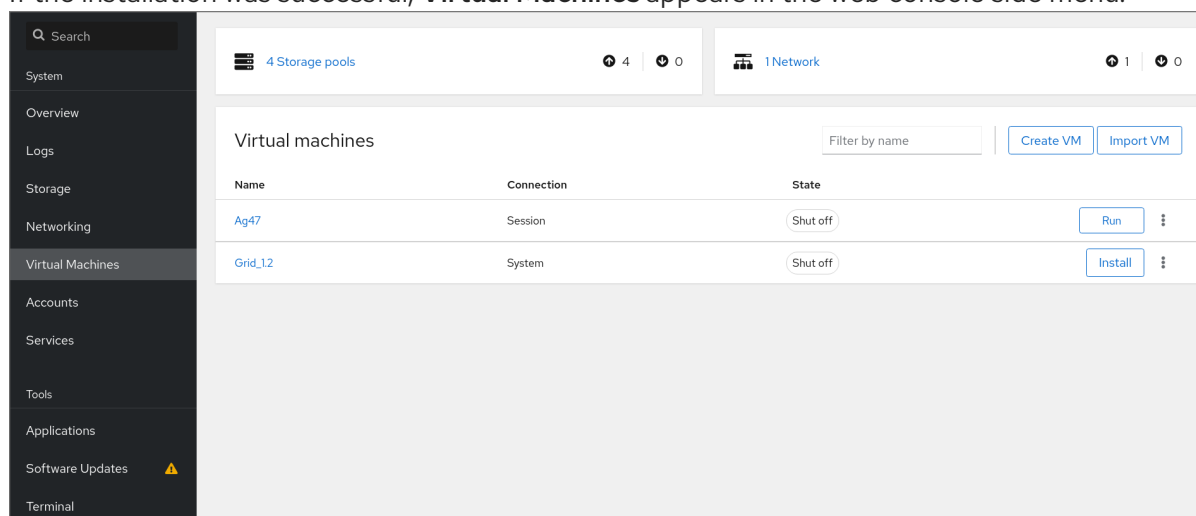
Additional resources

- Managing systems by using the RHEL 8 web console

28.3. RENAMING VIRTUAL MACHINES BY USING THE WEB CONSOLE

You might require renaming an existing virtual machine (VM) to avoid naming conflicts or assign a new unique name based on your use case. To rename the VM, you can use the RHEL web console.

Prerequisites

- The web console VM plug-in is installed on your system.

- The VM is shut down.

Procedure

1. In the **Virtual Machines** interface, click the Menu button ⋮ of the VM that you want to rename. A drop-down menu appears with controls for various VM operations.

2. Click **Rename**.
   The **Rename a VM** dialog appears.

3. In the **New name** field, enter a name for the VM.

4. Click **Rename**.

Verification

- Check that the new VM name has appeared in the **Virtual Machines** interface.

28.4. VIRTUAL MACHINE MANAGEMENT FEATURES AVAILABLE IN THE WEB CONSOLE

By using the RHEL 8 web console, you can perform the following actions to manage the virtual machines (VMs) on your system.

Table 28.1. VM tasks that can be performed in the RHEL 8 web console

Task	For details, see:
Create a VM and install it with a guest operating system	Creating virtual machines and installing guest operating systems by using the web console
Delete a VM.	Deleting virtual machines by using the web console
Start, shut down, and restart the VM	Starting virtual machines by using the web console and Shutting down and restarting virtual machines by using the web console
Connect to and interact with a VM using a variety of consoles	Interacting with virtual machines by using the web console
View a variety of information about the VM	Viewing virtual machine information by using the web console

Task	For details, see:
Adjust the host memory allocated to a VM	Adding and removing virtual machine memory by using the web console
Manage network connections for the VM	Using the web console for managing virtual machine network interfaces
Manage the VM storage available on the host and attach virtual disks to the VM	Managing storage for virtual machines
Configure the virtual CPU settings of the VM	Managing virtual CPUs by using the web console
Live migrate a VM	Live migrating a virtual machine by using the web console
Rename a VM	Renaming virtual machines by using the web console
Share files between the host and the VM	Sharing files between the host and its virtual machines
Manage host devices	Managing virtual devices by using the web console
Manage virtual optical drives	Managing virtual optical drives
Attach watchdog device	Attaching a watchdog device to a virtual machine by using the web console

28.5. DIFFERENCES BETWEEN VIRTUALIZATION FEATURES IN VIRTUAL MACHINE MANAGER AND THE WEB CONSOLE

The Virtual Machine Manager (**virt-manager**) application is supported in RHEL 8, but has been deprecated. The web console is intended to become its replacement in a subsequent major release. It is, therefore, recommended that you get familiar with the web console for managing virtualization in a GUI.

However, in RHEL 8, some VM management tasks can only be performed in **virt-manager** or the command line. The following table highlights the features that are available in **virt-manager** but not available in the RHEL 8.0 web console.

If a feature is available in a later minor version of RHEL 8, the minimum RHEL 8 version appears in the *Support in web console introduced* column.

Table 28.2. VM managemennt tasks that cannot be performed using the web console in RHEL 8.0

Task	Support in web console introduced	Alternative method by using CLI

Task	Support in web console introduced	Alternative method by using CLI
Setting a virtual machine to start when the host boots	RHEL 8.1	**virsh autostart**
Suspending a virtual machine	RHEL 8.1	**virsh suspend**
Resuming a suspended virtual machine	RHEL 8.1	**virsh resume**
Creating file-system directory storage pools	RHEL 8.1	**virsh pool-define-as**
Creating NFS storage pools	RHEL 8.1	**virsh pool-define-as**
Creating physical disk device storage pools	RHEL 8.1	**virsh pool-define-as**
Creating LVM volume group storage pools	RHEL 8.1	**virsh pool-define-as**
Creating partition-based storage pools	*CURRENTLY UNAVAILABLE*	**virsh pool-define-as**
Creating GlusterFS-based storage pools	*CURRENTLY UNAVAILABLE*	**virsh pool-define-as**
Creating vHBA-based storage pools with SCSI devices	*CURRENTLY UNAVAILABLE*	**virsh pool-define-as**
Creating Multipath-based storage pools	*CURRENTLY UNAVAILABLE*	**virsh pool-define-as**
Creating RBD-based storage pools	*CURRENTLY UNAVAILABLE*	**virsh pool-define-as**
Creating a new storage volume	RHEL 8.1	**virsh vol-create**
Adding a new virtual network	RHEL 8.1	**virsh net-create** or **virsh net-define**
Deleting a virtual network	RHEL 8.1	**virsh net-undefine**
Creating a bridge from a host machine's interface to a virtual machine	*CURRENTLY UNAVAILABLE*	**virsh iface-bridge**
Creating a snapshot	*CURRENTLY UNAVAILABLE*	**virsh snapshot-create-as**

Task	Support in web console introduced	Alternative method by using CLI
Reverting to a snapshot	*CURRENTLY UNAVAILABLE*	**virsh snapshot-revert**
Deleting a snapshot	*CURRENTLY UNAVAILABLE*	**virsh snapshot-delete**
Cloning a virtual machine	RHEL 8.4	**virt-clone**
Migrating a virtual machine to another host machine	RHEL 8.5	**virsh migrate**
Attaching a host device to a VM	RHEL 8.5	**virt-xml --add-device**
Removing a host device from a VM	RHEL 8.5	**virt-xml --remove-device**

Additional resources

- [Getting started with Virtual Machine Manager in RHEL 7](*Deprecated in RHEL 8 and later*)

CHAPTER 29. MANAGING REMOTE SYSTEMS IN THE WEB CONSOLE

Connect to the remote systems and manage them in the RHEL 8 web console.

The following chapter describes:

- The optimal topology of connected systems.

- How to add and remove remote systems.

- When, why, and how to use SSH keys for remote system authentication.

- How to configure a web console client to allow a user authenticated with a smart card to **SSH** to a remote host and access services on it.

Prerequisites

- Opened the SSH service on remote systems.

29.1. REMOTE SYSTEM MANAGER IN THE WEB CONSOLE

Using the RHEL 8 web console to manage remote systems in the network requires considering the topology of connected servers.

For optimal security, Red Hat recommends the following connection setup:

- Use one system with the web console as a bastion host. The bastion host is a system with opened HTTPS port.

- All other systems communicate through SSH.

With the web interface running on the bastion host, you can reach all other systems through the SSH protocol using port 22 in the default configuration.
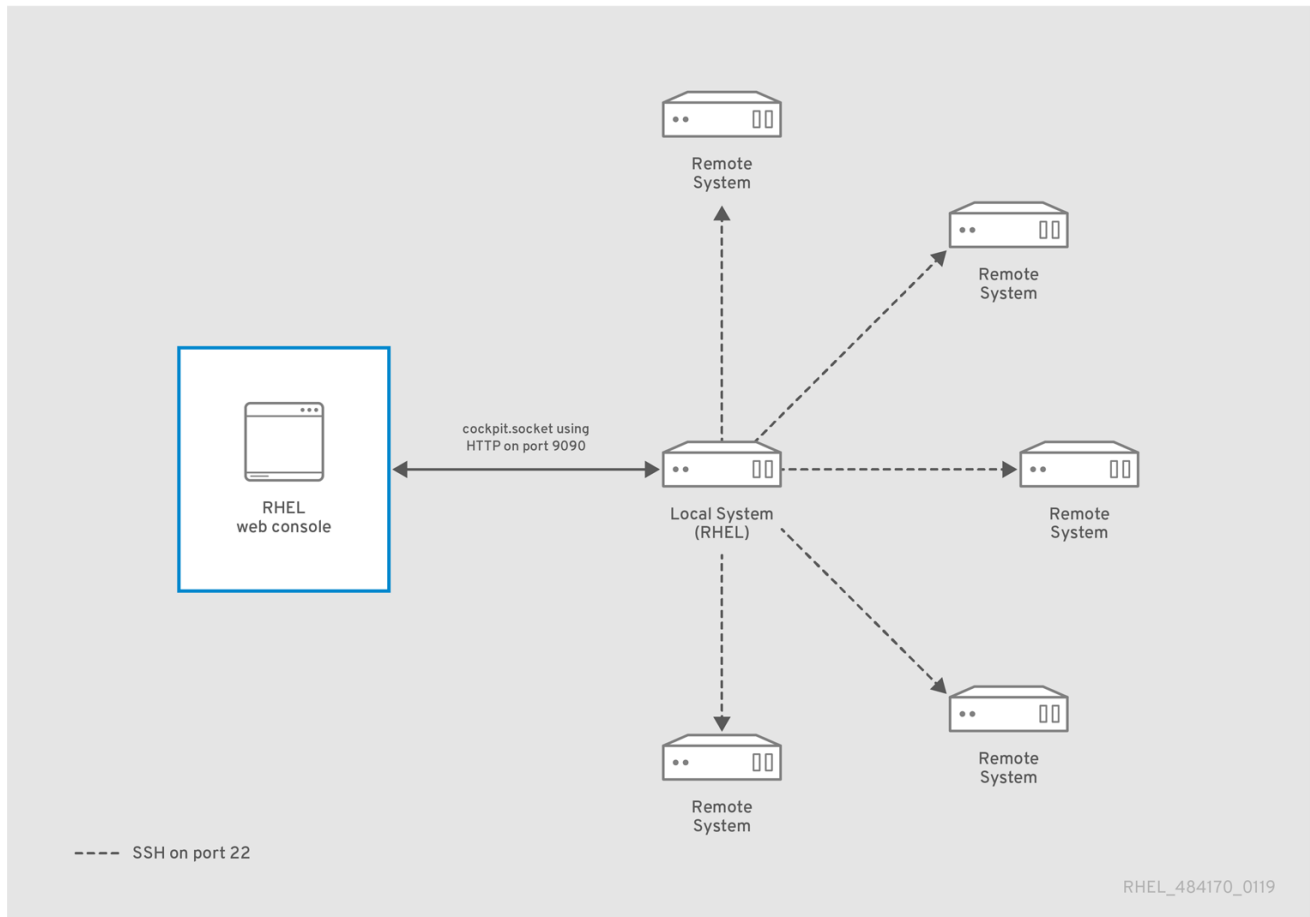
29.2. ADDING REMOTE HOSTS TO THE WEB CONSOLE

You can connect other systems with a user name and password.

Prerequisites

- You need to be logged into the web console with administration privileges. For details, see Logging in to the web console .

Procedure

1. In the RHEL 8 web console, click on your **username@hostname** in the top left corner of the **Overview** page.
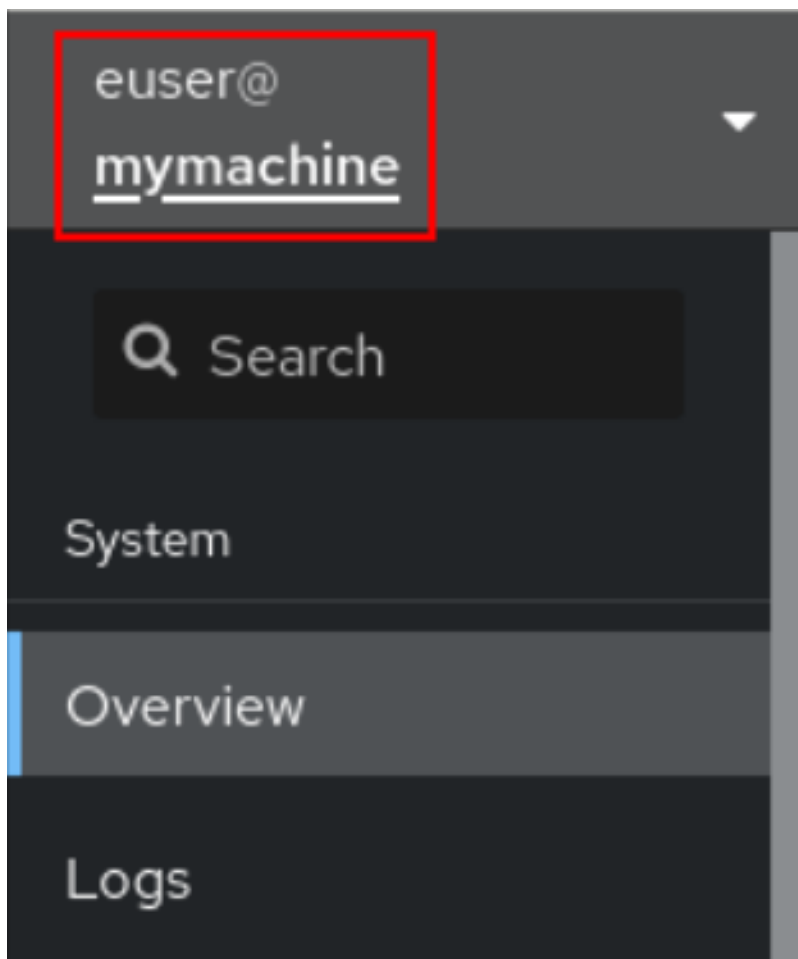
2. In the drop down menu, click the **Add new host** button.
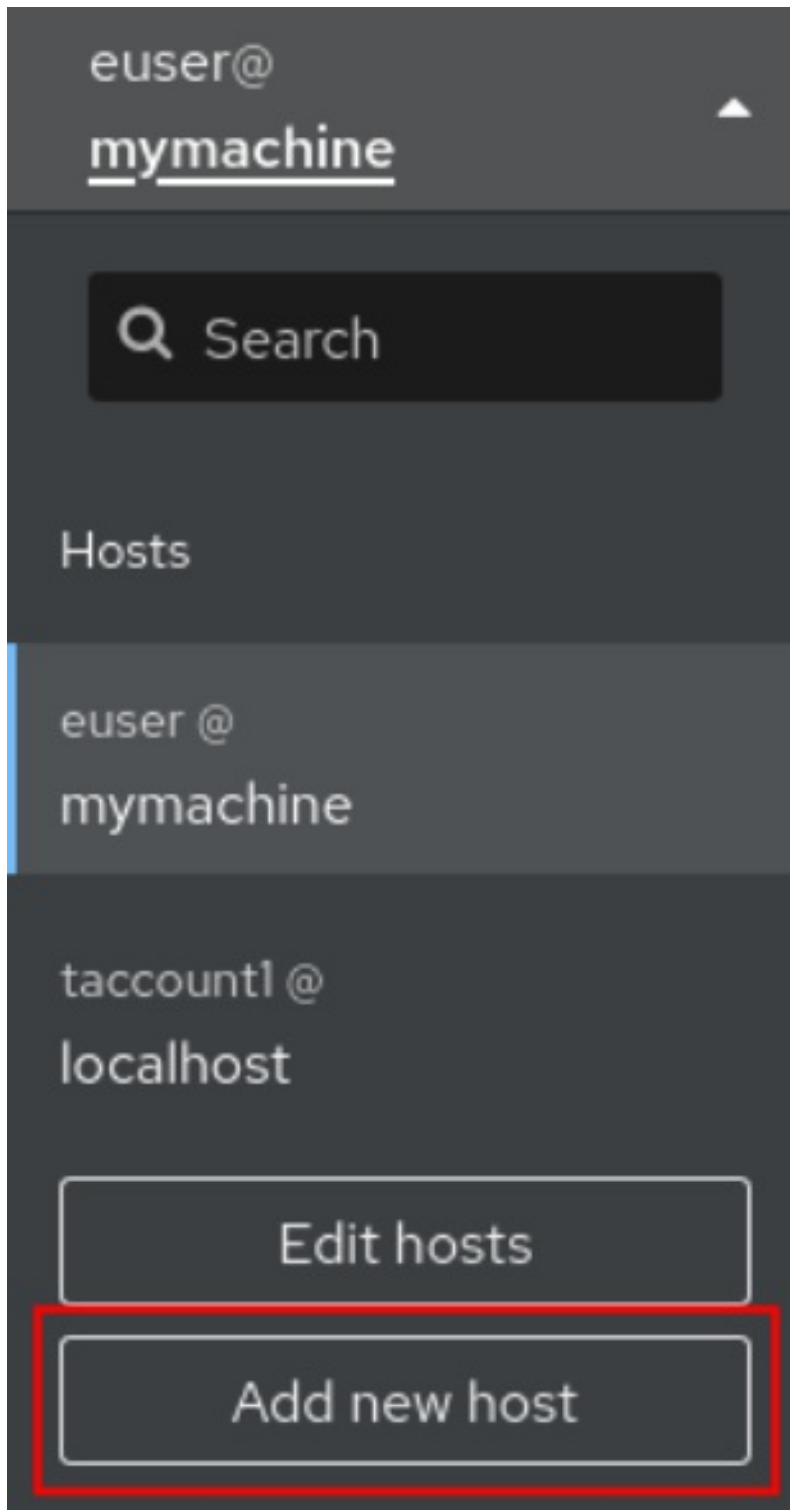
3. In the **Add new host** dialog box, specify the host you want to add.

4. (Optional) Add the user name for the account to which you want to connect.
   You can use any user account of the remote system. However, if you use credentials of a user account without administration privileges, you will not be able to perform administration tasks.

   If you use the same credentials as for your local system, the web console will authenticate remote systems automatically every time you log in. However, using the same credentials on more machines could be a potential security risk.

5. (Optional) Click the **Color** field to change the color of the system.

6. Click **Add**.

The new host will appear in the list of hosts in the **username@hostname** drop down menu.

> **NOTE**
>
> The web console does not save passwords used to log in to remote systems which means that you have to log in again after each system restart. Next time you log in, click the **Log in** button placed on the main screen of the disconnected remote system to open the login dialog.
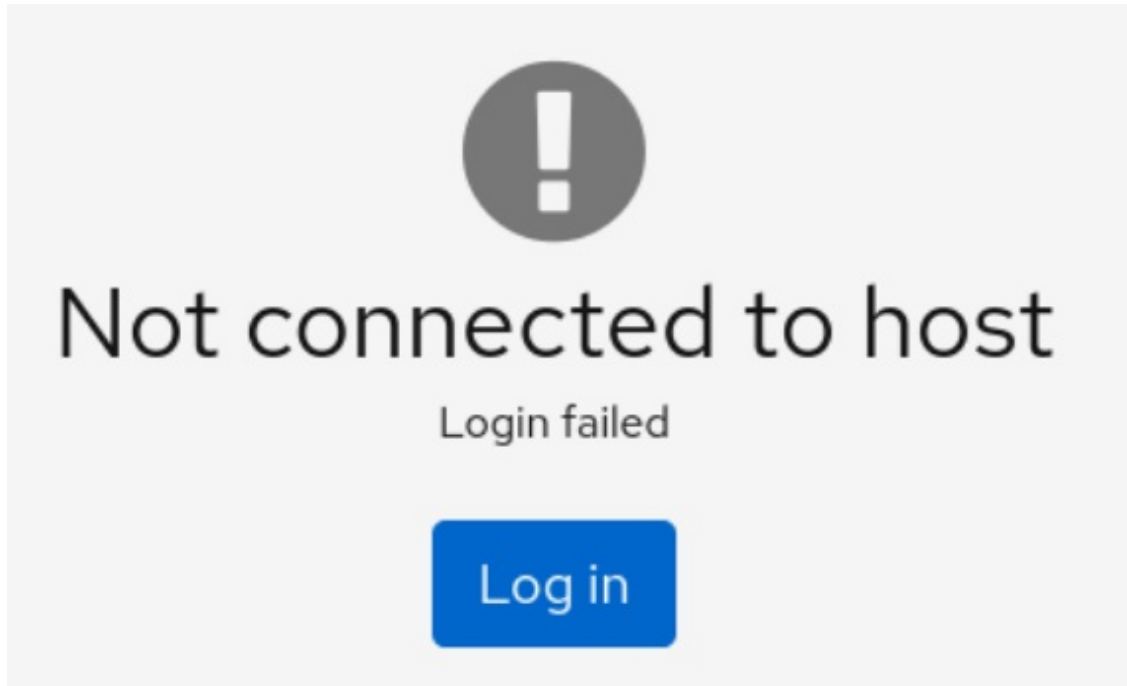
29.3. REMOVING REMOTE HOSTS FROM THE WEB CONSOLE

You can remove other systems from the web console.

Prerequisites

- Remote systems added.
  For details, see Section 29.2, "Adding remote hosts to the web console" .

- You must be logged into the web console with administrator privileges.
  For details, see Logging in to the web console .

Procedure

1. Log in to the RHEL 8 web console.

2. Click on your **username@hostname** in the top left corner of the **Overview** page.
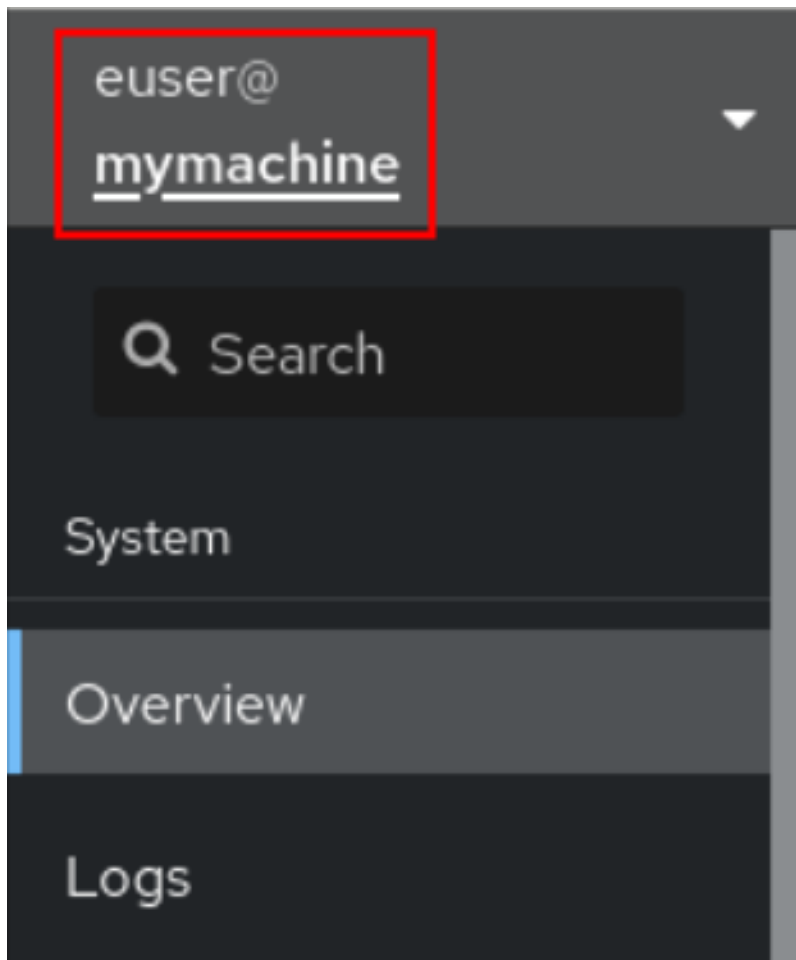
3. Click the **Edit hosts** icon.
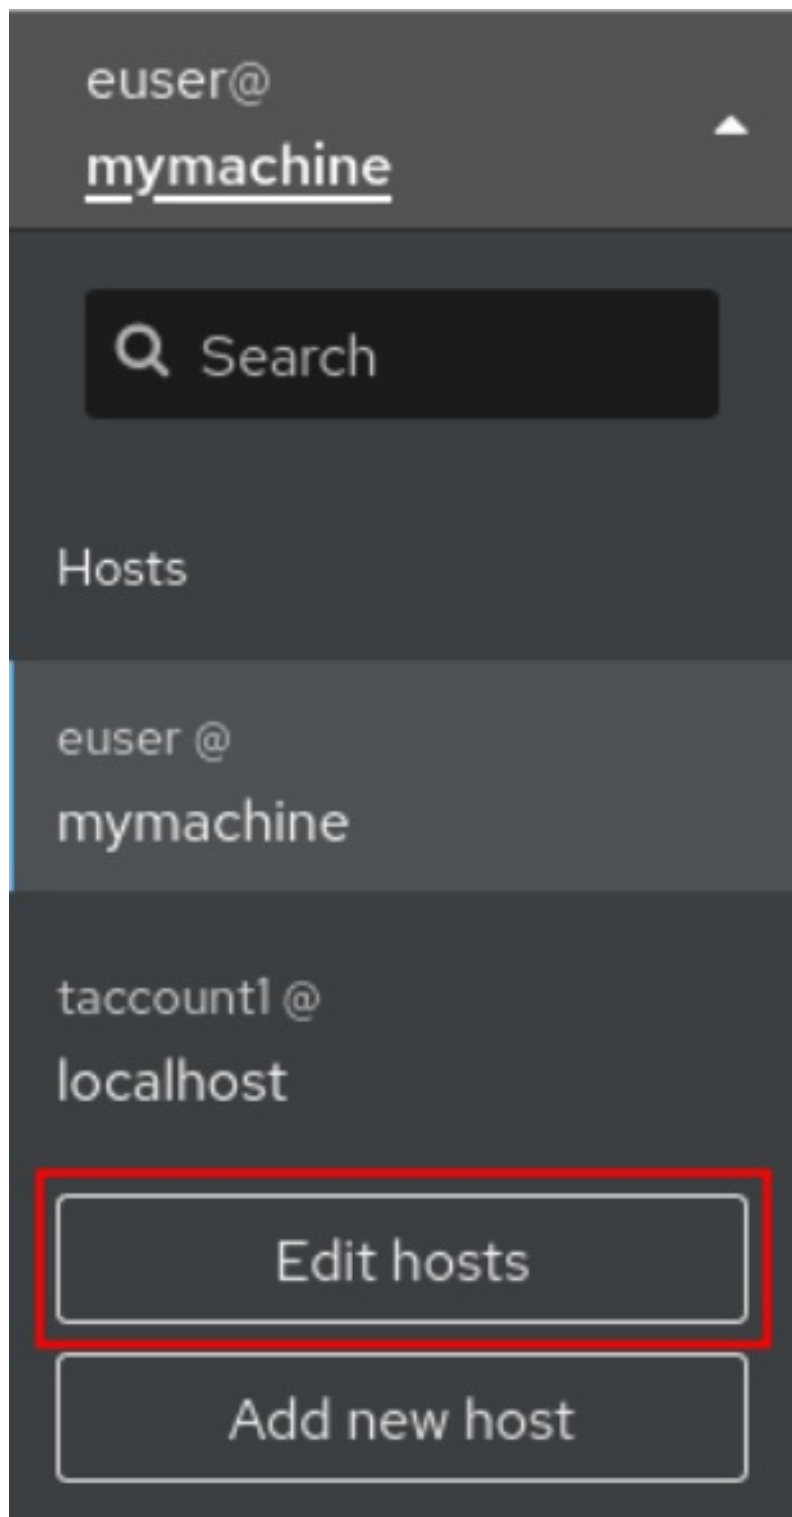
4. To remove a host from web console, click the red minus sign ▪ button next to its host name. Note that you cannot remove a host you are currently connected to.
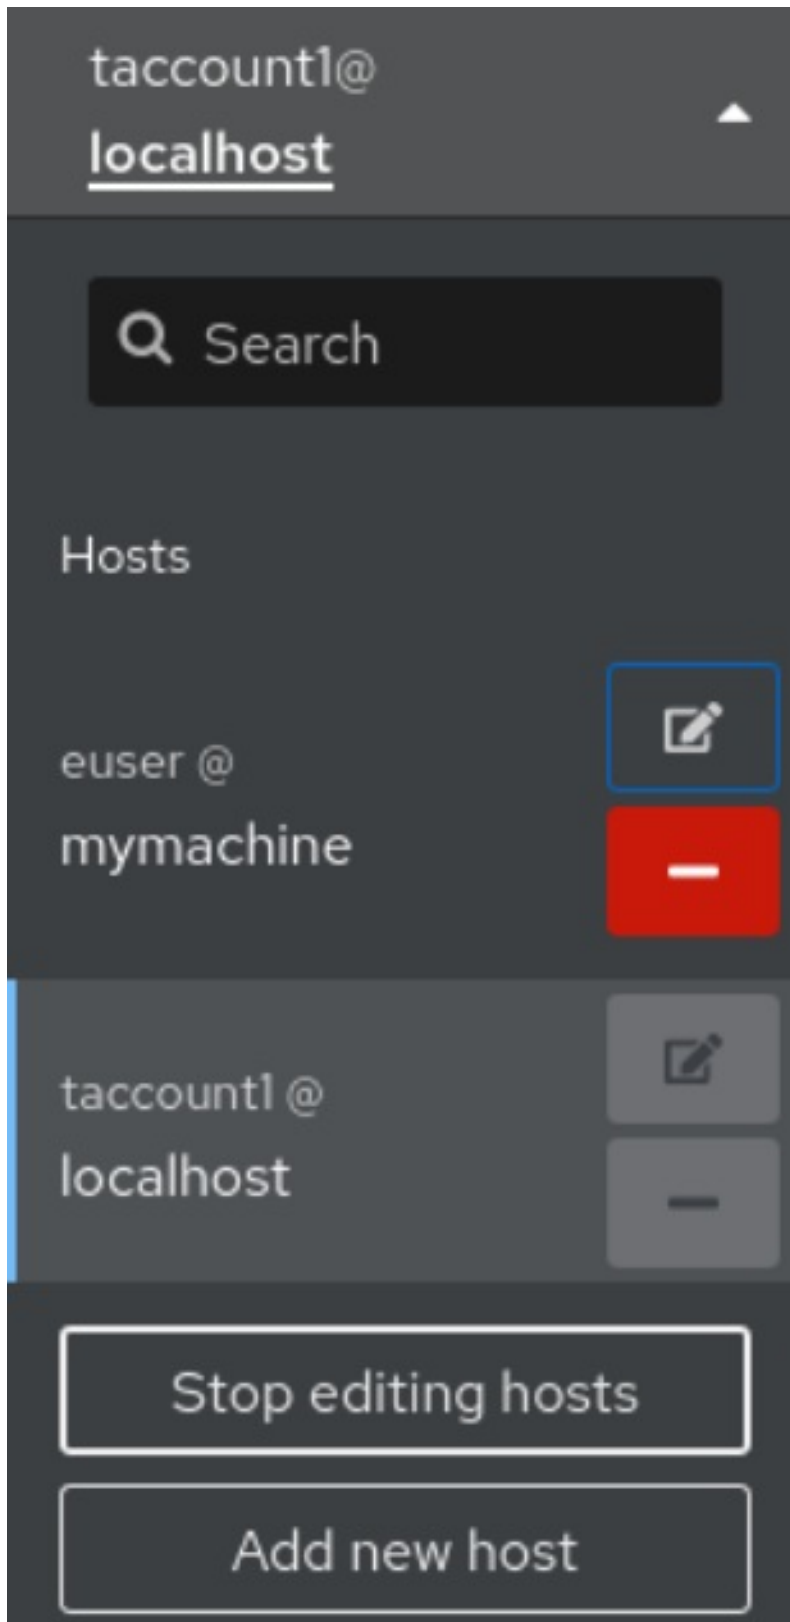
As a result, the server is removed from your web console.

29.4. ENABLING SSH LOGIN FOR A NEW HOST

When you add a new host you can also log into it with an SSH key. If you already have an SSH key on your system, the web console will use the existing one; otherwise, the web console can create a key.

Prerequisites

- You need to be logged into the web console with administration privileges.
  For details, see Logging in to the web console .

Procedure

1. In the RHEL 8 web console, click on your **username@hostname** in the top left corner of the **Overview** page.
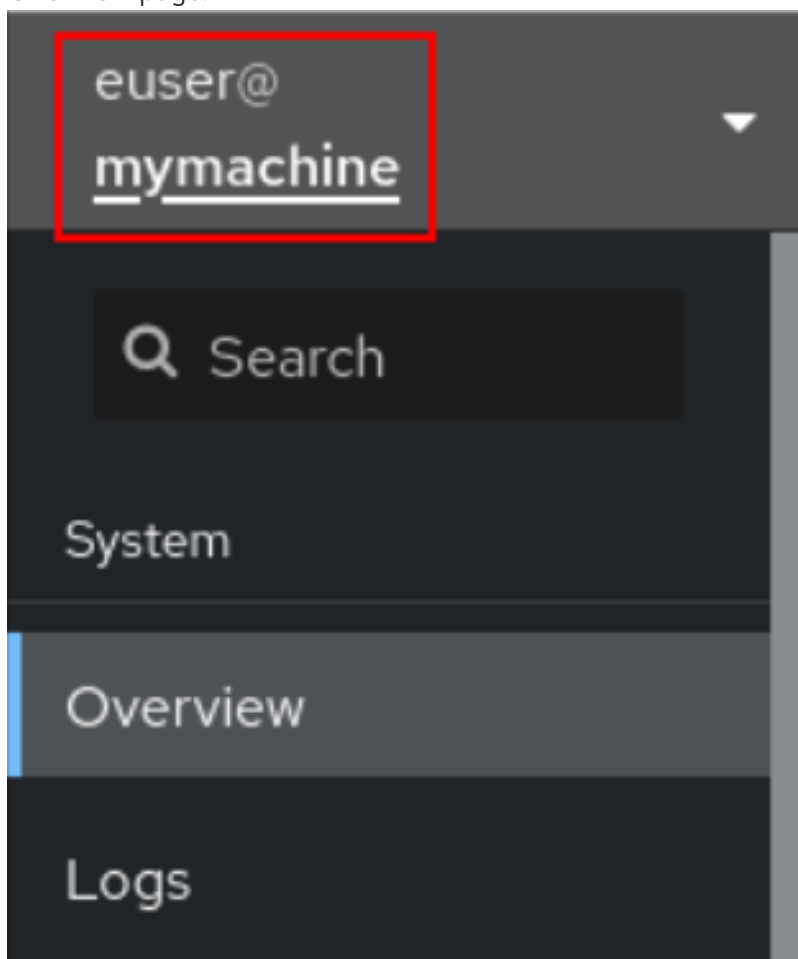
2. In the drop down menu, click the **Add new host** button.
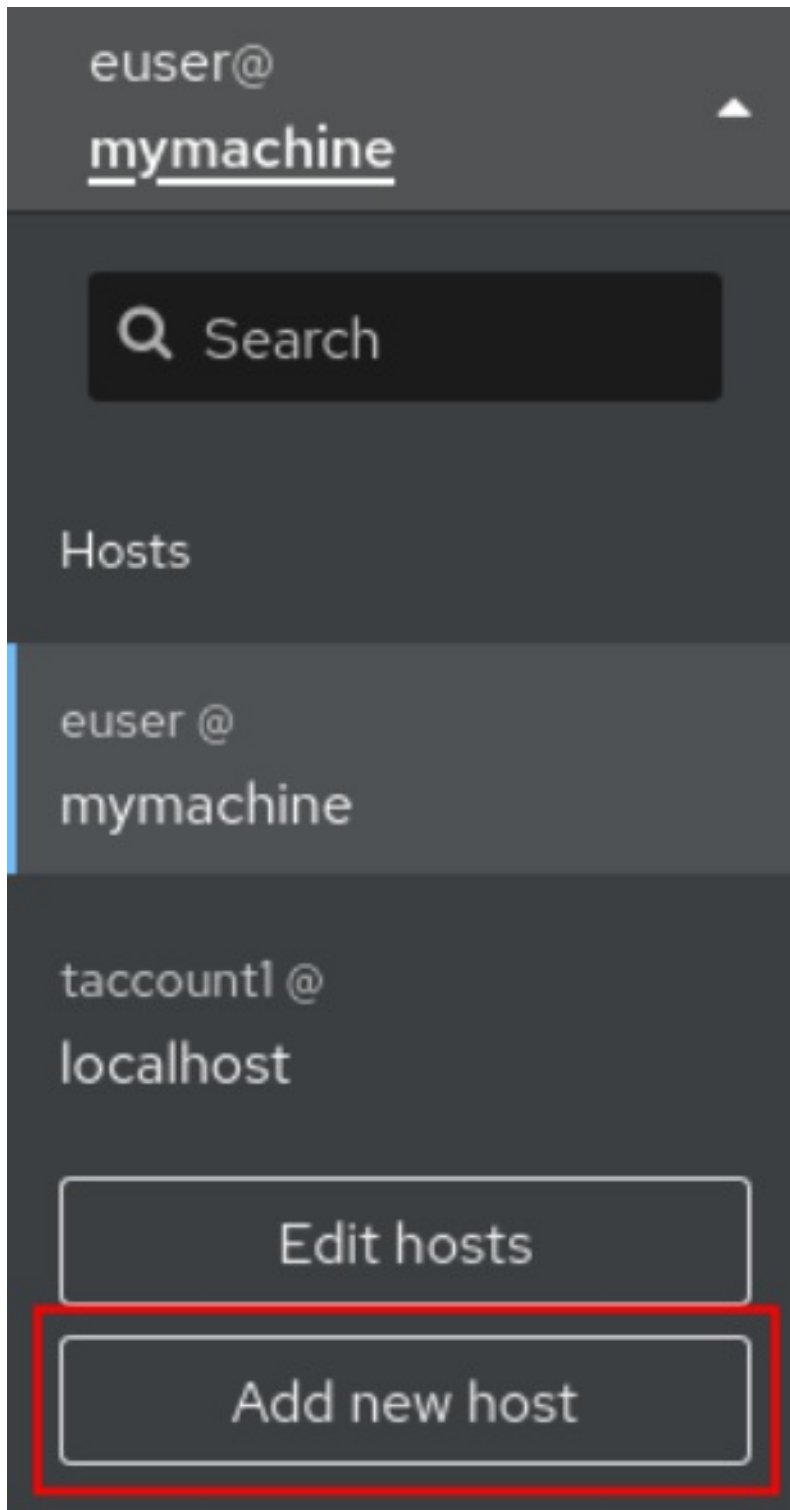
3. In the **Add new host** dialog box, specify the host you want to add.

4. Add the user name for the account to which you want to connect.
   You can use any user account of the remote system. However, if you use credentials of a user account without administration privileges, you will not be able to perform administration tasks.

5. (Optional) Click the **Color** field to change the color of the system.

6. Click **Add**.
   A new dialog window will appear asking for a password.

7. Enter the user account password.

8. Check **Authorize ssh key** if you already have an SSH key.

9. Check **Create a new SSH key and authorize it** if you do not have an SSH key. The web console will create it for you.

a. Add a password for the SSH key.

b. Confirm the password.

10. Click **Log in**
    The new host will appear in the list of hosts in the **username@hostname** drop down menu.

Verification steps

1. Log out.

2. Log back in.

3. Click **Log in** in the **Not connected to host** screen.

4. Select **SSH key** as your authentication option.

5. Enter your key password.

6. Click **Log in**.

Additional resources

- Using secure communications between two systems with OpenSSH

29.5. CONSTRAINED DELEGATION IN IDENTITY MANAGEMENT

The Service for User to Proxy (**S4U2proxy**) extension provides a service that obtains a service ticket to another service on behalf of a user. This feature is known as **constrained delegation**. The second service is typically a proxy performing some work on behalf of the first service, under the authorization context of the user. Using constrained delegation eliminates the need for the user to delegate their full ticket-granting ticket (TGT).

Identity Management (IdM) traditionally uses the Kerberos **S4U2proxy** feature to allow the web server framework to obtain an LDAP service ticket on the user's behalf. The IdM-AD trust system also uses constrained delegation to obtain a **cifs** principal.

You can use the **S4U2proxy** feature to configure a web console client to allow an IdM user that has authenticated with a smart card to achieve the following:

- Run commands with superuser privileges on the RHEL host on which the web console service is running without being asked to authenticate again.

- Access a remote host using **SSH** and access services on the host without being asked to authenticate again.

Additional resources

- Using Ansible to configure a web console to allow a user authenticated with a smart card to SSH to a remote host without being asked to authenticate again

- Using Ansible to configure a web console to allow a user authenticated with a smart card to run sudo without being asked to authenticate again

- S4U2proxy

- Service constrained delegation

29.6. CONFIGURING A WEB CONSOLE TO ALLOW A USER AUTHENTICATED WITH A SMART CARD TO SSH TO A REMOTE HOST WITHOUT BEING ASKED TO AUTHENTICATE AGAIN

After you have logged in to a user account on the RHEL web console, as an Identity Management (IdM) system administrator you might need to connect to remote machines by using the **SSH** protocol. You can use the constrained delegation feature to use **SSH** without being asked to authenticate again.

Follow this procedure to configure the web console to use constrained delegation. In the example below, the web console session runs on the **myhost.idm.example.com** host and it is being configured to access the **remote.idm.example.com** host by using **SSH** on behalf of the authenticated user.

Prerequisites

- You have obtained an IdM **admin** ticket-granting ticket (TGT).

- You have **root** access to **remote.idm.example.com**.

- The web console service is present in IdM.

- The **remote.idm.example.com** host is present in IdM.

- The web console has created an **S4U2Proxy** Kerberos ticket in the user session. To verify that this is the case, log in to the web console as an IdM user, open the **Terminal** page, and enter:

> $ **klist**
> Ticket cache: FILE:/run/user/1894000001/cockpit-session-3692.ccache
> Default principal: user@IDM.EXAMPLE.COM
>
> Valid starting   Expires   Service principal

> **07/30/21 09:19:06 07/31/21 09:19:06**
> **HTTP/myhost.idm.example.com@IDM.EXAMPLE.COM**
> 07/30/21 09:19:06 07/31/21 09:19:06 krbtgt/IDM.EXAMPLE.COM@IDM.EXAMPLE.COM
>   for client HTTP/myhost.idm.example.com@IDM.EXAMPLE.COM

Procedure

1. Create a list of the target hosts that can be accessed by the delegation rule:

   a. Create a service delegation target:

      > $ **ipa servicedelegationtarget-add cockpit-target**

   b. Add the target host to the delegation target:

      > $ **ipa servicedelegationtarget-add-member cockpit-target \ --**
      > **principals=host/remote.idm.example.com@IDM.EXAMPLE.COM**

2. Allow **cockpit** sessions to access the target host list by creating a service delegation rule and adding the **HTTP** service Kerberos principal to it:

   a. Create a service delegation rule:

      > $ **ipa servicedelegationrule-add cockpit-delegation**

   b. Add the web console client to the delegation rule:

      > $ **ipa servicedelegationrule-add-member cockpit-delegation \ --**
      > **principals=HTTP/myhost.idm.example.com@IDM.EXAMPLE.COM**

   c. Add the delegation target to the delegation rule:

      > $ **ipa servicedelegationrule-add-target cockpit-delegation \ --**
      > **servicedelegationtargets=cockpit-target**

3. Enable Kerberos authentication on the **remote.idm.example.com** host:

   a. **SSH** to **remote.idm.example.com** as **root**.

   b. Open the **/etc/ssh/sshd_config** file for editing.

   c. Enable **GSSAPIAuthentication** by uncommenting the **GSSAPIAuthentication no** line and replacing it with **GSSAPIAuthentication yes**.

4. Restart the **SSH** service on **remote.idm.example.com** so that the above changes take effect immediately:

   > $ **systemctl try-restart sshd.service**

Additional resources

- Logging in to the web console with smart cards

- Constrained delegation in Identity Management

29.7. USING ANSIBLE TO CONFIGURE A WEB CONSOLE TO ALLOW A USER AUTHENTICATED WITH A SMART CARD TO SSH TO A REMOTE HOST WITHOUT BEING ASKED TO AUTHENTICATE AGAIN

After you have logged in to a user account on the RHEL web console, as an Identity Management (IdM) system administrator you might need to connect to remote machines by using the **SSH** protocol. You can use the constrained delegation feature to use **SSH** without being asked to authenticate again.

Follow this procedure to use the **servicedelegationrule** and **servicedelegationtarget ansible-freeipa** modules to configure a web console to use constrained delegation. In the example below, the web console session runs on the **myhost.idm.example.com** host and it is being configured to access the **remote.idm.example.com** host by using **SSH** on behalf of the authenticated user.

Prerequisites

- The IdM **admin** password.

- **root** access to **remote.idm.example.com**.

- The web console service is present in IdM.

- The **remote.idm.example.com** host is present in IdM.

- The web console has created an **S4U2Proxy** Kerberos ticket in the user session. To verify that this is the case, log in to the web console as an IdM user, open the **Terminal** page, and enter:

  ```
  $ klist
  Ticket cache: FILE:/run/user/1894000001/cockpit-session-3692.ccache
  Default principal: user@IDM.EXAMPLE.COM

  Valid starting     Expires          Service principal
  07/30/21 09:19:06 07/31/21 09:19:06
  HTTP/myhost.idm.example.com@IDM.EXAMPLE.COM
  07/30/21 09:19:06  07/31/21 09:19:06  krbtgt/IDM.EXAMPLE.COM@IDM.EXAMPLE.COM
          for client HTTP/myhost.idm.example.com@IDM.EXAMPLE.COM
  ```

- You have configured your Ansible control node to meet the following requirements:

  - You are using Ansible version 2.14 or later.

  - You have installed the **ansible-freeipa** package on the Ansible controller.

  - The example assumes that in the *~/MyPlaybooks/* directory, you have created an Ansible inventory file with the fully-qualified domain name (FQDN) of the IdM server.

  - The example assumes that the **secret.yml** Ansible vault stores your **ipaadmin_password**.

Procedure

1. Navigate to your *~/MyPlaybooks/* directory:

   ```
   $ cd ~/MyPlaybooks/
   ```

2. Create a **web-console-smart-card-ssh.yml** playbook with the following content:

   a. Create a task that ensures the presence of a delegation target:

a. Create a task that ensures the presence of a delegation target.

```
---
- name: Playbook to create a constrained delegation target
  hosts: ipaserver

  vars_files:
  - /home/user_name/MyPlaybooks/secret.yml
  tasks:
  - name: Ensure servicedelegationtarget web-console-delegation-target is present
    ipaservicedelegationtarget:
      ipaadmin_password: "{{ ipaadmin_password }}"
      name: web-console-delegation-target
```

b. Add a task that adds the target host to the delegation target:

```
  - name: Ensure servicedelegationtarget web-console-delegation-target member
principal host/remote.idm.example.com@IDM.EXAMPLE.COM is present
    ipaservicedelegationtarget:
      ipaadmin_password: "{{ ipaadmin_password }}"
      name: web-console-delegation-target
      principal: host/remote.idm.example.com@IDM.EXAMPLE.COM
      action: member
```

c. Add a task that ensures the presence of a delegation rule:

```
  - name: Ensure servicedelegationrule delegation-rule is present
    ipaservicedelegationrule:
      ipaadmin_password: "{{ ipaadmin_password }}"
      name: web-console-delegation-rule
```

d. Add a task that ensures that the Kerberos principal of the web console client service is a member of the constrained delegation rule:

```
  - name: Ensure the Kerberos principal of the web console client service is added to the
servicedelegationrule web-console-delegation-rule
    ipaservicedelegationrule:
      ipaadmin_password: "{{ ipaadmin_password }}"
      name: web-console-delegation-rule
      principal: HTTP/myhost.idm.example.com
      action: member
```

e. Add a task that ensures that the constrained delegation rule is associated with the web–console–delegation–target delegation target:

```
  - name: Ensure a constrained delegation rule is associated with a specific delegation
target
    ipaservicedelegationrule:
      ipaadmin_password: "{{ ipaadmin_password }}"
      name: web-console-delegation-rule
      target: web-console-delegation-target
      action: member
```

3. Save the file.

4. Run the Ansible playbook. Specify the playbook file, the file storing the password protecting the **secret.yml** file, and the inventory file:

> $ **ansible-playbook --vault-password-file=password_file -v -i inventory web-console-smart-card-ssh.yml**

5. Enable Kerberos authentication on **remote.idm.example.com**:

   a. **SSH** to **remote.idm.example.com** as **root**.

   b. Open the **/etc/ssh/sshd_config** file for editing.

   c. Enable **GSSAPIAuthentication** by uncommenting the **GSSAPIAuthentication no** line and replacing it with **GSSAPIAuthentication yes**.

Additional resources

- Logging in to the web console with smart cards

- Constrained delegation in Identity Management

- **README-servicedelegationrule.md** and **README-servicedelegationtarget.md** in the **/usr/share/doc/ansible-freeipa/** directory

- Sample playbooks in the **/usr/share/doc/ansible-freeipa/playbooks/servicedelegationtarget** and **/usr/share/doc/ansible-freeipa/playbooks/servicedelegationrule** directories

CHAPTER 30. CONFIGURING SINGLE SIGN-ON FOR THE RHEL 8 WEB CONSOLE IN THE IDM DOMAIN

Learn how to use Single Sign-on (SSO) authentication provided by Identity Management (IdM) in the RHEL 8 web console.

Advantages:

- IdM domain administrators can use the RHEL 8 web console to manage local machines.

- Users with a Kerberos ticket in the IdM domain do not need to provide login credentials to access the web console.

- All hosts known to the IdM domain are accessible via SSH from the local instance of the RHEL 8 web console.

- Certificate configuration is not necessary. The console's web server automatically switches to a certificate issued by the IdM certificate authority and accepted by browsers.

This chapter covers the following steps to configure SSO for logging into the RHEL web console:

1. Add machines to the IdM domain using the RHEL 8 web console.
   For details, see Joining a RHEL 8 system to an IdM domain using the web console .

2. If you want to use Kerberos for authentication, you need to obtain a Kerberos ticket on your machine.
   For details, see Logging in to the web console using Kerberos authentication .

3. Allow administrators on the IdM server to run any command on any host.
   For details, see Enabling admin sudo access to domain administrators on the IdM server .

Prerequisites

- The RHEL web console installed on RHEL 8 systems.
  For details, see Installing the web console .

- IdM client installed on systems with the RHEL web console.
  For details, see IdM client installation .

30.1. JOINING A RHEL 8 SYSTEM TO AN IDM DOMAIN USING THE WEB CONSOLE

You can use the web console to join the Red Hat Enterprise Linux 8 system to the Identity Management (IdM) domain.

Prerequisites

- The IdM domain is running and reachable from the client you want to join.

- You have the IdM domain administrator credentials.

Procedure

1. Log into the RHEL web console.

For details, see Logging in to the web console .

2. In the **Configuration** field of the **Overview** tab click **Join Domain**.

3. In the **Join a Domain** dialog box, enter the host name of the IdM server in the **Domain Address** field.

4. In the **Domain administrator name** field, enter the user name of the IdM administration account.

5. In the **Domain administrator password**, add a password.

6. Click **Join**.

Verification steps

1. If the RHEL 8 web console did not display an error, the system has been joined to the IdM domain and you can see the domain name in the **System** screen.

2. To verify that the user is a member of the domain, click the Terminal page and type the **id** command:

```
$ id
euid=548800004(example_user) gid=548800004(example_user)
groups=548800004(example_user) context=unconfined_u:unconfined_r:unconfined_t:s0-
s0:c0.c1023
```

Additional resources

- Planning Identity Management

- Installing Identity Management

- Managing IdM users, groups, hosts, and access control rules

30.2. LOGGING IN TO THE WEB CONSOLE USING KERBEROS AUTHENTICATION

The following procedure describes steps on how to set up the RHEL 8 system to use Kerberos authentication.

IMPORTANT

With SSO you usually do not have any administrative privileges in the web console. This only works if you configured passwordless sudo. The web console does not interactively ask for a sudo password.

Prerequisites

- IdM domain running and reachable in your company environment.
  For details, see Joining a RHEL 8 system to an IdM domain using the web console .

- Enable the **cockpit.socket** service on remote systems to which you want to connect and manage them with the RHEL web console.

For details, see Installing the web console .

- If the system does not use a Kerberos ticket managed by the SSSD client, try to request the ticket with the **kinit** utility manually.

Procedure

Log in to the RHEL web console with the following address: **https://dns_name:9090**.

At this point, you are successfully connected to the RHEL web console and you can start with configuration.
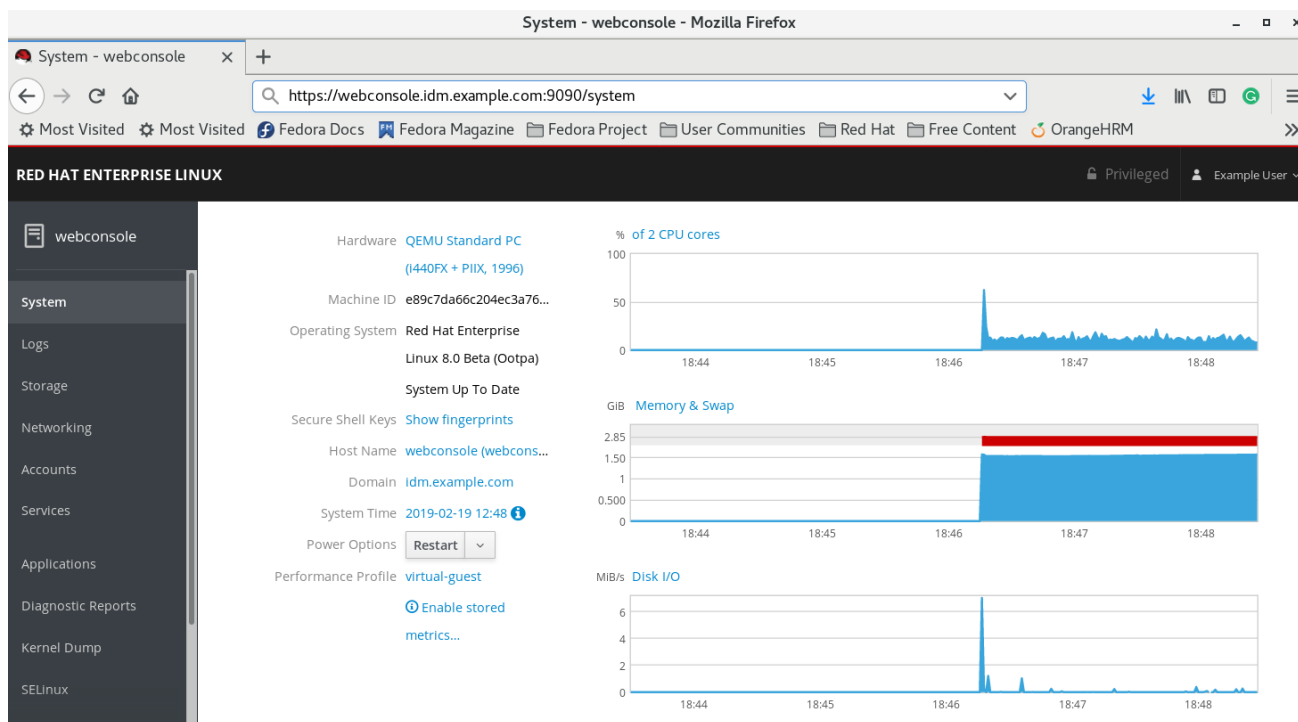
30.3. ENABLING ADMIN SUDO ACCESS TO DOMAIN ADMINISTRATORS ON THE IDM SERVER

The following procedure describes steps on how to allow domain administrators to run any command on any host in the Identity Management (IdM) domain.

To accomplish this, enable sudo access to the **admins** user group created automatically during the IdM server installation.

All users added to the **admins** group will have sudo access if you run **ipa-advise** script on the group.

Prerequisites

- The server runs IdM 4.7.1 or later.

Procedure

1. Connect to the IdM server.

2. Run the ipa-advise script:

   $ ipa-advise enable-admins-sudo | sh -ex

If the console did not display an error, the **admins** group have admin permissions on all machines in the IdM domain.

CHAPTER 31. CONFIGURING SMART CARD AUTHENTICATION WITH THE WEB CONSOLE FOR CENTRALLY MANAGED USERS

Configure smart card authentication in the RHEL web console for users who are centrally managed by:

- Identity Management

- Active Directory which is connected in the cross-forest trust with Identity Management

> **IMPORTANT**
>
> Smart card authentication does not elevate administrative privileges yet and the web console opens in the web browser in the read-only mode.
>
> You can run administrative commands in the built-in terminal with `sudo`.

Prerequisites

- The system for which you want to use the smart card authentication must be a member of an Active Directory or Identity Management domain.
  For details about joining the RHEL 8 system into a domain using the web console, see Joining a RHEL 8 system to an IdM domain using the web console.

- The certificate used for the smart card authentication must be associated with a particular user in Identity Management or Active Directory.
  For more details about associating a certificate with the user in Identity Management, see Adding a certificate to a user entry in the IdM Web UI   or Adding a certificate to a user entry in the IdM CLI.

31.1. SMART CARD AUTHENTICATION FOR CENTRALLY MANAGED USERS

A smart card is a physical device, which can provide personal authentication using certificates stored on the card. Personal authentication means that you can use smart cards in the same way as user passwords.

You can store user credentials on the smart card in the form of a private key and a certificate. Special software and hardware is used to access them. You insert the smart card into a reader or a USB socket and supply the PIN code for the smart card instead of providing your password.

Identity Management (IdM) supports smart card authentication with:

- User certificates issued by the IdM certificate authority. For more details, see Configuring Identity Management for smart card authentication.

- User certificates issued by the Active Directory Certificate Service (ADCS) certificate authority. For more details, see Configuring certificates issued by ADCS for smart card authentication in IdM.
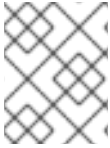
NOTE

If you want to start using smart card authentication, see the hardware requirements:
Smart Card support in RHEL8+ .

31.2. INSTALLING TOOLS FOR MANAGING AND USING SMART CARDS

Prerequisites

- The **gnutls-utils** package is installed.

- The **opensc** package is installed.

- The **pcscd** service is running.

Before you can configure your smart card, you must install the corresponding tools, which can generate
certificates and start the **pscd** service.

Procedure

1. Install the **opensc** and **gnutls-utils** packages:

   ```
   # {PackageManagerCommand} -y install opensc gnutls-utils
   ```

2. Start the **pcscd** service.

   ```
   # systemctl start pcscd
   ```

Verification steps

- Verify that the **pcscd** service is up and running

  ```
  # systemctl status pcscd
  ```

31.3. PREPARING YOUR SMART CARD AND UPLOADING YOUR CERTIFICATES AND KEYS TO YOUR SMART CARD

Follow this procedure to configure your smart card with the **pkcs15-init** tool, which helps you to
configure:

- Erasing your smart card

- Setting new PINs and optional PIN Unblocking Keys (PUKs)

- Creating a new slot on the smart card

- Storing the certificate, private key, and public key in the slot

- If required, locking the smart card settings as certain smart cards require this type of finalization

NOTE

The **pkcs15-init** tool may not work with all smart cards. You must use the tools that work with the smart card you are using.

Prerequisites

- The **opensc** package, which includes the **pkcs15-init** tool, is installed.
  For more details, see Installing tools for managing and using smart cards .

- The card is inserted in the reader and connected to the computer.

- You have a private key, a public key, and a certificate to store on the smart card. In this procedure, **testuser.key**, **testuserpublic.key**, and **testuser.crt** are the names used for the private key, public key, and the certificate.

- You have your current smart card user PIN and Security Officer PIN (SO-PIN).

Procedure

1. Erase your smart card and authenticate yourself with your PIN:

   > **$ pkcs15-init --erase-card --use-default-transport-keys**
   > Using reader with a card: *Reader name*
   > PIN [Security Officer PIN] required.
   > Please enter PIN [Security Officer PIN]:

   The card has been erased.

2. Initialize your smart card, set your user PIN and PUK, and your Security Officer PIN and PUK:

   > $ pkcs15-init --create-pkcs15 --use-default-transport-keys \ --pin *963214* --puk *321478* --so-pin *65498714* --so-puk *784123*
   > Using reader with a card: *Reader name*

   The **pcks15-init** tool creates a new slot on the smart card.

3. Set a label and the authentication ID for the slot:

   > $ pkcs15-init --store-pin --label testuser \ --auth-id 01 --so-pin *65498714* --pin *963214* --puk *321478*
   > Using reader with a card: *Reader name*

   The label is set to a human-readable value, in this case, **testuser**. The **auth-id** must be two hexadecimal values, in this case it is set to **01**.

4. Store and label the private key in the new slot on the smart card:

   > $ pkcs15-init --store-private-key testuser.key --label testuser_key \ --auth-id 01 --id 01 --pin *963214*
   > Using reader with a card: *Reader name*

NOTE

The value you specify for **--id** must be the same when storing your private key and storing your certificate in the next step. Specifying your own value for **--id** is recommended as otherwise a more complicated value is calculated by the tool.

5. Store and label the certificate in the new slot on the smart card:

   ```
   $ pkcs15-init --store-certificate testuser.crt --label testuser_crt \ --auth-id 01 --id 01 --format pem --pin 963214
   Using reader with a card: Reader name
   ```

6. Optional: Store and label the public key in the new slot on the smart card:

   ```
   $ pkcs15-init --store-public-key testuserpublic.key --label testuserpublic_key --auth-id 01 --id 01 --pin 963214
   Using reader with a card: Reader name
   ```

NOTE

If the public key corresponds to a private key or certificate, specify the same ID as the ID of the private key or certificate.

7. Optional: Certain smart cards require you to finalize the card by locking the settings:

   ```
   $ pkcs15-init -F
   ```

At this stage, your smart card includes the certificate, private key, and public key in the newly created slot. You have also created your user PIN and PUK and the Security Officer PIN and PUK.

31.4. ENABLING SMART CARD AUTHENTICATION FOR THE WEB CONSOLE

To be able to use smart card authentication in the web console, enable smart card authentication in the **cockpit.conf** file.

Additionally, you can disable password authentication in the same file.

Prerequisites

- The RHEL web console has been installed.
  For more details, see Installing the web console.

Procedure

1. Log in to the RHEL web console with administrator privileges.
   For details, see Logging in to the web console .

2. Click **Terminal**.

3. In the **/etc/cockpit/cockpit.conf**, set the **ClientCertAuthentication** to **yes**:

```
[WebService]
ClientCertAuthentication = yes
```

4. Optional: Disable password based authentication in **cockpit.conf** with:

```
[Basic]
action = none
```

This configuration disables password authentication and you must always use the smart card.

5. Restart the web console to ensure that the **cockpit.service** accepts the change:

```
# systemctl restart cockpit
```

31.5. LOGGING IN TO THE WEB CONSOLE WITH SMART CARDS

You can use smart cards to log in to the web console.

Prerequisites

- A valid certificate stored in your smart card that is associated to a user account created in a Active Directory or Identity Management domain.

- PIN to unlock the smart card.

- The smart card has been put into the reader.

Procedure

1. Open your web browser and add the web console's address in the address bar.
   The browser asks you to add the PIN protecting the certificate stored on the smart card.

2. In the **Password Required** dialog box, enter PIN and click **OK**.

3. In the **User Identification Request** dialog box, select the certificate stored in the smart card.

4. Select **Remember this decision**.
   The system does not open this window next time.
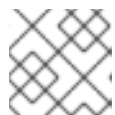
> **NOTE**
>
> This step does not apply to Google Chrome users.

5. Click **OK**.

You are now connected and the web console displays its content.

31.6. LIMITING USER SESSIONS AND MEMORY TO PREVENT A DOS ATTACK

A certificate authentication is protected by separating and isolating instances of the **cockpit-ws** web server against attackers who wants to impersonate another user. However, this introduces a potential denial of service (DoS) attack: A remote attacker could create a large number of certificates and send a

large number of HTTPS requests to **cockpit-ws** each using a different certificate.

To prevent this DoS, the collective resources of these web server instances are limited. By default, limits to the number of connections and to memory usage are set to 200 threads and a 75% (soft) / 90% (hard) memory limit.

The following procedure describes resource protection by limiting the number of connections and memory.

Procedure

1. In the terminal, open the **system-cockpithttps.slice** configuration file:

   ```
   # systemctl edit system-cockpithttps.slice
   ```

2. Limit the **TasksMax** to *100* and **CPUQuota** to *30%*:

   ```
   [Slice]
   # change existing value
   TasksMax=100
   # add new restriction
   CPUQuota=30%
   ```

3. To apply the changes, restart the system:

   ```
   # systemctl daemon-reload
   # systemctl stop cockpit
   ```

Now, the new memory and user session limits protect the **cockpit-ws** web server from DoS attacks.

31.7. ADDITIONAL RESOURCES

- Configuring Identity Management for smart card authentication .

- Configuring certificates issued by ADCS for smart card authentication in IdM .

- Configuring and importing local certificates to a smart card .

CHAPTER 32. MANAGING CONTAINER IMAGES BY USING THE RHEL WEB CONSOLE

You can use the RHEL web console web-based interface to pull, prune, or delete your container images.

32.1. PULLING CONTAINER IMAGES IN THE WEB CONSOLE

You can download container images to your local system and use them to create your containers.

Prerequisites

- The web console is installed and accessible. For more information, see Installing the web console and Logging in to the web console .

- The **cockpit-podman** add-on is installed:

  ```
  # yum install cockpit-podman
  ```

Procedure

1. Click **Podman containers** in the main menu.

2. In the **Images** table, click the overflow menu in the upper-right corner and select **Download new image**.

3. The **Search for an image** dialog box appears.

4. In the **Search for** field, enter the name of the image or specify its description.

5. In the **in** drop-down list, select the registry from which you want to pull the image.

6. Optional: In the **Tag** field, enter the tag of the image.

7. Click **Download**.

Verification

- Click **Podman containers** in the main menu. You can see the newly downloaded image in the **Images** table.

> **NOTE**
>
> You can create a container from the downloaded image by clicking the **Create container** in the **Images** table. To create the container, follow steps 3-8 in Creating containers in the web console.

32.2. PRUNING CONTAINER IMAGES IN THE WEB CONSOLE

You can remove all unused images that do not have any containers based on it.

Prerequisites

- At least one container image is pulled.

- The web console is installed and accessible. For more information, see Installing the web console and Logging in to the web console .

- The **cockpit-podman** add-on is installed:

  ```
  # yum install cockpit-podman
  ```

Procedure

1. Click **Podman containers** in the main menu.

2. In the **Images** table, click the overflow menu in the upper-right corner and select **Prune unused images**.

3. The pop-up window with the list of images appears. Click **Prune** to confirm your choice.

Verification

- Click **Podman containers** in the main menu. The deleted images should not be listed in the **Images** table.

32.3. DELETING CONTAINER IMAGES IN THE WEB CONSOLE

You can delete a previously pulled container image using the web console.

Prerequisites

- At least one container image is pulled.

- The web console is installed and accessible. For more information, see Installing the web console and Logging in to the web console .

- The **cockpit-podman** add-on is installed:

  ```
  # yum install cockpit-podman
  ```

Procedure

1. Click **Podman containers** in the main menu.

2. In the **Images** table, select the image you want to delete and click the overflow menu and select **Delete**.

3. The window appears. Click **Delete tagged images** to confirm your choice.

Verification

- Click the **Podman containers** in the main menu. The deleted container should not be listed in the **Images** table.

CHAPTER 33. MANAGING CONTAINERS BY USING THE RHEL WEB CONSOLE

You can use the Red Hat Enterprise Linux web console to manage your containers and pods. With the web console, you can create containers as a non-root or root user.

- As a *root* user, you can create system containers with extra privileges and options.

- As a *non-root* user, you have two options:

    - To only create user containers, you can use the web console in its default mode - **Limited access**.

    - To create both user and system containers, click **Administrative access** in the top panel of the web console page.

For details about differences between root and rootless containers, see Special considerations for rootless containers.

33.1. CREATING CONTAINERS IN THE WEB CONSOLE

You can create a container and add port mappings, volumes, environment variables, health checks, and so on.

Prerequisites

- The web console is installed and accessible. For more information, see Installing the web console and Logging in to the web console .

- The **cockpit-podman** add-on is installed:

    ```
    # yum install cockpit-podman
    ```

Procedure

1. Click **Podman containers** in the main menu.

2. Click **Create container**.

3. In the **Name** field, enter the name of your container.

4. Provide desired info in the **Details** tab.

    - *Available only with the administrative access*: Select the Owner of the container: System or User.

    - In the **Image** drop down list select or search the container image in selected registries.

        - Optional: Check the **Pull latest image** checkbox to pull the latest container image.

    - The **Command** field specifies the command. You can change the default command if you need.

        - Optional: Check the **With terminal** checkbox to run your container with a terminal.

- The **Memory limit** field specifies the memory limit for the container. To change the default memory limit, check the checkbox and specify the limit.

- *Available only for system containers*: In the **CPU shares field**, specify the relative amount of CPU time. Default value is 1024. Check the checkbox to modify the default value.

- *Available only for system containers*: In the **Restart policy** drop down menu, select one of the following options:

  - **No** (default value): No action.

  - **On Failure**: Restarts a container on failure.

  - **Always**: Restarts a container when exits or after rebooting the system.

5. Provide the required information in the **Integration** tab.

   - Click **Add port mapping** to add port mapping between the container and host system.

     - Enter the *IP address*, *Host port*, *Container port* and *Protocol*.

   - Click **Add volume** to add volume.

     - Enter the *host path*, *Container path*. You can check the **Writable** option checkbox to create a writable volume. In the SELinux drop down list, select one of the following options: **No Label**, **Shared** or **Private**.

   - Click **Add variable** to add environment variable.

     - Enter the *Key* and *Value*.

6. Provide the required information in the **Health check** tab.

   - In the **Command** fields, enter the 'healthcheck' command.

   - Specify the healthcheck options:

     - **Interval** (default is 30 seconds)

     - **Timeout** (default is 30 seconds)

     - **Start period**

     - **Retries** (default is 3)

     - When unhealthy: Select one of the following options:

       - **No action** (default): Take no action.

       - **Restart**: Restart the container.

       - **Stop**: Stop the container.

       - **Force stop**: Force stops the container, it does not wait for the container to exit.

7. Click **Create and run** to create and run the container.

> **NOTE**
>
> You can click **Create** to only create the container.

Verification

- Click **Podman containers** in the main menu. You can see the newly created container in the **Containers** table.

33.2. INSPECTING CONTAINERS IN THE WEB CONSOLE

You can display detailed information about a container in the web console.

Prerequisites

- The container was created.

- The web console is installed and accessible. For more information, see Installing the web console and Logging in to the web console .

- The **cockpit-podman** add-on is installed:

  ```
  # yum install cockpit-podman
  ```

Procedure

1. Click **Podman containers** in the main menu.

2. Click the **>** arrow icon to see details of the container.

   - In the **Details** tab, you can see container ID, Image, Command, Created (timestamp when the container was created), and its State.

     - *Available only for system containers*: You can also see IP address, MAC address, and Gateway address.

   - In the **Integration** tab, you can see environment variables, port mappings, and volumes.

   - In the **Log** tab, you can see container logs.

   - In the **Console** tab, you can interact with the container using the command line.

33.3. CHANGING THE STATE OF CONTAINERS IN THE WEB CONSOLE

In the Red Hat Enterprise Linux web console, you can start, stop, restart, pause, and rename containers on the system.

Prerequisites

- The container was created.

- The web console is installed and accessible. For more information, see Installing the web console and Logging in to the web console .

- The **cockpit-podman** add-on is installed:

> # yum install cockpit-podman

Procedure

1. Click **Podman containers** in the main menu.

2. In the **Containers** table, select the container you want to modify and click the overflow menu and select the action you want to perform:

   - Start

   - Stop

   - Force stop

   - Restart

   - Force restart

   - Pause

   - Rename

33.4. COMMITTING CONTAINERS IN THE WEB CONSOLE

You can create a new image based on the current state of the container.

Prerequisites

- The container was created.

- The web console is installed and accessible. For more information, see Installing the web console and Logging in to the web console .

- The **cockpit-podman** add-on is installed:

  > # yum install cockpit-podman

Procedure

1. Click **Podman containers** in the main menu.

2. In the **Containers** table, select the container you want to modify and click the overflow menu and select **Commit**.

3. In the **Commit container** form, add the following details:

   - In the **New image name** field, enter the image name.

   - Optional: In the **Tag** field, enter the tag.

   - Optional: In the **Author** field, enter your name.

   - Optional: In the **Command** field, change command if you need.

- Optional: Check the **Options** you need:

    - Pause container when creating image: The container and its processes are paused while the image is committed.

    - Use legacy Docker format: if you do not use the Docker image format, the OCI format is used.

4. Click **Commit**.

Verification

- Click the **Podman containers** in the main menu. You can see the newly created image in the **Images** table.

33.5. CREATING A CONTAINER CHECKPOINT IN THE WEB CONSOLE

Using the web console, you can set a checkpoint on a running container or an individual application and store its state to disk.

> **NOTE**
>
> Creating a checkpoint is available only for system containers.

Prerequisites

- The container is running.

- The web console is installed and accessible. For more information, see Installing the web console and Logging in to the web console .

- The **cockpit-podman** add-on is installed:

    ```
    # yum install cockpit-podman
    ```

Procedure

1. Click **Podman containers** in the main menu.

2. In the **Containers** table, select the container you want to modify and click the overflow icon menu and select **Checkpoint**.

3. Optional: In the **Checkpoint container** form, check the options you need:

    - Keep all temporary checkpoint files: keep all temporary log and statistics files created by CRIU during checkpointing. These files are not deleted if checkpointing fails for further debugging.

    - Leave running after writing checkpoint to disk: leave the container running after checkpointing instead of stopping it.

    - Support preserving established TCP connections

4. Click **Checkpoint**.

Verification

- Click the **Podman containers** in the main menu. Select the container you checkpointed, click the overflow menu icon and verify that there is a **Restore** option.

33.6. RESTORING A CONTAINER CHECKPOINT IN THE WEB CONSOLE

You can use data saved to restore the container after a reboot at the same point in time it was checkpointed.

> **NOTE**
>
> Creating a checkpoint is available only for system containers.

Prerequisites

- The container was checkpointed.

- The web console is installed and accessible. For more information, see Installing the web console and Logging in to the web console .

- The **cockpit-podman** add-on is installed:

  ```
  # yum install cockpit-podman
  ```

Procedure

1. Click **Podman containers** in the main menu.

2. In the **Containers** table, select the container you want to modify and click the overflow menu and select **Restore**.

3. Optional: In the **Restore container** form, check the options you need:

   - **Keep all temporary checkpoint files**: Keep all temporary log and statistics files created by CRIU during checkpointing. These files are not deleted if checkpointing fails for further debugging.

   - **Restore with established TCP connections**

   - **Ignore IP address if set statically**: If the container was started with IP address the restored container also tries to use that IP address and restore fails if that IP address is already in use. This option is applicable if you added port mapping in the Integration tab when you create the container.

   - **Ignore MAC address if set statically**: If the container was started with MAC address the restored container also tries to use that MAC address and restore fails if that MAC address is already in use.

4. Click **Restore**.

Verification

- Click the **Podman containers** in the main menu. You can see that the restored container in the **Containers** table is running.

33.7. DELETING CONTAINERS IN THE WEB CONSOLE

You can delete an existing container using the web console.

Prerequisites

- The container exists on your system.

- The web console is installed and accessible. For more information, see Installing the web console and Logging in to the web console .

- The **cockpit-podman** add-on is installed:

  ```
  # yum install cockpit-podman
  ```

Procedure

1. Click **Podman containers** in the main menu.

2. In the **Containers** table, select the container you want to delete and click the overflow menu and select **Delete**.

3. The pop-up window appears. Click **Delete** to confirm your choice.

Verification

- Click the **Podman containers** in the main menu. The deleted container should not be listed in the **Containers** table.

33.8. CREATING PODS IN THE WEB CONSOLE

You can create pods in the RHEL web console interface.

Prerequisites

- The web console is installed and accessible. For more information, see Installing the web console and Logging in to the web console .

- The **cockpit-podman** add-on is installed:

  ```
  # yum install cockpit-podman
  ```

Procedure

1. Click **Podman containers** in the main menu.

2. Click **Create pod**.

3. Provide desired information in the **Create pod** form:

   - *Available only with the administrative access*: Select the Owner of the container: System or User.

   - In the **Name** field, enter the name of your container.

- Click **Add port mapping** to add port mapping between container and host system.

  - Enter the IP address, Host port, Container port and Protocol.

- Click **Add volume** to add volume.

  - Enter the host path, Container path. You can check the Writable checkbox to create a writable volume. In the SELinux drop down list, select one of the following options: No Label, Shared or Private.

4. Click **Create**.

Verification

- Click **Podman containers** in the main menu. You can see the newly created pod in the **Containers** table.

33.9. CREATING CONTAINERS IN THE POD IN THE WEB CONSOLE

You can create a container in a pod.

Prerequisites

- The web console is installed and accessible. For more information, see Installing the web console and Logging in to the web console .

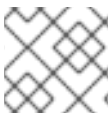- The **cockpit-podman** add-on is installed:

  ```
  # yum install cockpit-podman
  ```

Procedure

1. Click **Podman containers** in the main menu.

2. Click **Create container in pod**.

3. In the **Name** field, enter the name of your container.

4. Provide the required information in the **Details** tab.

   - *Available only with the administrative access*: Select the Owner of the container: System or User.

   - In the **Image** drop down list select or search the container image in selected registries.

     - Optional: Check the **Pull latest image** checkbox to pull the latest container image.

   - The **Command** field specifies the command. You can change the default command if you need.

     - Optional: Check the **With terminal** checkbox to run your container with a terminal.

   - The **Memory limit** field specifies the memory limit for the container. To change the default memory limit, check the checkbox and specify the limit.

- *Available only for system containers*: In the **CPU shares field**, specify the relative amount of CPU time. Default value is 1024. Check the checkbox to modify the default value.

- *Available only for system containers*: In the **Restart policy** drop down menu, select one of the following options:

  - **No** (default value): No action.

  - **On Failure**: Restarts a container on failure.

  - **Always**: Restarts container when exits or after system boot.

5. Provide the required information in the **Integration** tab.

   - Click **Add port mapping** to add port mapping between the container and host system.

     - Enter the *IP address*, *Host port*, *Container port* and *Protocol*.

   - Click **Add volume** to add volume.

     - Enter the *host path*, *Container path*. You can check the **Writable** option checkbox to create a writable volume. In the SELinux drop down list, select one of the following options: **No Label**, **Shared**, or **Private**.

   - Click **Add variable** to add environment variable.

     - Enter the *Key* and *Value*.

6. Provide the required information in the **Health check** tab.

   - In the **Command** fields, enter the healthcheck command.

   - Specify the healthcheck options:

     - **Interval** (default is 30 seconds)

     - **Timeout** (default is 30 seconds)

     - **Start period**

     - **Retries** (default is 3)

     - When unhealthy: Select one of the following options:

       - **No action** (default): Take no action.

       - **Restart**: Restart the container.

       - **Stop**: Stop the container.

       - **Force stop**: Force stops the container, it does not wait for the container to exit.

> **NOTE**
>
> The owner of the container is the same as the owner of the pod.

> **NOTE**
>
> In the pod, you can inspect containers, change the status of containers, commit containers, or delete containers.

Verification

- Click **Podman containers** in the main menu. You can see the newly created container in the pod under the **Containers** table.

33.10. CHANGING THE STATE OF PODS IN THE WEB CONSOLE

You can change the status of the pod.

Prerequisites

- The pod was created.

- The web console is installed and accessible. For more information, see Installing the web console and Logging in to the web console .

- The **cockpit-podman** add-on is installed:

  ```
  # yum install cockpit-podman
  ```

Procedure

1. Click **Podman containers** in the main menu.

2. In the **Containers** table, select the pod you want to modify and click the overflow menu and select the action you want to perform:

   - Start

   - Stop

   - Force stop

   - Restart

   - Force restart

   - Pause

33.11. DELETING PODS IN THE WEB CONSOLE

You can delete an existing pod using the web console.

Prerequisites

- The pod exists on your system.

- The web console is installed and accessible. For more information, see Installing the web console and Logging in to the web console .

- The **cockpit-podman** add-on is installed:

  ```
  # yum install cockpit-podman
  ```

Procedure

1. Click **Podman containers** in the main menu.

2. In the **Containers** table, select the pod you want to delete and click the overflow menu and select **Delete**.

3. In the following pop-up window, click **Delete** to confirm your choice.

> **WARNING**
>
> You remove all containers in the pod.

Verification

- Click the **Podman containers** in the main menu. The deleted pod should not be listed in the **Containers** table.