Administering SUSE Linux Enterprise Server Using Cockpit

Cockpit is a Web-based graphical interface that allows administrators to manage SUSE Linux Enterprise Server servers from a centralized location. It eliminates the need for complex command-line operations by providing an intuitive dashboard to manage storage, networking, containers, system updates and user accounts, thereby streamlining system administration and improving operational efficiency.

Publication Date: 05 Nov 2026

Contents

1 About Cockpit 3

2 Cockpit installation 3

3 Accessing Cockpit 4

4 Configuring servers using Cockpit 10

5 Filtering Cockpit logs 12

6 Managing storage using Cockpit 15

7 Managing networking using Cockpit 22

8 Managing virtual machines using Cockpit 30

9 Working with containers 50

10 User administration using Cockpit 55

11 Managing services using Cockpit 58

12 Managing applications in Cockpit 59

13 Managing packages using Cockpit 60

14 Managing subscriptions using Cockpit 61

15 SELinux mode and policy 62

16 Legal Notice 66

A GNU Free Documentation License 67

1 About Cockpit

Cockpit is a Web-based graphical interface that enables you to manage most administration tasks from one place. You do not need to create credentials for Cockpit as, by default, Cockpit uses the same credentials that you use to log in to your server. Cockpit uses APIs that already exist on the system without adding a layer to the system.

Cockpit enables you to perform the following tasks:

- download container images and run containers
- manage the server storage
- inspect and change network settings
- manage user accounts
- view system logs
- inspect and interact with `systemd` services
- switch between SELinux modes
- use a terminal on a remote server in your Web browser

2 Cockpit installation

Cockpit can be installed during the system installation using Agama, or you can install it later from the running system. To verify if Cockpit is installed on your system, run:

```
> sudo zypper se -i cockpit
```

If Cockpit is not installed, proceed as described in *Section 2.1, "Installing Cockpit"*.

2.1 Installing Cockpit

If Cockpit is not present on your system, you can install it by following the procedure below:

1. Run the following command to install the Cockpit pattern:

   ```
   > sudo zypper in -t pattern cockpit
   ```

2. If the Cockpit instance is intended to serve as a primary one, you need to enable the Cockpit socket in `systemd` by running:

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

After running the command, the server exposes the default *9090* port and `systemd` starts the `cockpit-ws` service that listens on the 9090 port.

3. In case you have enabled the firewall, proceed as follows:

   a. Open the firewall for Cockpit

```
> sudo firewall-cmd --add-service=cockpit --permanent && sudo firewall-cmd —-
reload
```

   b. Install SELinux policies for Cockpit and restart `cockpit.socket` by running:

```
> sudo zypper in -f cockpit-selinux-policies && sudo systemctl restart
 cockpit.socket
```

4. Now you can access the Cockpit Web interface by opening the following address in your Web browser:

```
https://IP_ADDRESS_OF_MACHINE:9090
```

2.2 Cockpit plug-ins

Cockpit uses plug-ins to administer the system. The plug-ins are included in the installation pattern. However, depending on the technologies installed on your system, some plug-ins may not be visible to you. For example, if NFS is not present, the corresponding NFS panel is not visible.

3 Accessing Cockpit

Cockpit enables you to log in directly to each machine that can expose the 9090 port. This machine is sometimes referred to as the primary server. It is the primary server that runs the `cockpit-ws` through which connections to additional servers are established. By default, Cockpit listens on both HTTP and HTTPS connections. However, most of the HTTP connections are redirected to HTTPS, with exceptions like local host access.

If the port cannot be accessed on the particular machine, you can still use Cockpit to administer the machine by using it as a secondary server. For the procedure of adding a server as secondary, refer to *Procedure 2, "Adding a server as secondary"*.



Note: A limited number of secondary servers

The number of secondary servers that you can administer from one primary server is limited to 20. If you need to administer more servers, add other primary servers or use another tool for cluster administration.

3.1 TLS certificates

By default, Cockpit loads `.cert` or `.crt` certificates from the directory `/etc/cockpit/ws-certs.d`. The corresponding private key must be a separate file with the same file name but with the `.key` suffix. Make sure the key is not encrypted.

If no certificate is found in the directory, Cockpit generates a self-signed certificate (`0-self-signed.cert`) to establish a secure connection.

To check which certificate Cockpit uses, run the command:

```
> sudo /usr/libexec/cockpit-certificate-ensure --check
```

3.2 Authentication

You do not need separate credentials to log in to Cockpit. Use the same credentials that you use to log in to SUSE Linux Enterprise Server. However, on new installations, login using `root` is not allowed by default. Either enable `root` login with a password as described in *Section 3.2.2, "Enabling `root` to log in using a password"*, or create an unprivileged user to access Cockpit. On instances upgraded from a previous release, `root` login is still allowed. In all cases, we recommend enhancing the security by adding 2FA as described in *Section 3.2.1, "Enabling 2FA authorization"*.

At first, you log in to Cockpit with limited access. To perform administrative tasks, click *Limited access* in the upper-right menu and unlock the administrative mode by entering your password.

3.2.1 Enabling 2FA authorization

To set up 2FA on SUSE Linux Enterprise Server, you need an available TOTP application of your choice. Then run a command to configure the authorization. The following sections provide details on how to proceed with the configuration of 2FA and also give instructions in situations when your 2FA fails.

3.2.1.1 Applications providing TOTP 2FA

The following applications providing 2FA are supported on SUSE Linux Enterprise Server.

Using cloud storage

- PSONO (https://psono.com/) - available for Firefox, Chrome, Docker, iOS, Android
- Google Authenticator - available on Android, iOS and Wear OS
- Okta Verify (https://help.okta.com/en-us/content/topics/mobile/okta-verify-overview.htm) - available on Android, iOS, macOS and Windows

Using only local storage

- Yubico Authenticator (https://www.yubico.com/products/yubico-authenticator/) - with a hardware key
- KeePassXC (https://keepassxc.org/) - available on Linux desktops, Windows and macOS
- KeePassDX (https://www.keepassdx.com/) - available on Android
- FreeOTP Plus (https://github.com/helloworld1/FreeOTPPlus) - for Android
- FreeOTP (https://github.com/freeotp/freeotp-ios) - for iOS

3.2.1.2 Setting up 2FA

Each user can configure their own 2FA, or `root` can configure it for any regular user on the system. To set up 2FA for a user from a running system, proceed as follows.

1. Run the command:

```
> sudo /sbin/jeos-config otp
```

2. Scan the code to any TOTP application mentioned above.
3. Confirm the process by entering an OTP code.

3.2.1.3 Recovering access

Setting up 2FA is optional. However, once set, the second factor is mandatory to log in to Cockpit. If the second factor becomes unavailable, you can change it or disable it. Even without the second factor, you can still log in to the machine using SSH or directly from a console. After login, you can use the following two options:

Change the second factor

Run the command as `root`:

```
> sudo /sbin/jeos-config otp
```

Disable the 2FA

Remove the file `.pam_oath_usersfile` from the affected user's home directory.

3.2.2 Enabling `root` to log in using a password



Warning: `root` login with a password is not secure

We strongly discourage you from enabling `root` login with a password for security reasons.

In new SLES installations, `root` login using a password is disabled by default due to security reasons. To allow `root` login with password, proceed as follows:

1. Open the `/etc/cockpit/disallowed-users` file.
2. Remove `root` from the file.

3.3 Logging in to the primary server directly

Whenever you have direct network access to the 9090 port, you can directly log in to the server using your credentials. To do so, follow *Procedure 1, "Logging in to the primary server"*.



Note: No dedicated credentials for Cockpit needed

By default, access is controlled by a Cockpit-specific PAM stack located at `/usr/lib/pam.d/cockpit`. The default configuration allows logging in with the same user names and passwords that are used for any local account on the system.

PROCEDURE 1: LOGGING IN TO THE PRIMARY SERVER

1. Go to the Cockpit login page by opening the following address in a browser:

   ```
   https://IP_ADDRESS_OF_MACHINE:9090
   ```

2. Enter the credentials.

3.4 Logging in to secondary servers

If your machine does not have direct access to the 9090 port, you can use this machine as a secondary server. Ensure that Cockpit is installed on the machine.

There are two ways of logging in to a secondary server: you can log in to a secondary server directly or you can use the primary server.

3.4.1 Logging in to secondary servers directly

You can log in to any secondary server without logging in to the primary server first. This solution can be useful when you do not have credentials for the primary server. The primary server will be used as a bridge, and you will be connected to the secondary server using SSH.

To connect to the secondary server, proceed as follows:

1. Go to the Cockpit login page by opening the following address in a browser:

   ```
   https://IP_ADDRESS_OF_MACHINE:9090
   ```

2. Fill in the credentials for the secondary server.

3. Expand *Other options* on the login screen.

4. Fill in the IP address of the secondary server.

5. Proceed by clicking *Login*.

6. If you are trying to log in for the first time, you are asked to verify the fingerprint. After this, click *Accept and connect*.

3.4.2 Accessing secondary servers from the primary server

If you have credentials for the primary server, you can access secondary servers from the primary one. You must add the secondary servers first, as described in *Procedure 2, "Adding a server as secondary"*.

PROCEDURE 2: ADDING A SERVER AS SECONDARY

1. Log in to the primary server using the account with the *system administrator* role.

2. Click the *USERNAME*@*HOSTNAME* in the upper-left corner.

3. Click *Add new host*.

4. Fill in the host identification and optionally user name that will be used to log in to the server. You can assign a color to the machine. When the details are complete, click *Add*.

5. Verify the fingerprint on the server you want to add. If the fingerprint matches or if you have not set up the SSH connection, click *Trust and add host* to proceed.

6. Fill in the password and, if needed, check *Automatic login*. Cockpit will generate a new SSH key for the user, and next time you will be logged in automatically.

3.5 Logging in without a home directory

If a user does not have a `$HOME` directory—which frequently occurs during the initial login of network users managed via SSSD—the login attempt fails. To allow users without home directories to log into Cockpit, override the default PAM configuration to disable the look-up for the second-factor authentication file:

PROCEDURE 3: LOGGING IN WITHOUT A HOME DIRECTORY

Important

By following this procedure, you will not be able to use a second-factor authentication token via `pam_oath`.

1. Copy the system default Cockpit PAM configuration to the administrator override directory:

```
> sudo  cp /usr/lib/pam.d/cockpit /etc/pam.d/cockpit
```

2. Open `/etc/pam.d/cockpit` in an editor and locate the line referencing `pam_oath.so` at the end of the file.

3. Comment out or remove the line entirely:

```
#  auth [user_unknown=ignore success=ok] pam_oath.so usersfile=
${HOME}/.pam_oath_usersfile no_usersfile_okay window=20 digits=6
```

4. Save and close the file.

3.6 Switching to the administration mode

By default, a regular user can log in to Cockpit with limited access that does not enable the user to perform administration tasks like managing user accounts, updating the system, and so on.

To switch to administrative access, proceed as follows:

1. Click the *Limited access* button.

2. Fill in the `root` password.

3. Click *Authenticate* to confirm.

To turn off administrative mode, proceed as follows:

1. Click *Administrative access*.

2. To confirm, click *Limit access*.

4 Configuring servers using Cockpit

Using the Cockpit *Overview* part, you can perform changes to the default server configuration or the configuration you provided during the manual installation. In this part, you can change the host name or change the system date or time zone.

4.1 Viewing system health

To view the system health, click *Overview* on the left pane. You can view the following details in the *Health* section.

If a service fails, you can click the *1 service failed* to view the details of the service failure.

Click *View login history* to view your login history. You can modify the user credentials in the user page.

4.2 Viewing metrics and history

You can view the PCP metrics and usage history under the *Usage* section in the *Overview* page.

- Click *Overview* on the left pane.
  The following details appear in the *Metrics and History* page.
- *CPU*: The number of CPUs, the CPU load and a list of services that use the CPU.
- *Memory*: The RAM and Swap memory usage, and the related services.
  Cockpit reads memory usage data from the `virtio-balloon` driver inside the VM to estimate the memory VM uses. The specific value `currentMemory` shows how much memory the host has currently assigned to that VM. This number can change dynamically when the balloon inflates or deflates. If the VM does not use `virtio-balloon`, then Cockpit does not know the actual memory being used inside the VM. Hence the memory displayed could be different. The memory displayed is an approximation and not the actual memory usage inside the VM.
- *Disks*: A list of used disks and partitions. Click *View per-disk throughput* to display a table showing the data transfer rates for each disk.
- *Network* - Displays the list of networks.
- *Metrics history*: When present, Cockpit integrates with the Performance Co-Pilot (PCP) framework to collect and manage system performance metrics. These metrics are used to render historical data visualizations on the *Metrics and History* interface. Cockpit can also use the archive logging sub-system of PCP to retrieve and display time-shifted historical data. In environments where PCP is not installed, Cockpit falls back to its own metrics collection mechanism; however, this mode does not support data archival or retrieval of historical metrics.
  Archival of metrics depends on whether the `pmlogger.service` systemd unit is running or not. You can enable the *Collect metrics* option to enable and start the service.

4.3 Changing the server host name

To change the host name, proceed as follows:

PROCEDURE 4: CHANGING THE HOST NAME

1. Navigate to the *Overview* page.
2. In the *Configuration* section, click *edit* under *Hostname*.
3. Fill in the following:
   - *Pretty host name*—a user-defined free-form host name
   - *Real host name*—the name of the device in the network

4.4 Changing the system time or time zone

To change the system time or time zone, proceed as follows:

PROCEDURE 5: CHANGING SYSTEM TIME OR TIME ZONE

1. Navigate to the *Overview* page.
2. In the *Configuration* section, click the current time value under *System time*.
3. In the pop-up window you can change the following:
   - *Time zone*—the value set during the manual installation or, in case of raw images, set to UTC.
   - *Set time*—by default, NTP is used for time synchronization. You can set the time manually or, if you defined alternative NTP servers, you can use those NTP servers for time synchronization.
   - Click *Change*.

5 Filtering Cockpit logs

Click *Logs* on the left navigation pane and select the required filters. You can filter the logs according to the following criteria:

- *Time*. For details, refer to *Section 5.1, "Filtering according to time"*.
- *Priority*. For details, refer to *Section 5.2, "Filtering according to priority"*.

- *Identifier*. You can filter logs for a particular service, daemon or process. Available identifiers are parsed from the logs currently displayed according to the set filters.
- Free-form filters. For details, refer to *Section 5.3, "Logs filters"*.



Note: The filter criteria are combined

When changing any of the time, priority or identifier criteria, the others are still applied. For example, if you change the time criterion to *Last 24 hours*, the priority and identifier criteria remain the same.

5.1 Filtering according to time

To filter the logs according to a specific time, you can choose from the following values:

Current boot

Displays logs for the current boot only. The *Resume* button enables continuous refreshing of currently displayed logs.

Previous boot

Displays logs relevant to the previous boot.

Last 24 hours

Displays logs that were recorded within the last 24 hours.

Last 7 days

Displays logs that were recorded within the last 7 days.

5.2 Filtering according to priority

The standard **`syslog`** severity levels are used (sorted from most to least severe):

Only emergency

The system is unusable. This is a panic condition.

Alert and above

This log requires your immediate action.

Critical and above

Failures in the primary systems. You should correct the problem immediately.

Error and above

Not an urgent error, but should be handled within a specific time.

Warning and above

Not an error, but indicates that an error might occur if no action is taken.

Notice and above

Unusual events that are not errors. No immediate action is required.

Info and above

Normal operational messages that serve as a confirmation that the system works properly.

Debug and above

These messages are used just for debugging the system.

5.3 Logs filters

You can refine the logs view here according to the following criteria:

Since

Logs for the specified date or newer will be displayed. You can specify the time in the following way:

- using the absolute date in the format *YYYY-MM-DD*
- using any of the terms: `yesterday`, `today`, `tomorrow` and `now`
- using relative time by prefixing the value with `-` or `+` and specifying units. You can use the following units: `seconds` or `s`, `minutes` or `min`, `hours` or `h`, `days` or `d`, `weeks` or `w`, `months` or `m`, and `years` or `y`.

Until

Logs for the specified date or older will be displayed. You can specify the time in the following way:

- using the absolute date in the format *YYYY-MM-DD*
- using any of the terms: `yesterday`, `today`, `tomorrow` and `now`
- using relative time by prefixing the value with `-` or `+` and specifying units. You can use the following units: `seconds` or `s`, `minutes` or `min`, `hours` or `h`, `days` or `d`, `weeks` or `w`, `months` or `m`, and `years` or `y`.

Boot

Enter an integer: 0 means the current boot, -1 is for the previous boot, 1 for the first boot, 2 for the second, etc.

Unit

Specify the `systemd` unit for which you want to display logs. Use one of the formats:

- `_SYSTEMD_UNIT=`*`NAME`*`.service`
- `COREDUMP_UNIT=`*`NAME`*`.service`
- `UNIT=`*`NAME`*`.service`

Free-form search

Enter the string that you want to find in the log messages. You can also use PERL-compatible regular expressions (https://www.freedesktop.org/software/systemd/man/journalctl.html#-g). Alternatively, you can filter messages according to message log fields in the format *`FIELD=VALUE`*. For example, `CODE_LINE=349` displays logs with this value.

6 Managing storage using Cockpit

The *Storage* page enables you to monitor traffic on your drives, repartition your system, manage NFS mounts, view storage logs, and create RAIDs or LVM.

6.1 Monitoring data flow on disks

The graphs on the *Storage* page display reading and writing data flow to devices. Each device in the graph has a different color. Hover over the displayed data flow peak to identify the device name.

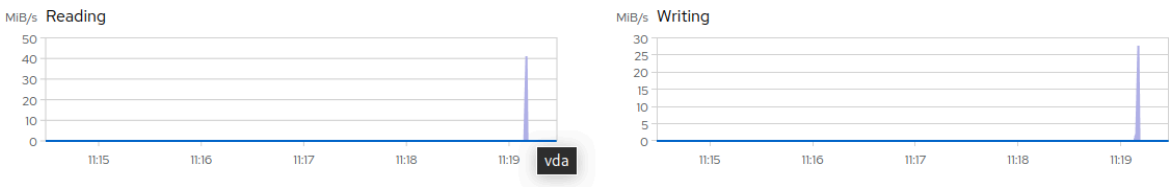


FIGURE 1: DATA FLOW VIEW

6.2 Managing file systems

The *Filesystems* view enables you to create a partition table and to format or mount file systems. You can sort the mounted partition according to *Name* or *Mount point*.

6.2.1 Formatting partitions using Cockpit

To format the partition, proceed as follows:

PROCEDURE 6: FORMATTING PARTITIONS

1. Navigate to the *Storage* page.
2. In the *Filesystem* view, click the partition you want to format.
3. Click *Format* next to the particular partition description to open the format window.
4. Enter a unique name for the partition.
5. In the *Mount point*, specify the directory to which the partition will be mounted. The *Mount point* field is mandatory.
6. In *Type*, select the file system type. Btrfs is mandatory for the `/` partition.
7. If needed, configure the encryption:

   Passphrase and Confirm

   Enter a passphrase to unlock the encrypted partition.

   Store passphrase

   The passphrase is stored in `/etc/luks-keys` and you are not asked for the passphrase on next boot.

   Encryption options

   You can pass a list of options described in supported encrypted options (https://www.man7.org/linux/man-pages/man5/crypttab.5.html#SUPPORTED_OPTIONS).
8. Select the *Mount options*. In the *Custom mount options* text field, you can enter a comma-separated list of options. For common options, refer to File system Independent Mount Options (https://linux.die.net/man/8/mount). These options are used in the `options` part of the `/etc/fstab` file.

6.2.2 Mounting partitions using Cockpit



Note: The partition must be formatted

Before you try to mount a partition or disk, you need to format the device first. For details, refer to *Procedure 6, "Formatting partitions"*.

To mount a partition, proceed as follows:

1. Navigate to the *Storage* page.
2. In the *Filesystems* view, click the device to mount.
3. Click *Mount* to open the *Mount filesystem* window.
4. Specify the *Mount point*.
5. To select mount options in the *Custom mount options* text field, you can enter a comma-separated list of options. For common options, refer to File system Independent Mount Options (https://linux.die.net/man/8/mount). These options are used in the `options` part of the `/etc/fstab` file.
6. Select at which booting stage the partition must be mounted.
7. Click *Mount* to proceed.

6.3 Managing NFS mount points

The *NFS mounts* view under the *Storage* page enables you to add, edit or delete NFS mounts.

6.3.1 Adding an NFS mount point

To add an NFS mount point, proceed as follows:

1. Navigate to the *Storage* page.
2. From the three-line menu, select *New NFS mounts* view.

3. Specify the following values:

Server address

Provide the IP address or name of the NFS server.

Path on server

Select the available path on the NFS server that can be mounted.

Local mount point

Specify a directory on the local system where the path will be mounted.

Mount options

Check any of the options:

- *Mount at boot* – to mount the path automatically after each start or restart of the system.
- *Mount read only* – you will not be able to perform changes to the data on the NFS path.
- *Custom mount options* is a comma-separated list of the **`mount`** command options.

6.3.2 Editing existing NFS mount points

To edit an NFS mount, proceed as follows:

1. Navigate to the *Storage* page.
2. In the *NFS mounts* view, click on the particular NFS mount.
3. On the next screen, click *Edit* and specify the details described in *NFS mount details*.

6.4 Managing RAID using Cockpit

Using Cockpit, you can create or modify software RAID of different levels.

6.4.1 Creating RAID using Cockpit



Note: Sufficient number of disks

Make sure that you have enough disks available according to the RAID level.

To create a software RAID, proceed as follows:

PROCEDURE 7: CREATING A RAID

1. Navigate to the *Storage* page.
2. Select the *Create RAID device* option in the three-line menu in the *Devices* view.
3. Enter the following parameters of the RAID:

   Name

   Enter a unique name for the RAID.

   RAID level

   Select one of the RAID levels.

   Chunk size

   The size of chunks in KBs. A chunk is the minimum amount of data read or written to each data disk in the array during a single read/write operation.

   Disks

   Select the disks that should be included in the RAID. The required number of disks depends on the selected RAID level.
4. Confirm the parameters by clicking *Create*. The RAID then appears in the *Devices* part.

6.4.2 Modifying RAID

Using the *Storage* plug-in of Cockpit, you can stop or delete a RAID. Here you can also remove or add disks to the array.

To modify an existing RAID, proceed as follows:

1. Navigate to the *Storage* page.
2. Click the RAID in *Devices* to open the RAID details view.
3. In the detailed view, you can stop or delete the RAID, add or remove disks and format the device. With certain RAID levels, you can switch on the *Bitmap* option that enables you to synchronize only the changes after a disk is temporarily disconnected. If the *Bitmap* is off, all data on the disk will be synchronized.



Note: Removing or adding disks

After any change in the disk number of the array, the system undergoes resynchronization that may take some time. Keep in mind that each RAID level requires a minimum number of disks. Therefore, Cockpit does not allow removing the disks that are required by the particular RAID level.

6.5 Managing volume groups and LVM

6.5.1 Creating volume groups

To create a volume group of disks, proceed as follows:

1. Click *Storage*.
2. Under the three-line menu in *Devices*, select *Create LVM2 volume group*.
3. Enter the volume group name.
4. Select the disks that will be part of the volume group.
5. Confirm the data with *Create*. The volume group appears in the *Devices* view.

6.5.2 Creating logical block volumes

If you have a volume group, you can create a logical block volume from it. To do so, proceed as follows:

1. Navigate to the *Storage* page.
2. In the *Devices*, click the volume group you want to use.
3. Click *Create new logical volume*
4. Specify a logical volume name. Select a block device and choose the size to use.
5. Select the *Block device for filesystems*.
6. Select the size to use.
7. Click *Create* to confirm the details.

8. Format the block volume by clicking *Format* and filling the details as described in *Step 4*.

6.5.3 Creating a thin logical volume

If you have a volume group, you can create a thin logical volume as described below:

PROCEDURE 8: CREATING A THIN LOGICAL VOLUME

1. Navigate to the *Storage* page.
2. Click the volume group in *Devices*.
3. In the volume group details, click *Create new logical volume*.
4. Specify a logical volume name.
5. Select a pool of thinly provisioned volumes.
6. Select the size to use.
7. Click *Create* to confirm the details.
8. Create a thin volume by clicking *Create thin volume*.
9. Enter a unique name.
10. Select the size of the volume.
11. Click *Create* to confirm the thin volume.
12. You can create several volumes of the particular volume group by clicking *Create thin volume* again and repeating the steps above.
13. Format the volumes by clicking *Format* and filling the details as described in *Step 4*.

6.5.4 Managing logical volumes

To perform any administration task on an existing logical volume, perform the following steps:

1. Navigate to the *Storage* page.
2. In the *Filesystems* view, click the logical volume.

3. Here you can perform the following actions with existing logical volumes:

Deactivate/Activate

In the three-dot menu, select *Deactivate* or *Activate*.

Mount

By clicking *Mount* and filling in the mount point and options, the volume will be mounted.

Shrink/Grow

The shrink/grow function is not available for all file systems.

In the expanded details about the volume, click *Shrink* or *Grow*.

Delete

In the three-dot menu, select *Delete*.

7 Managing networking using Cockpit

After clicking *Networking*, you can view traffic on your system, manage the firewall, manage network interfaces, or view network logs.

7.1 Managing firewall rules and zones

Cockpit enables you to create new zones or update the existing ones. In the firewall settings, you can add services to a zone or allow access to ports.



Note: Cockpit service is mandatory

Do not remove the Cockpit service from the default firewall zone as the Cockpit service may get blocked, and you may get disconnected from the server.

7.1.1 Adding firewall zones

The *public zone* is the default firewall zone. To add a new zone, proceed as follows:

PROCEDURE 9: ADDING NEW FIREWALL ZONES

1. Navigate to the *Networking* page.

2. Click *Edit rules and zones*.
3. Click *Add new zone*.
4. Select *Trust level*. Each trust level of network connections has a predefined set of included services (the Cockpit service is included in all trust levels). The description for each trust level appears in the *Description* section.
5. Define allowed addresses within the zone. Select one of the values:
    - *Entire subnet* to allow all addresses in the subnet.
    - *Range*—a comma-separated list of IP addresses with the routing prefix, for example, 192.0.2.0/24, 2001:db8::/32.
6. Click *Add zone*.

7.1.2 Adding allowed services and ports to a zone

You can add services to an existing firewall zone as described below:

PROCEDURE 10: ADDING SERVICES TO A FIREWALL ZONE

1. Navigate to the *Networking* page.
2. Click *Edit rules and zones*.
3. Click *Add services*.
4. To add a service, select *Services* and choose the services from the list.
5. To allow custom ports, select *Custom ports* and specify the port value for UDP and/or TCP. You can assign an identifier to this port.
6. To confirm the changes, click *Add services* or *Add ports*, respectively.

7.2 About network bonds

A bond interface is a combination of several network interfaces into one bond. Depending on the *Mode* (described further), network bonding can improve performance by increasing the network throughput and bandwidth. Network bonding can also increase fault tolerance by keeping overall connectivity even if some of the bonded interfaces stopped working.

7.2.1 Managing bonds

7.2.1.1 Adding bonds



Warning: Using DHCP when creating bonds may cause disconnection from Cockpit

When you try to create a network bond, you will be disconnected from Cockpit as its IP address changes if the following conditions are met:

- DHCP is used on the server running Cockpit
- the currently used network interface is added to the new network bond

To add a bond, proceed as follows:

1. Navigate to the *Networking* page.
2. Click *Add bond*.
3. Specify the following parameters of the bond interface:

   Name

   Enter a unique name for the interface.

   Interfaces

   Select which network interfaces should be grouped in the bond.

   MAC

   You can either select a specific MAC address of the underlying interface, or you can use any of the following options:

   Permanent

   Use the permanent hardware address if the device has a MAC address.

   Preserve

   During the bond activation, the MAC address is not changed.

   Random

   A random MAC address is created on each connection attempt.

Stable

Creates a hashed MAC address.

Mode

Keep the default mode or select any of the following modes:

Round Robin

Transfers packets from the first available interface to the last. The mode offers fault tolerance and load balancing.

Active Backup

Only one interface in the bonding is active. If the active interface fails, the backup will be activated.

XOR

Balancing using a transmit hash policy. The default is a modulo device count. To select a different policy, specify the `xmit_hash_policy` option in the *Option* field.

Broadcast

Everything is transmitted on all interfaces.

Adaptive Transmit Load Balancing

Channel bonding that does not require any special switch support. The outgoing traffic is distributed according to the current load on each interface.

Adaptive Load Balancing

Includes adaptive transmit load balancing and receive load balancing; no special switch support is required.

Primary

This selection is available only in the *Active Backup* mode. You can select a particular interface that will be used as primary, while other interfaces in the bond are used as secondary.

Link monitoring

Select the type of link monitoring.

Monitoring interval

Specifies the intervals at which the particular link monitor performs checks. The value is in ms.

Link up delay

Define the time in ms for how long the bond is disabled after a link has been activated. The value should be a multiple of the *Monitoring interval* value. Otherwise, it will be rounded to the nearest value. Available only for the MII link monitor.

Link down delay

Define the time in ms for how long the bond is disabled if a link failure has been detected. The value should be a multiple of the *Monitoring interval* value. Otherwise, it will be rounded to the nearest value. Available only for the MII link monitor.

Monitoring targets

Specify the list of host IP addresses that you want to monitor. Available only for the ARP link monitor.

4. Click *Apply*.

7.2.1.2 Modifying bonds

To modify a bond, proceed as follows:

1. Navigate to the *Networking* page.
2. Click on the particular bond name to open the details.
3. You can modify the following bond parameters:

Bond

Click the bond name and the bond details appear. Modify the details and click *Save*.

Connect automatically

The bond connects automatically by default. Clear the box to disable the automatic connection.

IPv4 and IPv6

After clicking *edit*, you can set an IP address and configure a specific DNS, DNS search domain and routes.

MTU

By default, *Automatic* is selected. After clicking *edit*, you can specify a particular value of the maximum transmission unit in bytes or select *Automatic*.

Bond

After clicking *edit*, you can edit the same parameters as when you were creating the bond interface.

7.3 Managing network bridges

A network bridge is a device that creates a single aggregated network from multiple networks.

7.3.1 Creating network bridges



Warning: Using DHCP when creating network bridges may cause disconnection from Cockpit

When you try to create a network bridge, you will be disconnected from Cockpit as its IP address changes if the following conditions are met:

- DHCP is used on the server running Cockpit
- the currently used network interface is added to the new network bridge

To create a network bridge, proceed as follows:

1. Navigate to the *Networking* page.
2. In the *Interfaces* view, click *Add bridge*.
3. Specify the following:

   Name

   Specify a unique name for the bridge.

   Ports

   Select the interfaces to be included in the bridge.

Spanning Tree Protocol (STP)

STP is a network protocol used for Ethernet networks that prevents bridge loops by setting a preferred link whenever network switches are connected with several links. This preferred link is used for all Ethernet traffic unless it fails. In that case, a redundant link is used instead. For details regarding STP, see STP (https://en.wikipedia.org/wiki/Spanning_Tree_Protocol).
If you enable the STP protocol, you can edit the following settings:

STP priority

The lower the priority, the higher the probability of the switch becoming the root switch.

STP forward delay

Specify the time spent in the listening and learning state (in seconds). The default value is 15 s, but you can use any value between 4 and 30 s.

STP hello time

Specify the time between each bridge protocol data unit (BDPU) that is sent on a port (in seconds). The default value is 2 s, but the recommended range is 1 to 10 s.

STP maximum message age

Specify the maximum length of time that passes before a bridge port saves its configuration BPDU information.

4. Click *Add*

7.3.2 Modifying or deleting existing bridges

To modify or delete a bridge, proceed as follows:

1. Navigate to the *Networking* page.
2. In the *Interfaces* view, click the bridge name to open the details.
3. To delete the bridge, click *Delete*.
   To modify, change any of the following details and click *Save*.

General

The bridge connects automatically by default. To disable the automatic connection, uncheck the option.

IPv4 and IPv6

After clicking *edit*, you can set the IP address and configure a specific DNS, DNS search domain and routes.

Bridge

By clicking *edit*, you can edit all parameters of the bridge.

7.4 Managing VLANs using Cockpit

A virtual local area network is a logical subnetwork that groups devices from different physical LANs.

7.4.1 Creating a virtual local area network

To add a VLAN, proceed as follows:

1. Navigate to the *Networking* page.
2. In the *Interfaces* view, click *Add VLAN*.
3. Fill in the VLAN details:

   Parent

   Select the parent network interface.

   VLAN ID

   Specify an ID in the range 1–4094.

   Name

   Enter the name of the VLAN.
4. Click *Add VLAN*.

7.4.2 Modifying or deleting existing VLANs

To modify or delete an existing VLAN, proceed as follows:

1. Navigate to the *Networking* page.

2. In the *Interface* view, click the VLAN name.

3. To delete the VLAN, click *Delete* or modify any of the following details and click *Save*.

Parent

Select the parent network interface.

VLAN ID

Specify an ID in the range 1–4094.

Name

Enter the name of the VLAN.

8 Managing virtual machines using Cockpit

The *Virtual machines* section allows you to view, create and manage virtual machines, storage pools and networks, and also to import virtual machines.



Note: *Virtual machines* not visible

If you do not see the plug-in in your Cockpit instance, you need to install the `kvm_server` pattern.

```
> sudo  zypper install -t pattern kvm_server
> sudo  systemctl enable --now libvirtd
```

Using the *Virtual machines* screen, you can perform the following tasks:

- Install virtual machines in *Section 8.2, “Installing a virtual machine using Cockpit”*
- Create virtual machines in *Section 8.1, “Creating virtual machines using Cockpit”*
- Modify virtual machine configurations in *Section 8.3, “Modifying virtual machines using Cockpit”*
- Import a virtual machine in *Section 8.3.10, “Importing a virtual machine”*
- Clone a virtual machine in *Section 8.3.7, “Cloning a virtual machine”*
- Rename a virtual machine in *Section 8.3.9, “Renaming a virtual machine”*
- Delete a virtual machine in *Section 8.3.8, “Deleting a virtual machine”*

8.1 Creating virtual machines using Cockpit

The Cockpit Web console allows you to create virtual machines.

To create a virtual machine using Cockpit:

1. Navigate to the *Virtual machines* page.
2. Click *Create VM*.
3. Select the *Connection* as either *System* or *User session*.
   For more information on the connections, click the information icon next to *Connection*.
4. Select an installation type from the drop-down list and proceed with creating the VM.

   - *Section 8.1.1, "Download an OS"*
     For this installation type, Cockpit provides a list of popular Linux distributions (Fedora, Ubuntu, Rocky Linux, etc.). It automatically downloads the official ISO and starts the installer.
   - *Section 8.1.2, "Cloud base image"*
     For this installation type, you create a virtual machine from the cloud image file on the host's file system. You must provide the path to the file for this. This installation type supports unattended installation, where you can pre-fill a username, password, and SSH key in Cockpit. The VM then boots and configures itself automatically.
   - *Section 8.1.3, "Local installation media (ISO image or distro install tree)"*
     For this installation type, you manually provide a path to an .iso file already stored on your server's hard drive.
   - *Section 8.1.4, "URL (ISO image or distro install tree)"*
     For this installation type, you provide a Web link to an installation tree (common for automated network installs).
   - *Section 8.1.5, "Network boot (PXE)"*
     For this installation type, the system boots the VM from a network server. This is typically used in enterprise environments with existing deployment infrastructure.

8.1.1 Download an OS

You can select *Download an OS* to create a VM with your preferred OS. Cockpit provides a list of popular Linux distributions (like OpenSUSE Leap, Fedora, Ubuntu, or CentOS). When you select the distribution, Cockpit downloads the official ISO image from the internet and creates a VM based on the given configurations.

1. Select the distribution from the *Operating system* drop-down list.
2. Define the *Storage limit* and *Memory*.
   The available storage limit and memory are displayed so that you can configure the storage limit and memory accordingly.
3. To run the VM after creation, click *Create and run*.
   The virtual machine has been installed.
4. To edit the VM parameters after creation, click *Create and edit*.
   Modify the parameters and click *Install*.
   For more information on modifying the parameters, see *Section 8.3, "Modifying virtual machines using Cockpit"*.

8.1.2 Cloud base image

You can create a virtual machine from the cloud image file on the host's file system. You must provide the path to the file for this. This installation type supports unattended installation, where you can pre-fill a username, password, and SSH key in Cockpit. The VM then boots and configures itself automatically.

1. Select the path to the cloud image file on the host's file system from the *Installation source* drop-down list.
2. Select the distribution from the *Operating system* drop-down list.
3. Define the *Storage limit* and *Memory*.
   The available storage limit and memory are displayed so that you can configure the storage limit and memory accordingly.
4. To run the VM after creation, click *Create and run*.
   The virtual machine has been installed.
5. To edit the VM parameters after creation, click *Create and edit*.
   Modify the parameters and click *Install*.

For more information on modifying the parameters, see *Section 8.3, "Modifying virtual machines using Cockpit"*.

8.1.3 Local installation media (ISO image or distro install tree)

You can create a VM from the local installation media using the path to the ISO file on the host's file system. In this, you must provide the direct file path in which the ISO image is stored, and the VM boots from that local file and begins the installation.

1. Select the path to the ISO file on the host's file system from the *Installation source* drop-down list.
2. Select the distribution from the *Operating system* drop-down list.
3. Define the *Storage limit* and *Memory*.
   The available storage limit and memory are displayed so that you can configure the storage limit and memory accordingly.
4. To run the VM after creation, click *Create and run*.
   The virtual machine has been installed.
5. To edit the VM parameters after creation, click *Create and edit*.
   Modify the parameters and click *Install*.
   For more information on modifying the parameters, see *Section 8.3, "Modifying virtual machines using Cockpit"*.

8.1.4 URL (ISO image or distro install tree)

You can create a virtual machine by downloading the OS from a URL of the ISO image or distro install tree. In this type of installation, you can point the VM to a remote installation source via a web link (HTTP/HTTPS/FTP). It is particularly useful for network-based installations where the host doesn't have the image saved locally but can pull the necessary boot files directly from a public or private mirror.

1. Specify the URL of the ISO file in the *Installation source*.
2. Select the distribution from the *Operating system* drop-down list.
3. Define the *Storage limit* and *Memory*.
   The available storage limit and memory are displayed so that you can configure the storage limit and memory accordingly.
4. To run the VM after creation, click *Create and run*.

The virtual machine has been installed.

5. To edit the VM parameters after creation, click *Create and edit*.
   Modify the parameters and click *Install*.
   For more information on modifying the parameters, see *Section 8.3, "Modifying virtual machines using Cockpit"*.

8.1.5 Network boot (PXE)

You can create a VM by installing the OS over the network using the Network boot PXE (Preboot eXecution Environment). In this type of installation, the VM starts without a local OS and requests boot files from a PXE server over the network. The PXE server supplies the bootloader, kernel, initrd, and OS installation files, allowing the operating system to install directly from the network.

1. Select the installation source from the *Installation source* drop-down list.
2. Select the distribution from the *Operating system* drop-down list.
3. Define the *Storage limit* and *Memory*.
   The available storage limit and memory are displayed so that you can configure the storage limit and memory accordingly.
4. To run the VM after creation, click *Create and run*.
   The virtual machine has been installed.
5. To edit the VM parameters after creation, click *Create and edit*.
   Modify the parameters and click *Install*.
   For more information on modifying the parameters, see *Section 8.3, "Modifying virtual machines using Cockpit"*.

8.2 Installing a virtual machine using Cockpit

To install an already created VM, navigate to the *Virtual machines* page and click *Install* corresponding to the VM.



Note

You can also create a VM and install it.

For more information on creating a VM, see *Section 8.1, "Creating virtual machines using Cockpit"*.

8.3 Modifying virtual machines using Cockpit

The Cockpit web console allows you to view, filter, and manage virtual machines (VMs), storage pools, and networks. You can also use this interface to import existing VMs and activate, deactivate, and delete your virtualization resources.

You can dynamically modify the VM through Cockpit. You can add disks to the VM for additional block storage for data isolation or LVM expansion, add network interfaces to enable traffic segregation between public and private networks, add host devices to allow the VM to access physical hardware directly, and add shared directories to facilitate file transfers between the host and guest file systems. You can also take snapshots before and in between to create a recovery point, ensuring the VM can be restored if a configuration change leads to a boot failure or OS corruption.

8.3.1 Modifying virtual machine

To modify a virtual machine's configuration,

1. Navigate to the *Virtual machines* page.
2. Click the required VM name.
3. Adjust the VM attributes as described in the following sections.

8.3.1.1 Modifying basic VM attributes

In the *Overview* section, you can change the following attributes of the VM:

- To modify the *Memory*, click *edit* and change the storage and memory, and then click *Save*.
- To modify the *Boot order*, click *edit* and select the up and down arrows to set the boot order, and then click *Save*.
- To enable or disable running the VM when the host boots, click the *Autostart* toggle button.
- To modify the *Watchdog* device type, click *edit*, select the action, and then click *Save*.
  Watchdogs act when systems stop responding. To use this virtual watchdog device, the guest system also needs to have an additional driver and a running watchdog service.

The available options are:

- *Reset* performs a hard reset of the VM.
- *Power off* shuts down the VM. Use this if a hung VM might cause data corruption on shared storage.
- *Inject a non-maskable interrupt* when the watchdog timer expires. The hypervisor sends a non-maskable interrupt (NMI) signal to the virtual processor.
- *Pause* suspends the VM in its current state. This allows an administrator to verify the VM for debugging before manual recovery.
- *Do nothing* records the event in the host logs but takes no physical action on the VM.

- To modify the *Vsock*, click *edit* and modify or remove the custom identifier.
  Virtual socket support enables communication between the host and guest over a socket. It still requires special vsock-aware software to communicate over the socket.
- To modify the *Firmware*, click the firmware name, select the required firmware from the drop-down list, and click *Save*.

8.3.2 Managing disks

You can add a new disk to the VM, add an existing disk, or attach a virtual disk that already exists as a file on the host machine's file system.

8.3.2.1 Adding disks

1. In the *Disks* section, click *Add disk*.
2. To create a new disk source, select *Create new* in *Source*.
   - Select the pool name from the *Pool* drop-down list.
   - Specify the new volume name in the *Name* field.
   - Select the disk size and format.
   - Click *Show additional options* and select the cache and bus from the drop-down lists.

The available options for cache are:

- *None*: Bypass the guest's cache; the VM talks directly to the host's storage.
- *Writeback*: Writes data to the host's page cache and reports it as “done” immediately.
- *Writethrough*: Writes data to both the cache and the disk before the write is confirmed.
- *DirectSync*: This is similar to *None* but bypasses even more layers of the host's file system cache. This is useful for specific database workloads that handle their own caching.

The *Bus* type determines the protocol used to communicate between the VM and the storage. The following bus options are available:

- *VirtIO*: VirtIO is a paravirtualization standard. Instead of the host machine pretending to be a specific piece of physical hardware (which is slow), the guest OS uses a special driver that knows it is running in a virtual environment. This is best used for standard production and high-performance.
- *SATA/IDE*: These emulate older physical hardware. Use these only if you are installing an older OS (like Windows XP) that doesn't have VirtIO drivers out of the box. This can be used for OS installation or troubleshooting.
- *SCSI*: Often used in enterprise environments, specifically with the `virtio-scsi` controller, which allows for advanced features like TRIM/Discard (to reclaim space on SSDs). This can be used for enterprise storage or databases.
- *USB*: The USB option allows you to pass through physical devices from your host machine (like a thumb drive, webcam, or security dongle) directly into the VM. Use USB only when you need to connect a specific physical device or if you are troubleshooting a system that cannot load VirtIO drivers.

- Specify the disk identifier in *Disk identifier*.

3. To use an existing disk source, select *Use existing* in *Source*.

- Select the pool name from the *Pool* drop-down list.
- Select the new volume name from the *Volume* drop-down list.
- Click *Show additional options* and select the cache and bus from the drop-down lists.

The available options for cache are:

- *None*: Bypass the guest's cache; the VM talks directly to the host's storage.
- *Writeback*: Writes data to the host's page cache and reports it as “done” immediately.
- *Writethrough*: Writes data to both the cache and the disk before the write is confirmed.
- *DirectSync*: This is similar to *None* but bypasses even more layers of the host's file system cache. This is useful for specific database workloads that handle their own caching.

The *Bus* type determines the protocol used to communicate between the VM and the storage. The available options for bus are:

- *VirtIO*: VirtIO is a paravirtualization standard. Instead of the host machine pretending to be a specific piece of physical hardware (which is slow), the Guest OS uses a special driver that knows it is running in a virtual environment. This is best used for standard production and high-performance.
- *SATA/IDE*: These emulate older physical hardware. Use these only if you are installing an older OS (like Windows XP) that doesn't have VirtIO drivers out of the box. This can be used for OS installation or troubleshooting.
- *SCSI*: Often used in enterprise environments, specifically with the `virtio-scsi` controller, which allows for advanced features like TRIM/Discard (to reclaim space on SSDs). This can be used for enterprise storage or databases.
- *USB*: The USB option allows you to pass through physical devices from your host machine (like a thumb drive, webcam, or security dongle) directly into the VM. Use USB only when you need to connect a specific physical device or if you are troubleshooting a system that cannot load VirtIO drivers.

- Specify the disk identifier in *Disk identifier*.

4. To use a custom path for the disk source, select the *Custom path* option in *Source*.

   - Select the custom path from the *Custom path* drop-down list.
   - Select the device as *Disk image file* or *CD/DVD disc* from the *Device* drop-down list.
   - Click *Show additional options* and select the cache and bus from the drop-down lists.

The available options for cache are:

- *None*: Bypass the guest's cache; the VM talks directly to the host's storage.
- *Writeback*: Writes data to the host's page cache and reports it as “done” immediately.
- *Writethrough*: Writes data to both the cache and the disk before the write is confirmed.
- *DirectSync*: This is similar to *None* but bypasses even more layers of the host's file system cache. This is useful for specific database workloads that handle their own caching.

The *Bus* type determines the protocol used to communicate between the VM and the storage. The available options for bus are:

- *VirtIO*: VirtIO is a paravirtualization standard. Instead of the host machine pretending to be a specific piece of physical hardware (which is slow), the Guest OS uses a special driver that knows it is running in a virtual environment. This is best used for standard production and high-performance.
- *SATA/IDE*: These emulate older physical hardware. Use these only if you are installing an older OS (like Windows XP) that doesn't have VirtIO drivers out of the box. This can be used for OS installation or troubleshooting.
- *SCSI*: Often used in enterprise environments, specifically with the `virtio-scsi` controller, which allows for advanced features like TRIM/Discard (to reclaim space on SSDs). This can be used for enterprise storage or databases.
- *USB*: The USB option allows you to pass through physical devices from your host machine (like a thumb drive, webcam, or security dongle) directly into the VM. Use USB only when you need to connect a specific physical device or if you are troubleshooting a system that cannot load VirtIO drivers.

- Specify the disk identifier in *Disk identifier*.

5. Click *Add*.
   The new disk has been added to the virtual machine.

8.3.2.2 Modifying disks

1. In the *Disks* section, click *Edit* corresponding to the disk you want to modify.

2. Select the access type as *Read-only* or *Writeable*.
3. Select the cache and bus from the drop-down lists.
   The available options for cache are:

   - *None*: Bypass the guest's cache; the VM talks directly to the host's storage.
   - *Writeback*: Writes data to the host's page cache and reports it as “done” immediately.
   - *Writethrough*: Writes data to both the cache and the disk before the write is confirmed.
   - *DirectSync*: This is similar to *None* but bypasses even more layers of the host's file system cache. This is useful for specific database workloads that handle their own caching.

   The *Bus* type determines the protocol used to communicate between the VM and the storage.
   The available options for bus are:

   - *VirtIO*: VirtIO is a paravirtualization standard. Instead of the host machine pretending to be a specific piece of physical hardware (which is slow), the Guest OS uses a special driver that knows it is running in a virtual environment. This is best used for standard production and high-performance.
   - *SATA/IDE*: These emulate older physical hardware. Use these only if you are installing an older OS (like Windows XP) that doesn't have VirtIO drivers out of the box. This can be used for OS installation or troubleshooting.
   - *SCSI*: Often used in enterprise environments, specifically with the `virtio-scsi` controller, which allows for advanced features like TRIM/Discard (to reclaim space on SSDs). This can be used for enterprise storage or databases.
   - *USB*: The USB option allows you to pass through physical devices from your host machine (like a thumb drive, webcam, or security dongle) directly into the VM. Use USB only when you need to connect a specific physical device or if you are troubleshooting a system that cannot load VirtIO drivers.

4. Click *Save*.
   The attributes are updated.

8.3.2.3 Removing disks

1. In the *Disks* section, click the three-dot icon corresponding to the disk you want to remove.

2. Click *Remove*.
3. To detach the disk from the VM so the disk is not visible to the guest OS, click *Remove*.
4. To detach the disk from the virtual machine and delete the corresponding file on the VM host, click *Remove and delete file*.

8.3.3 Managing network interfaces

You can add a new virtual NIC to your VM using one of the following network interfaces:

- `Virtual network`, the recommended type for general guest connectivity on hosts with dynamic / wireless networking configurations. Provides a connection whose details are described by the named network definition.
- `Bridge to LAN`, the recommended type for general guest connectivity on hosts with static wired networking configurations. Provides a bridge from the guest virtual machine directly onto the LAN. This needs a bridge device on the host with one or more physical NICs.
- `Direct attachment`, the recommended type for high performance or enhanced security.

8.3.3.1 Adding network interfaces

1. In the *Network interfaces* section, click *Add network interface*
2. Select the interface type as *Virtual network*, *Bridge to LAN*, or *Direct attachment* from the drop-down list.
    - For a *Virtual network*:
        - Select the source from the *Source* list.
        - Select the model from the *Model* list.
        - To generate the MAC address automatically, in *MAC address*, select *Generate automatically*.

To set the MAC address manually, in *MAC address*, select *Set manually*, and specify the required MAC address.

- For a *Bridge to LAN*:
  - Select the model from the *Model* list.
  - If you want to generate the MAC address automatically, in *MAC address*, select *Generate automatically*.
    If you want to set the MAC address manually, in *MAC address*, select *Set manually*, and specify the required MAC address.
- For a *Direct attachment*, select *Direct attachment* and select the source from the *Source* list.
  - Select the *Mode* as *VEPA*, *Bridge*, *Private*, or *Passthrough*.
    For more information on mode, see https://libvirt.org/formatdomain.html#direct-attachment-to-physical-interface↗.
  - Select the model from the *Model* list.
  - To generate the MAC address automatically, in *MAC address*, select *Generate automatically*.
    To set the MAC address manually, in *MAC address*, select *Set manually*, and specify the required MAC address.

3. Click *Add*.
   The new network interface has been added to the virtual machine.

8.3.3.2 Modifying network interfaces

1. In the *Network interfaces* section, click *Edit* corresponding to the network interface you want to modify.
2. Select the interface type from the drop-down list.
3. For a *Virtual network*, select *Virtual network*.
   - Select the source from the *Source* list.
   - Select the model from the *Model* list.
   - Modify the *MAC address*.

4. For a *Bridge to LAN*, select *Bridge to LAN*.
    - Select the model from the *Model* list.
    - Modify the *MAC address*.
5. For a *Direct attachment*, select *Direct attachment* and select the source from the *Source* list.
    - Select the *Mode* as *VEPA*, *Bridge*, *Private*, or *Passthrough*.
      For more information on mode, see https://libvirt.org/formatdomain.html#direct-attachment-to-physical-interface.
    - Select the model from the *Model* list.
    - Modify the *MAC address*.
6. Click *Save*.
   The network interface is updated.

8.3.3.3 Plugging or unplugging network interfaces

You can plug or unplug network interfaces from the virtual machine while troubleshooting or to ensure security. You can also unplug a virtual network interface to move it to a different VLAN or bridge without shutting down the entire networking service.

In the *Network interfaces* section, click the *Plug* or *Unplug* button corresponding to the network interface.

The network interface is plugged into or unplugged from the virtual machine as selected.

8.3.3.4 Removing network interfaces

1. In the *Network interfaces* section, click the three-dot icon corresponding to the network interface you want to remove.
2. Click *Remove*.
   The network interface is removed from the virtual machine.

8.3.4 Managing host devices

You can add host devices to give the VM direct access to physical hardware currently plugged into the host.

8.3.4.1 Adding host devices

1. In the *Host devices* section, click *Add host device*.
2. Select the host device type as *USB* or *PCI*.
3. Select the required devices to be added from the given list.
4. Click *Add*.
   The selected host devices are added to the virtual machine.

8.3.4.2 Removing host devices

In the *Host devices* section, click *Remove* corresponding to the host device you want to remove.

8.3.5 Snapshots

A snapshot captures the state of a virtual machine at a specific point in time, providing a recovery point for the operating system and its data. Using snapshots, you can revert to the configuration that was set when the snapshot was taken.

8.3.5.1 Creating snapshots

1. In the *Snapshots* section, click *Create snapshot*.
2. Modify the snapshot name if required.
   The default snapshot name is the date and time at which the snapshot was taken.
3. Enter a description for the snapshot.
4. Click *Create*.
   A snapshot is created with the given name and description.

8.3.5.2 Deleting or reverting snapshots

You can revert or delete snapshots. Reverting to the snapshot takes the VM back to the time of the snapshot and the current state is lost, along with any data not captured in a snapshot. When you delete a snapshot, the snapshot and all of its captured content are lost.

To revert a snapshot, click *Revert* corresponding to the snapshot you want to revert and click *Revert* to confirm.

To delete a snapshot, click *Delete* corresponding to the snapshot you want to delete and click *Delete* to confirm.

8.3.6 Managing shared directories

Shared host directories must be manually mounted inside the VM. You can mount the shared folder using the command **`mount -t virtiofs hostshare [mount point]`** in which the `mount point` is the mount point inside the guest.

8.3.6.1 Adding shared directories

1. Add the shared host directories manually:

   ```
   mount -t virtiofs [mount tag] [mount point]
   ```

   In this, the *mount tag* is a unique name you give the shared folder so the VM can identify it, while the *mount point* is the actual folder path inside the VM where those files will appear.
2. In the *Shared directories* section, click *Add shared directory*.
3. Select the host path that is to be exported from the *Source path* drop-down list.
4. Specify the tag name to be used by the guest to mount this export point in *Mount tag*.
5. To use extended attributes on files and directories, click *Show additional options* and select the checkbox.
6. Click *Share*.

8.3.6.2 Removing shared directories

1. Click *Remove* corresponding to the shared directory you want to remove.
2. Click *Remove* to confirm removing the shared directory.

8.3.7 Cloning a virtual machine

To clone a virtual machine:

1. Navigate to the *Virtual machines* page and click the three-dot menu corresponding to the VM.

2. Click *Clone*.
3. Specify a name for the cloned VM and click *Clone*.
   A clone of the VM has been created.

8.3.8 Deleting a virtual machine

To delete a virtual machine:

1. Navigate to the *Virtual machines* page and click the three-dot menu corresponding to the VM.
2. Click *Delete*.
3. Select the associated storage files to be deleted with the VM and click *Delete*.
   The VM and the associated storage files have been deleted.

8.3.9 Renaming a virtual machine

To rename a virtual machine,

1. Navigate to the *Virtual machines* page and click the three-dot menu corresponding to the VM.
2. Click *Rename*.
3. Specify the new name and click *Rename*.

8.3.10 Importing a virtual machine

Import a virtual machine if you have a disk image with an already installed system: `raw`, `qcow2` and so on. To import a virtual machine, proceed as follows:

1. Navigate to the *Virtual machines* page and click *Import VM*.
2. In the *Details* section, specify the following details:
   - Enter a name for the VM.
   - Select the *Connection* as either *System* or *User session*.



Note

Click the information icon next to *Connection* to view details on the available options.

- Select the disk image located on the host file system.
- Select the appropriate operating system and memory.

3. Define the *Storage parameters*, including the storage limit and memory.
4. To run the VM after importing, click *Import and run*.
5. To edit the VM parameters after importing, click *Import and edit*.
   Modify the parameters and click *Import and run*.

8.4 Managing storage pools using Cockpit

Storage pools are locations (like directories or partitions) where VM images are stored.

8.4.1 Creating a storage pool

1. Navigate to *Virtual machines* and click the *Storage pools* link.
2. Click *Create storage pool*.
3. In the *Create storage pool* page, select the *Connection* as either *System* or *User session*.
4. Specify a storage pool name and select the storage pool type from the drop-down list.
5. For storage pool type `File system directory`, specify the target path as the path on the host's file system.
6. For storage pool type `Network file system`, fill in the following attributes:
   - In *Target path*, specify the path on the host's file system.
   - In *Host*, specify the host name.
   - In *Source path*, specify the directory on the server that is being exported.

7. For storage pool type *iSCSI*, fill in the following attributes:
    - In *Target path*, specify the path on the host's file system.
    - In *Host*, specify the host name.
    - In *Source path*, specify the iSCSI target IQN.
8. For storage pool type *Physical disk device*, fill in the following attributes:
    - In *Target path*, specify the path on the host's file system.
    - In *Source path*, specify the physical disk device on the host.
    - Select the format from the drop-down list.
9. For storage pool type *Pre-formatted block device*, fill in the following attributes:
    - In *Target path*, specify the path on the host's file system.
    - In *Source path*, specify the path to the file.
    - Select the format from the drop-down list.
10. For storage pool type *iSCSI direct target*, fill in the following attributes:
    - In *Host name*, specify the name of the host.
    - In *Source path*, specify the iSCSI target IQN.
    - In *Initiator*, specify the iSCSI initiator IQN.
11. To configure the pool to start when the host boots, select the checkbox *Start pool when host boots*.
12. Click *Create*.

8.4.2 Managing a storage pool

To perform any administration task on an existing storage pool, perform the following steps:

1. Navigate to the *Virtual machines* page and select the *Networks* link.
2. Here you can perform the following actions on existing storage pools:

    Activate

    Click the *Activate* button next to the required storage pool.

Deactivate

Click the *Deactivate* button next to the storage pool name.

Delete

In the three-dot menu, select *Delete*.

8.5 Managing networks using Cockpit

Cockpit enables you to create virtual networks and manage them.

8.5.1 Creating a network

1. Navigate to the *Virtual machines* page and click the *Networks* link.
2. Click *Create virtual network*.
   In the *Create virtual network* page, enter the following details:
   - Enter a network name.
   - Select *Forward mode* and *Device* from the drop-down lists.
   - Select the *IP configuration*, then provide the IPv4 and/or IPv6 address and prefix length.
   - To configure a DHCP range, select Set DHCP range and enter the start and end values.
3. Click *Create*.

8.5.2 Managing a network

To perform any administration task on an existing network, perform the following steps:

1. Navigate to the *Virtual machines* page and select the *Networks* link.
2. Here you can perform the following actions on existing storage pools:

   Activate

   Click the *Activate* button next to the required network.

   Deactivate

   Click *Deactivate* next to the network name.

Delete

In the three-dot menu, select *Delete*.

9 Working with containers

After the first login to Cockpit, you need to start Podman. Keep the default check box selected to start Podman automatically on each boot.

The *Podman containers* page enables you to pull images from registries and manage your containers. You can also filter the view by entering a filter criterion into the filter field.

9.1 Managing container images



Note: openSUSE registry and Docker Hub not enabled by default

The openSUSE registry and Docker Hub are not configured in the default installation. To download container images from those registries, you need to add the registries to the `/etc/containers/registries.conf` file as follows:

```
unqualified-search-registries = ["registry.suse.com", "registry.opensuse.org",
 "docker.io"]
```

In the *Images* view, you can download, update or delete already pulled images. Each function is available through the three-dot menu. After clicking the menu, there are the following options:

- *Download new image*: How to proceed with downloading an image is described in *Procedure 11, "Downloading a new image"*.
- *Pull all images*: Cockpit pulls new versions of the container images you already downloaded.
- *Prune all unused images*: All images that are not used by any container will be removed.

PROCEDURE 11: DOWNLOADING A NEW IMAGE

1. In the *Podman containers* › *Images* view, open the three-dot menu and select *Download new image*.

2. Select the *Owner* to define who can see the downloaded image. The *System* restricts the image visibility to users with administrative access. The image downloaded under the *User* owner is visible to the regular user and also to all other users with administrative access.
3. Choose a preferred image registry or proceed with `All registries`.
4. Define the *Tag*. The default value is `latest`.
5. Fill in the image name or description in the *Search for* field to start the search.
   Cockpit suggests possible images according to the entered name, registry and tag.
6. Select the desired image and click *Download*.

9.2 Managing containers using Cockpit

9.2.1 Running new containers from images



Note: Image required to run a container

To run a container, you need a container image. The image can be pulled using Podman or Cockpit. When using Cockpit, you can pull an image in advance as described in *Procedure 11, "Downloading a new image"*, or you can pull the image directly from the *Create container* form as described below.

To run a new container, proceed as follows:

1. Navigate to the *Podman containers* page.
2. If you pulled an image in advance:
   a. In the *Images* view, click *Show images*.
   b. Click *Create container* next to the image you want to use.
3. If you do not have the image, click *Create container* in the *Containers* view.
4. In the *Create container* window, enter the container details as described below. Note that some options are available only for system administrators.

In the *Details* tab, enter the following details:

Owner

Select whether the container will be visible only for users with **sudo** privileges by selecting *system*. The *user* defines that the container is visible to privileged users and regular users.

Name

Specify a unique name for the container.

Image

This field is enabled if you do not have an image. After you start typing the image name, Cockpit makes suggestions of images in the configured registries.

Pull the latest image

The checkbox is available if you are creating the container from an already downloaded image. If selected, the latest image version is pulled before the container is started.

Command

You can specify a command to run in the container.

With terminal

Select the option to gain access to the container using a terminal. If not selected, the container will be in the detached state.

Memory limit

You can limit the maximum memory consumption of the container by checking the box and specifying the limit.

CPU shares

Specify the weight of the container to use CPU time. The default weight is 1024. The weight applies only if containers are under high load. If the tasks in one container are idle, other containers may use its CPU time.

If you have four containers, two of them have CPU shares of 512 and the other two have 1024. Thus, under high load, the containers with lower CPU shares get only 16,5% of CPU time, while those with 1024 CPU shares get 33% of CPU time.

In the *Integration* tab, you can enter the following parameters:

Port mapping

After you click the *Add port mapping* button, specify the host IP address, the host port to map the container port onto, the container port and select the protocol. If you do not set the host IP address or set the value to 0.0.0.0, the port is bound to ALL host IP addresses. If you omit the host port, a random one is used for the mapping.

Volumes

This field maps a path in a container onto a path on the host machine. Fill in the host path, the container path and select the SELinux label.

The SELinux label *private* defines that the volume is accessible only from the particular container. The *shared* label means that all containers can access the volume.

Environment variables

To define environment variables in the container, click *Add variable* and fill in *Key* and *Value*. You can enter multiple variables by adding more lines.

In the *Health check* tab, you can set the interval for triggering commands to check the container's status. Fill in the following parameters:

Command

Specify the command that is triggered to check the container status.

Interval

Specify the interval of checks in seconds.

Timeout

The maximum time in seconds to wait before the interval is considered failed.

Start period

The time interval after the container is started when the health check is not performed.

Retries

Specify how many times the check can be performed before the status is considered unhealthy.

When unhealthy

Select the action to take after a container is considered unhealthy.

5. To create the container, click *Create* or *Create and run* to create and start the container.

9.2.2 Further actions with running containers

Under the three-dot menu, you can perform the following actions:

- delete the container
- pause the container
- commit changes performed to the container, for example, installing packages to the container
- checkpoint the container—write the state of the container to disk and stop the container
- restart the container either by regular *Restart*, where processes running inside the container are stopped, or by *Force restart*, where the processes are killed, and you may lose data
- stop the container, either by regular *Stop*, *Force stop* or *Checkpoint*. When using *Checkpoint*, the state of all processes in the container is written to the disk, and after the next start, the container is restored to the same point before stopping.

By expanding the container details, you can access the container's terminal in the *Console* tab and view its information in other tabs.

9.3 Pods management

9.3.1 Creating pods

Cockpit enables you to create pods in which you can then create containers. To create a pod, follow the steps:

1. Navigate to the *Podman containers* page.
2. Click *Create pod*.
3. Fill in the pod details:

   Name

   Enter a unique name for the pod.

   Owner

   Specify whether the pod will be visible only under `root` privileges or also to regular users.

Port mapping

After clicking *Add port mapping*, you can map a pod port onto a host port. Specify the container's port, assign the desired host port and IP address. If the host IP address is not set or is set to 0.0.0.0, the port is bonded to all host IP addresses. If you omit the host port number, a random port number is assigned to the mapping.

Volumes

After clicking *Add volume*, you can map a directory on the host onto the container's volume. Select the host path, enter the path in containers and select the SELinux labeling.

4. Click *Create* to confirm the pod creation.

9.3.2 Creating containers in pods



Important: Existing containers cannot be added to pods

During the planning, note that only new containers can be run in a pod. You cannot add an already created container that has not been run under a pod to any pod.

To create containers in a pod, follow the steps:

1. Navigate to the *Podman containers* page.
2. In the desired pod group, click *Create container in pod*.
3. Fill in the container details as described in *Section 9.2.1, "Running new containers from images"*. Remember that the owner of new containers is the same as the owner of the particular pod.

10 User administration using Cockpit

The Cockpit *Accounts* screen enables you to administer user accounts and groups. You can search and view the users, manage the user access and delete users.



Note: User administration only for server administrators

Only users with *Administrative access* can edit other users.

Using the *Accounts* Cockpit screen, you can perform the following tasks:

- Creating new users of the system as described in *Section 10.2, "Creating user accounts using Cockpit"*
- Assigning **sudo** privileges to user accounts as described in *Section 10.1, "Modifying existing user accounts"*
- Forcing a change of a user's password as described in *Section 10.1, "Modifying existing user accounts"*
- Locking a particular user account as described in *Section 10.1, "Modifying existing user accounts"*

10.1 Modifying existing user accounts

To modify a user account, proceed as follows:

1. Navigate to the *Accounts* page.
2. Click the account you want to modify.
3. In the user details view, you can perform the following actions:

   Delete the user

   Click *Delete* to remove the user from the system.

   Terminate user's session

   By clicking *Terminate session*, you can log a particular user out of the system.

   Manage access to the account

   You can set a date for when the account will expire. The default is to never expire.

   You can disallow the user to use their password to log in. The user then must use a different method of authentication.

   Manage the user's password

   Click *Set password* to set a new password for the account.

   By clicking *Force change*, the user will have to change the password on the next login.

   Click *edit* to set whether and when the password expires.

   Add SSH key

   You can add an SSH key for passwordless authentication via SSH. Click *Add key*, paste the contents of the public SSH key and confirm it by clicking *Add*.

Change shell

Click the shell name and select the shell from the list.

View login history

You can view the login history of the user.

10.2 Creating user accounts using Cockpit

To add a new user to the system, proceed as follows:

1. Navigate to the *Accounts* page.
2. Click *Create new account* to open the window that enables you to add a new user.
3. Fill in the user account details. You can assign a different home directory to the user in the *Home directory* text box. If you do not specify a directory, the standard `/home/USERNAME` path is used. If you select *Disallow password authentication*, the user will have to use an authentication method other than filling in a password, for example, an SSH login.
4. Click *Create* to confirm the account.
5. To add an SSH key to the account, you need to modify the account as described in *Section 10.1, "Modifying existing user accounts"*.

10.3 Creating user groups

The topic covers the creation of user groups.

To create a user group, proceed as follows:

1. Navigate to the *Accounts* page.
2. Click *Create new group*.
3. Enter a unique name for the group and specify or leave the default ID.



Note

The already existing group ID cannot be overwritten. Group IDs under 1000 are usually reserved for system accounts, services, and so on. If you create a group with an ID less than 1000, the group cannot be later deleted using Cockpit.

11 Managing services using Cockpit

The following sections describe how to start, stop and restart a service, target, socket, timer or path.

11.1 Managing `systemd` units

To manage a `systemd` unit, proceed as follows:

1. Click the *Services* page.
2. Select the appropriate tab (*System services*, *Targets*, *Sockets*, *Timers* or *Paths*).
3. Click the unit you want to administer.
4. In the unit details, you can view relations to other `systemd` units, the status of the unit, or you can perform the following actions that can be found in the three-dot menu:
   - *Start* if the unit is not running
   - *Restart* the running unit
   - *Stop* the running unit
   - *Disallow running*—that will stop the service permanently, including all its dependencies. Keep in mind that the dependent service can be used by other units, and disallowing the unit may cause serious trouble for the system.

11.2 Creating new timers

`systemd` timers help you in automating recurring tasks. A `systemd` timer can control the triggering of `systemd` services and the handling of events.



Note: Overriding existing timers

The default set of `systemd` timers is stored in `/usr/lib/systemd`. If you create a timer with already existing names, the default unit file is not overwritten, but a new one is created in `/etc/systemd/system/` that overrides the default unit file. To restore the timer to the default one, delete the timer unit file in `/etc/systemd/system/`.

If you try to create a timer that already exists in the `/etc/systemd/system/` directory, the unit file will be overwritten, and the previous changes are lost.

To create a `systemd` timer using Cockpit, proceed as follows:

1. Navigate to *Services*.
2. In the *Timers* tab, click *Create timer*.
3. Fill in the details:

   Name

   The name of the timer that will be used in the unit name and in the service unit name as well. For example, specifying the name *example* will create the following unit files: `/etc/systemd/system/example.timer` and `/etc/systemd/system/example.service`.

   Description

   You can provide a short description of the timer.

   Command

   The command to be invoked when the timer is triggered.

   Trigger

   The timer can be triggered each time you reboot your machine or at a specific time. For the *After system boot* option, you can define the delay of the service invocation. For the *At specific time* option, specify when the service should be invoked.
4. Click *Save*.

12 Managing applications in Cockpit

The Cockpit Web interface supports modularity through plug-in applications. This section describes how to view, navigate and remove these applications using the *Applications* page.

Cockpit applications (or modules) extend the functionality of the Web interface. Adding an application creates a new entry in the sidebar, allowing you to manage complex system services—such as firewalls, virtual machines or SELinux—through a visual interface rather than the command line.

Commonly installed modules include: *Networking*, *Software Updates (Packages)*, *Repositories* and *SELinux*.

12.1 Viewing and navigating applications

You can access your installed modules by clicking the *Applications* tab in the left-hand navigation pane and clicking the application's name to switch to its dedicated management tab. To uninstall a module, locate the application in the list and click the *Remove* button.

12.2 Troubleshooting missing applications

If the *Applications* page is empty, or if a newly installed module does not appear in the list, you may need to refresh the Cockpit service.

Restart the socket: Run the following command in your terminal:

```
> sudo  systemctl restart cockpit.socket
```

Check for service failures: If a specific service fails to start, the *Applications* page will display the relevant systemd error logs directly on the screen to assist with debugging.

13 Managing packages using Cockpit

Cockpit enables you to view the list of packages available to install, reinstall and uninstall in a single Web interface.

The *Packages* tab communicates with `PackageKit` to install, uninstall, or search for packages and integrate logs.



Note: The *Packages* plug-in is not available for the immutable system.

Due to the nature of package management on immutable systems, where you need to reboot your system after installing or removing packages, this plug-in is not available on such instances.

The list of available packages appears in the *Packages* tab, and you can reinstall or uninstall the packages. You can view the package name, package version number, and a short description of the package in this tab. If an update fails because of a conflict, the *Packages* tab shows the technical details in the Web interface so that you need not manually verify the `/var/log/zypp/history`.

- To install a package by initiating zypper in the background and downloading the required packages, click the *Install* toggle button.
  Specify the package name in the *Find by name* search box. The list of packages matching the search appears.
  Click the *Install* or *Reinstall* button and then click the *Install* button again to confirm the installation.
- To uninstall a package, click the *Uninstall* toggle button and then click the *Uninstall* button corresponding to the package, and click *OK* to confirm the uninstallation of the package.

14 Managing subscriptions using Cockpit

Cockpit enables you to register new subscriptions, view the list of registered and available subscriptions and also deregister subscriptions.

14.1 Register new subscription

To register a new subscription, you must have a registration code and an e-mail address.

1. Navigate to the *Subscriptions* section.
2. In the *Register a new subscription* section, enter the details as described below.

   Registration code
   : Specify the registration code.

   Email
   : Specify the e-mail address.
3. Click *Register*.

14.2 Edit subscription settings

You can edit the subscription settings.

1. Navigate to the *Subscriptions* section.
2. Click *Edit Settings*.
3. Specify the proxy URI.
4. To allow insecure proxy addresses, select the *Allow insecure Proxies* check box.
5. Select the language code in *Language*.
6. Click *Save*.

14.3 Register available extension

Extensions that you can use are listed under *Available Extensions*. To register an extension, proceed as follows:

1. Navigate to the *Subscriptions* section.
2. In the *Available Extensions* section, click *Register* corresponding to the extension you want to register.

14.4 Deregistering subscriptions or extensions

You can deregister subscriptions or extensions using Cockpit.

1. Navigate to the *Subscriptions* section.
2. In the *Registered Subscriptions* section, click *De-register* corresponding to the subscription or extension you want to deregister.

15 SELinux mode and policy

The SELinux tool enables you to switch between SELinux modes and view current modifications of the SELinux policy.

Important: Missing SELinux module

The SELinux Cockpit module is visible only if SELinux is enabled on the system. If you cannot access the module, SELinux is probably disabled. To check that SELinux is enabled, run:

```
> sestatus
```

On SUSE Linux Enterprise Server, SELinux is in the enforcing mode by default. To temporarily switch to the permissive mode, click the button with the `Enforcing` label. Note that the change persists only until the next boot. If you need to perform a persistent change of the mode, edit the configuration file `/etc/selinux/config`.

The *System modifications* lists all modifications performed on the default SELinux policy. If you want to export the modifications and reuse them on different servers, click *View automation script*. In the new window, you can copy a shell script or an Ansible configuration file that can be applied on other servers.

15.1 Updating SLES using Cockpit

To update your system, proceed as follows:

1. Navigate to the *Software updates* page.
2. Click the *Check for Updates* icon to get a list of new package updates and patches available for your system. We recommend installing the patches marked as important as soon as possible.
3. Now you can update your system either with immediate reboot, or the reboot might be postponed:
    a. Click *Install all updates* to install all updates. By default, all available updates are selected.
    To install specific updates, select *Unselect all* and then select the required updates and click *Install selected updates*.
    After the installation is complete, you can reboot your system or click *Ignore* to postpone the reboot.

    ### Note

    When the Web Console is restarted, you will no longer see progress information. However, the update process will continue in the background. Reconnect to continue watching the update process.

b. To reboot the system after installation, enable the *Reboot after completion* toggle button. By default, the *Reboot after completion* toggle button is disabled. You must reboot the system to activate the snapshot with updates. If you perform further changes without rebooting the system beforehand, a new snapshot will be created from the same point as the snapshots with updates. Therefore, the new snapshot will not include the updates.

c. To view the installed package information, click *Package information*.

d. In the *Reboot* window, enter the message for the logged-in users and select the time delay for the reboot.

e. Click *Reboot*.

4. You can view the update history in the *Update history* section.

15.2 Managing Software Repositories using Cockpit

Cockpit enables you to view, add, edit, and delete software repositories.

15.2.1 Adding new repository

To add a new repository, you must have the repository URL.

1. Navigate to the *Software Repositories* section.
2. Click *Add repo*.
3. In the *Add repo* window, enter the repository details as described below.

Alias

Specify the repository alias.

Name

Specify a unique name for the repository.

Priority

Specify the priority number. The default priority is set to `99`.

You can set the priority within the range `1` to `99`. Where `99` denotes lowest priority and `1` the highest.

Enabled

The check box is enabled by default. To disable the repository, clear the check box.

Autorefresh

The check box is enabled by default. To disable auto-refresh for the repository, clear the check box.

GPG Check

The check box is enabled by default. To disable the GPG Check, clear the check box.

Uri

Specify the repository URI.

4. Click *Save*.

15.2.2 Refreshing repositories

You can refresh repositories manually.

1. Navigate to the *Software Repositories* section.
2. Click *Refresh repositories*.
   The repositories are refreshed, and a window appears stating that the repos refresh is successful.
3. Click *Ok*

15.2.3 Editing repositories

You can edit repositories. You cannot edit the repository URI using Cockpit.

1. Navigate to the *Software Repositories* section.
2. Click the three-dot menu and click *Edit*.
3. Modify the required details and click *Save*.

15.2.4 Deleting repositories

You can delete repositories using Cockpit.

1. Navigate to the *Software Repositories* section.
2. Click the three-dot menu and click *Delete*.
3. Click *Delete* to confirm deletion.

15.3 Solving SELinux access issues

In the *SELinux* page, you can view access denial messages from the audit log. On top of that, Cockpit provides possible ways of solving the access denial. To do so, follow the steps:

1. Navigate to the *SELinux* page.
2. In *SELinux access control errors*, expand the details regarding access denial.
3. To view the audit log record, click *Audit log*.
4. To view possible solutions, click *Solutions*. Some solutions may be applied directly through Cockpit by clicking *Apply this solution*.

16 Legal Notice

Copyright© 2006– 2026 SUSE LLC and contributors. All rights reserved.

Permission is granted to copy, distribute and/or modify this document under the terms of the GNU Free Documentation License, Version 1.2 or (at your option) version 1.3; with the Invariant Section being this copyright notice and license. A copy of the license version 1.2 is included in the section entitled “GNU Free Documentation License”.

For SUSE trademarks, see https://www.suse.com/company/legal/. All other third-party trademarks are the property of their respective owners. Trademark symbols (®, ™ etc.) denote trademarks of SUSE and its affiliates. Asterisks (*) denote third-party trademarks.

All information found in this book has been compiled with utmost attention to detail. However, this does not guarantee complete accuracy. Neither SUSE LLC, its affiliates, the authors, nor the translators shall be held liable for possible errors or the consequences thereof.

A GNU Free Documentation License

Copyright (C) 2000, 2001, 2002 Free Software Foundation, Inc. 51 Franklin St, Fifth Floor, Boston, MA 02110-1301 USA. Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

0. PREAMBLE

The purpose of this License is to make a manual, textbook, or other functional and useful document "free" in the sense of freedom: to assure everyone the effective freedom to copy and redistribute it, with or without modifying it, either commercially or non-commercially. Secondarily, this License preserves for the author and publisher a way to get credit for their work, while not being considered responsible for modifications made by others.

This License is a kind of "copyleft", which means that derivative works of the document must themselves be free in the same sense. It complements the GNU General Public License, which is a copyleft license designed for free software.

We have designed this License to use it for manuals for free software, because free software needs free documentation: a free program should come with manuals providing the same freedoms that the software does. But this License is not limited to software manuals; it can be used for any textual work, regardless of subject matter or whether it is published as a printed book. We recommend this License principally for works whose purpose is instruction or reference.

1. APPLICABILITY AND DEFINITIONS

This License applies to any manual or other work, in any medium, that contains a notice placed by the copyright holder saying it can be distributed under the terms of this License. Such a notice grants a world-wide, royalty-free license, unlimited in duration, to use that work under the conditions stated herein. The "Document", below, refers to any such manual or work. Any member of the public is a licensee, and is addressed as "you". You accept the license if you copy, modify or distribute the work in a way requiring permission under copyright law.

A "Modified Version" of the Document means any work containing the Document or a portion of it, either copied verbatim, or with modifications and/or translated into another language.

A "Secondary Section" is a named appendix or a front-matter section of the Document that deals exclusively with the relationship of the publishers or authors of the Document to the Document's overall subject (or to related matters) and contains nothing that could fall directly within that overall subject. (Thus, if the

Document is in part a textbook of mathematics, a Secondary Section may not explain any mathematics.) The relationship could be a matter of historical connection with the subject or with related matters, or of legal, commercial, philosophical, ethical or political position regarding them.

The "Invariant Sections" are certain Secondary Sections whose titles are designated, as being those of Invariant Sections, in the notice that says that the Document is released under this License. If a section does not fit the above definition of Secondary then it is not allowed to be designated as Invariant. The Document may contain zero Invariant Sections. If the Document does not identify any Invariant Sections then there are none.

The "Cover Texts" are certain short passages of text that are listed, as Front-Cover Texts or Back-Cover Texts, in the notice that says that the Document is released under this License. A Front-Cover Text may be at most 5 words, and a Back-Cover Text may be at most 25 words.

A "Transparent" copy of the Document means a machine-readable copy, represented in a format whose specification is available to the general public, that is suitable for revising the document straightforwardly with generic text editors or (for images composed of pixels) generic paint programs or (for drawings) some widely available drawing editor, and that is suitable for input to text formatters or for automatic translation to a variety of formats suitable for input to text formatters. A copy made in an otherwise Transparent file format whose markup, or absence of markup, has been arranged to thwart or discourage subsequent modification by readers is not Transparent. An image format is not Transparent if used for any substantial amount of text. A copy that is not "Transparent" is called "Opaque".

Examples of suitable formats for Transparent copies include plain ASCII without markup, Texinfo input format, LaTeX input format, SGML or XML using a publicly available DTD, and standard-conforming simple HTML, PostScript or PDF designed for human modification. Examples of transparent image formats include PNG, XCF and JPG. Opaque formats include proprietary formats that can be read and edited only by proprietary word processors, SGML or XML for which the DTD and/or processing tools are not generally available, and the machine-generated HTML, PostScript or PDF produced by some word processors for output purposes only.

The "Title Page" means, for a printed book, the title page itself, plus such following pages as are needed to hold, legibly, the material this License requires to appear in the title page. For works in formats which do not have any title page as such, "Title Page" means the text near the most prominent appearance of the work's title, preceding the beginning of the body of the text.

A section "Entitled XYZ" means a named subunit of the Document whose title either is precisely XYZ or contains XYZ in parentheses following text that translates XYZ in another language. (Here XYZ stands for a specific section name mentioned below, such as "Acknowledgements", "Dedications", "Endorsements", or "History".) To "Preserve the Title" of such a section when you modify the Document means that it remains a section "Entitled XYZ" according to this definition.

The Document may include Warranty Disclaimers next to the notice which states that this License applies to the Document. These Warranty Disclaimers are considered to be included by reference in this License, but only as regards disclaiming warranties: any other implication that these Warranty Disclaimers may have is void and has no effect on the meaning of this License.

2. VERBATIM COPYING

You may copy and distribute the Document in any medium, either commercially or non-commercially, provided that this License, the copyright notices, and the license notice saying this License applies to the Document are reproduced in all copies, and that you add no other conditions whatsoever to those of this License. You may not use technical measures to obstruct or control the reading or further copying of the copies you make or distribute. However, you may accept compensation in exchange for copies. If you distribute a large enough number of copies you must also follow the conditions in section 3.

You may also lend copies, under the same conditions stated above, and you may publicly display copies.

3. COPYING IN QUANTITY

If you publish printed copies (or copies in media that commonly have printed covers) of the Document, numbering more than 100, and the Document's license notice requires Cover Texts, you must enclose the copies in covers that carry, clearly and legibly, all these Cover Texts: Front-Cover Texts on the front cover, and Back-Cover Texts on the back cover. Both covers must also clearly and legibly identify you as the publisher of these copies. The front cover must present the full title with all words of the title equally prominent and visible. You may add other material on the covers in addition. Copying with changes limited to the covers, as long as they preserve the title of the Document and satisfy these conditions, can be treated as verbatim copying in other respects.

If the required texts for either cover are too voluminous to fit legibly, you should put the first ones listed (as many as fit reasonably) on the actual cover, and continue the rest onto adjacent pages.

If you publish or distribute Opaque copies of the Document numbering more than 100, you must either include a machine-readable Transparent copy along with each Opaque copy, or state in or with each Opaque copy a computer-network location from which the general network-using public has access to download using public-standard network protocols a complete Transparent copy of the Document, free of added material. If you use the latter option, you must take reasonably prudent steps, when you begin distribution of Opaque copies in quantity, to ensure that this Transparent copy will remain thus accessible at the stated location until at least one year after the last time you distribute an Opaque copy (directly or through your agents or retailers) of that edition to the public.

It is requested, but not required, that you contact the authors of the Document well before redistributing any large number of copies, to give them a chance to provide you with an updated version of the Document.

4. MODIFICATIONS

You may copy and distribute a Modified Version of the Document under the conditions of sections 2 and 3 above, provided that you release the Modified Version under precisely this License, with the Modified Version filling the role of the Document, thus licensing distribution and modification of the Modified Version to whoever possesses a copy of it. In addition, you must do these things in the Modified Version:

A. Use in the Title Page (and on the covers, if any) a title distinct from that of the Document, and from those of previous versions (which should, if there were any, be listed in the History section of the Document). You may use the same title as a previous version if the original publisher of that version gives permission.

B. List on the Title Page, as authors, one or more persons or entities responsible for authorship of the modifications in the Modified Version, together with at least five of the principal authors of the Document (all of its principal authors, if it has fewer than five), unless they release you from this requirement.

C. State on the Title page the name of the publisher of the Modified Version, as the publisher.

D. Preserve all the copyright notices of the Document.

E. Add an appropriate copyright notice for your modifications adjacent to the other copyright notices.

F. Include, immediately after the copyright notices, a license notice giving the public permission to use the Modified Version under the terms of this License, in the form shown in the Addendum below.

G. Preserve in that license notice the full lists of Invariant Sections and required Cover Texts given in the Document's license notice.

H. Include an unaltered copy of this License.

I. Preserve the section Entitled "History", Preserve its Title, and add to it an item stating at least the title, year, new authors, and publisher of the Modified Version as given on the Title Page. If there is no section Entitled "History" in the Document, create one stating the title, year, authors, and publisher of the Document as given on its Title Page, then add an item describing the Modified Version as stated in the previous sentence.

J. Preserve the network location, if any, given in the Document for public access to a Transparent copy of the Document, and likewise the network locations given in the Document for previous versions it was based on. These may be placed in the "History" section. You may omit a network location for a work that was published at least four years before the Document itself, or if the original publisher of the version it refers to gives permission.

K. For any section Entitled "Acknowledgements" or "Dedications", Preserve the Title of the section, and preserve in the section all the substance and tone of each of the contributor acknowledgements and/or dedications given therein.

L. Preserve all the Invariant Sections of the Document, unaltered in their text and in their titles. Section numbers or the equivalent are not considered part of the section titles.

M. Delete any section Entitled "Endorsements". Such a section may not be included in the Modified Version.

N. Do not retitle any existing section to be Entitled "Endorsements" or to conflict in title with any Invariant Section.

O. Preserve any Warranty Disclaimers.

If the Modified Version includes new front-matter sections or appendices that qualify as Secondary Sections and contain no material copied from the Document, you may at your option designate some or all of these sections as invariant. To do this, add their titles to the list of Invariant Sections in the Modified Version's license notice. These titles must be distinct from any other section titles.

You may add a section Entitled "Endorsements", provided it contains nothing but endorsements of your Modified Version by various parties--for example, statements of peer review or that the text has been approved by an organization as the authoritative definition of a standard.

You may add a passage of up to five words as a Front-Cover Text, and a passage of up to 25 words as a Back-Cover Text, to the end of the list of Cover Texts in the Modified Version. Only one passage of Front-Cover Text and one of Back-Cover Text may be added by (or through arrangements made by) any one entity. If the Document already includes a cover text for the same cover, previously added by you or by arrangement made by the same entity you are acting on behalf of, you may not add another; but you may replace the old one, on explicit permission from the previous publisher that added the old one.

The author(s) and publisher(s) of the Document do not by this License give permission to use their names for publicity for or to assert or imply endorsement of any Modified Version.

5. COMBINING DOCUMENTS

You may combine the Document with other documents released under this License, under the terms defined in section 4 above for modified versions, provided that you include in the combination all of the Invariant Sections of all of the original documents, unmodified, and list them all as Invariant Sections of your combined work in its license notice, and that you preserve all their Warranty Disclaimers.

The combined work need only contain one copy of this License, and multiple identical Invariant Sections may be replaced with a single copy. If there are multiple Invariant Sections with the same name but different contents, make the title of each such section unique by adding at the end of it, in parentheses, the name of the original author or publisher of that section if known, or else a unique number. Make the same adjustment to the section titles in the list of Invariant Sections in the license notice of the combined work.

In the combination, you must combine any sections Entitled "History" in the various original documents, forming one section Entitled "History"; likewise combine any sections Entitled "Acknowledgements", and any sections Entitled "Dedications". You must delete all sections Entitled "Endorsements".

6. COLLECTIONS OF DOCUMENTS

You may make a collection consisting of the Document and other documents released under this License, and replace the individual copies of this License in the various documents with a single copy that is included in the collection, provided that you follow the rules of this License for verbatim copying of each of the documents in all other respects.

You may extract a single document from such a collection, and distribute it individually under this License, provided you insert a copy of this License into the extracted document, and follow this License in all other respects regarding verbatim copying of that document.

7. AGGREGATION WITH INDEPENDENT WORKS

A compilation of the Document or its derivatives with other separate and independent documents or works, in or on a volume of a storage or distribution medium, is called an "aggregate" if the copyright resulting from the compilation is not used to limit the legal rights of the compilation's users beyond what the individual works permit. When the Document is included in an aggregate, this License does not apply to the other works in the aggregate which are not themselves derivative works of the Document.

If the Cover Text requirement of section 3 is applicable to these copies of the Document, then if the Document is less than one half of the entire aggregate, the Document's Cover Texts may be placed on covers that bracket the Document within the aggregate, or the electronic equivalent of covers if the Document is in electronic form. Otherwise they must appear on printed covers that bracket the whole aggregate.

8. TRANSLATION

Translation is considered a kind of modification, so you may distribute translations of the Document under the terms of section 4. Replacing Invariant Sections with translations requires special permission from their copyright holders, but you may include translations of some or all Invariant Sections in addition to the original versions of these Invariant Sections. You may include a translation of this License, and all the license notices in the Document, and any Warranty Disclaimers, provided that you also include the original English version of this License and the original versions of those notices and disclaimers. In case of a disagreement between the translation and the original version of this License or a notice or disclaimer, the original version will prevail.

If a section in the Document is Entitled "Acknowledgements", "Dedications", or "History", the requirement (section 4) to Preserve its Title (section 1) will typically require changing the actual title.

9. TERMINATION

You may not copy, modify, sublicense, or distribute the Document except as expressly provided for under this License. Any other attempt to copy, modify, sublicense or distribute the Document is void, and will automatically terminate your rights under this License. However, parties who have received copies, or rights, from you under this License will not have their licenses terminated so long as such parties remain in full compliance.

10. FUTURE REVISIONS OF THIS LICENSE

The Free Software Foundation may publish new, revised versions of the GNU Free Documentation License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns. See https://www.gnu.org/copyleft/.

Each version of the License is given a distinguishing version number. If the Document specifies that a particular numbered version of this License "or any later version" applies to it, you have the option of following the terms and conditions either of that specified version or of any later version that has been published (not as a draft) by the Free Software Foundation. If the Document does not specify a version number of this License, you may choose any version ever published (not as a draft) by the Free Software Foundation.

ADDENDUM: How to use this License for your documents

```
Copyright (c) YEAR YOUR NAME.
```

```
Permission is granted to copy, distribute and/or modify this document
under the terms of the GNU Free Documentation License, Version 1.2
or any later version published by the Free Software Foundation;
with no Invariant Sections, no Front-Cover Texts, and no Back-Cover Texts.
A copy of the license is included in the section entitled “GNU
Free Documentation License”.
```

If you have Invariant Sections, Front-Cover Texts and Back-Cover Texts, replace the “with...Texts.” line with this:

```
with the Invariant Sections being LIST THEIR TITLES, with the
Front-Cover Texts being LIST, and with the Back-Cover Texts being LIST.
```

If you have Invariant Sections without Cover Texts, or some other combination of the three, merge those two alternatives to suit the situation.

If your document contains nontrivial examples of program code, we recommend releasing these examples in parallel under your choice of free software license, such as the GNU General Public License, to permit their use in free software.