

Cockpit Guide

This guide describes the administration of SUSE Linux Enterprise Micro using Cockpit.

Publication Date: December 18, 2025

Contents

1	About Cockpit	2
2	Getting Cockpit	2
3	Cockpit authentication	3
4	System overview	5
5	SLE Micro logs	6
6	Managing storage	8
7	Managing networking	16
8	Working with containers	26
9	Users administration	31
10	Managing services	33
11	Managing SELinux mode and policy	35
12	Managing SLE Micro updates and snapshots	35
A	GNU licenses	37

1 About Cockpit

Cockpit is a web-based graphical interface that enables you to manage most of the administration tasks from one place by using plugins. 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
- update the server
- manage the server storage
- inspect and change network settings
- manage user accounts
- view system logs
- create and administer virtual machines
- inspect and interact with systemd services
- switch between SELinux modes
- use a terminal on a remote server in your web browser

2 Getting Cockpit

Cockpit is included in the delivered pre-built images, or can be installed if you are installing your own instances manually. For details regarding the manual installation, refer to *Book "Deployment Guide", Chapter 12 "Installation steps", Section 12.9.2 "Software"*.

Though Cockpit is present in the pre-built images by default, the plugin for administration of virtual machines needs to be installed manually. You can do so by installing the microos-cockpit pattern as described below. You can also use the following command in case Cockpit is not installed on your system.

```
# transactional-update pkg install -t pattern microos-cockpit
```

Reboot your machine to switch to the latest snapshot.



Note: Cockpit's plugins installed from the `microos-cockpit` pattern may differ according to technologies installed on your system

The plugin `Podman containers` is installed only if the *Containers Runtime for non-clustered systems* patterns are installed on your system. Similarly, the `Virtual Machines` plugin is installed only if the *KVM Virtualization Host* pattern is installed on your system.

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

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

After running the command, the server will expose the default `9090` port. `systemd` will start the `cockpit-ws` service that listens on the `9090` port.

In case you have enabled the firewall, you must also open the firewall for Cockpit as follows:

```
# firewall-cmd --permanent --zone=public --add-service=cockpit
```

And then reload the firewall configuration by running:

```
# firewall-cmd --reload
```

Now you can access the Cockpit web interface by opening the following address in your web browser:

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

3 Cockpit authentication

3.1 Authentication

Cockpit enables you to log in directly to each machine that can expose the `9090` port. 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 a procedure of adding a server as secondary, refer to [Section 3.1.1, “Adding secondary servers”](#).



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.

To log in to the primary server, enter your credentials on the Cockpit login page. To log in to secondary servers, there are two options. You can access secondary machines after logging in to the primary server, or you can access a secondary machine using the *Connect to* field on the login page.

3.1.1 Adding secondary servers

To add a server as secondary, proceed as follows:

1. Log in to the primary server using the account with the *system administrator* role.
2. Click on 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 a fingerprint on the server you want to add. If the fingerprint matches or if you have not set up the SSH connection, click *Accept and connect* 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.

3.1.2 Logging into 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, expand *Other options* on the login screen of the primary server. Fill in your user name, password and the IP address of the secondary server. If you are trying to log in for the first time, you are asked to verify the fingerprint. After this, click *Accept and connect*.

4 System overview

The *Overview* page provides information about the health of your server, CPU and memory usage and also hardware details. On this page, you can also change the host name, configure the system time and view SSH keys.

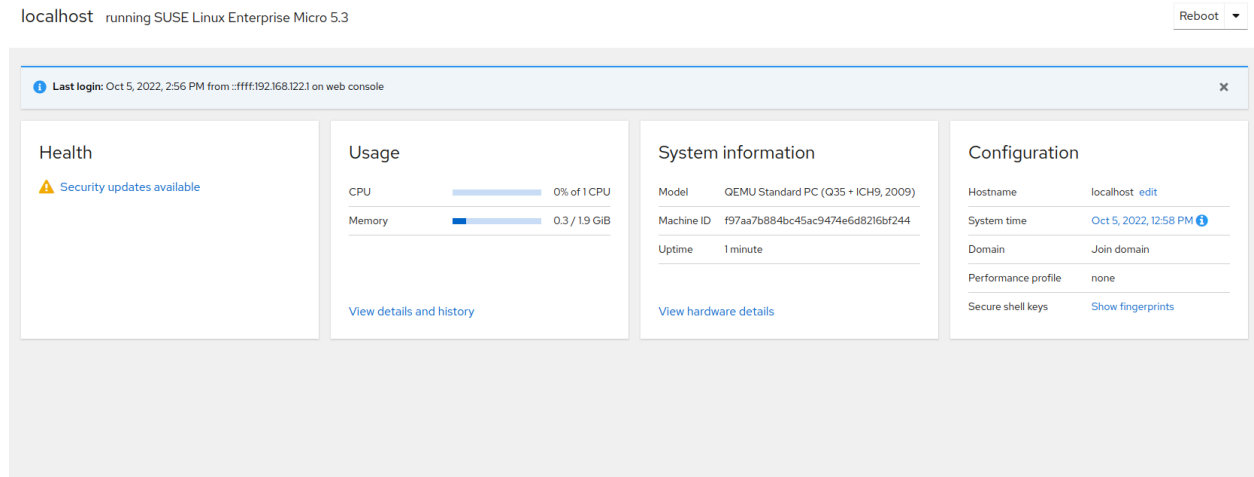


FIGURE 1: SYSTEM OVERVIEW

The *Health* part displays if your system is up to date or if there are any updates to install. Clicking the displayed statement takes you to *Software Updates*. For details, refer to [Section 12, “Managing SLE Micro updates and snapshots”](#).

The *Usage* part displays CPU and memory usage. For detailed performance metrics, click *View details and history*. Clicking on mount points in *Disks* takes you to *Storage*. Clicking on a particular service opens the service details under *Services*.

In the *Configuration* part, you can perform the following tasks:

- change the host name by clicking *edit* and entering the required host name
- change the system time by clicking on the date and time value. In the popup window, select a time zone or set the time, either manually or by synchronizing it with an NTP server.
- view your SSH keys by clicking *Show fingerprints*

5 SLE Micro logs

After clicking *Logs*, Cockpit enables you to view logs on your system. You can filter the logs according to the following criteria:

Time

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

- *Current boot*
- *Previous boot*
- *Last 24 hours*
- *Last 7 days*

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 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 actions are required.

Info and above

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

Debug and above

These messages are used just to debug the system.

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.

Filters

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

Since

Logs on the specified date or newer will be displayed. You can specify the time 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 on the specified date or older will be displayed. You can specify the time 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 a systemd unit which logs you want to display. Use one of the formats:

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

Free-form search

Enter a 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) (<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

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

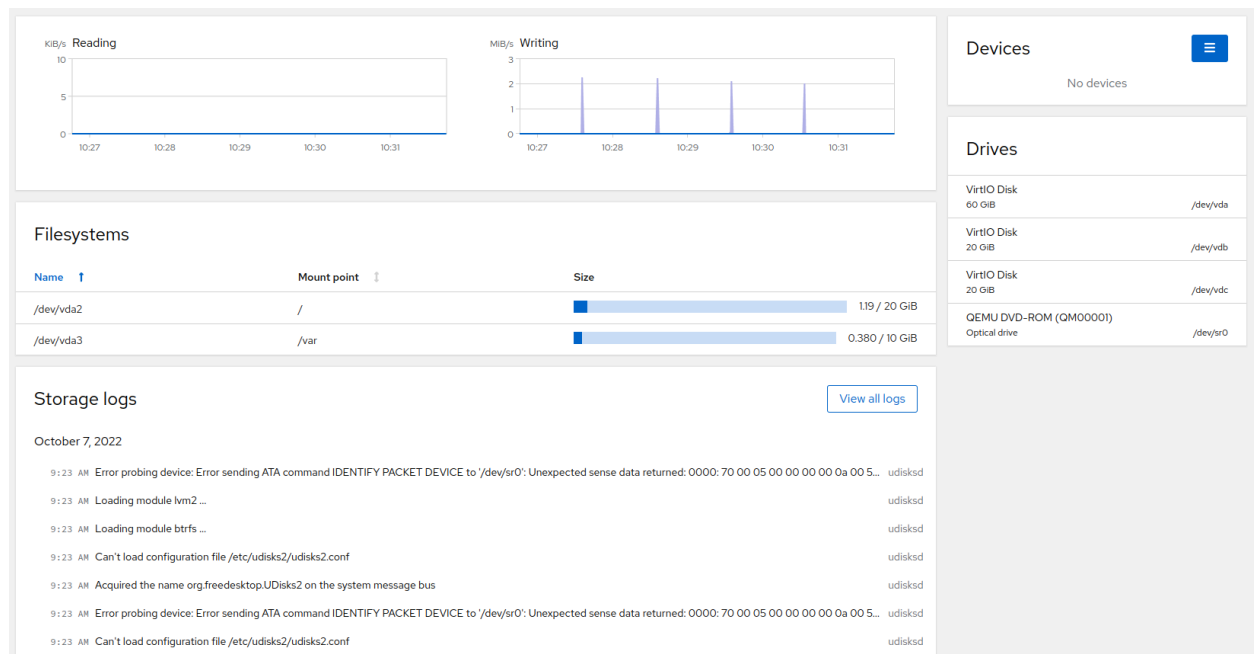


FIGURE 2: STORAGE VIEW

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 3: 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*. Click on the particular file system to open the details window.

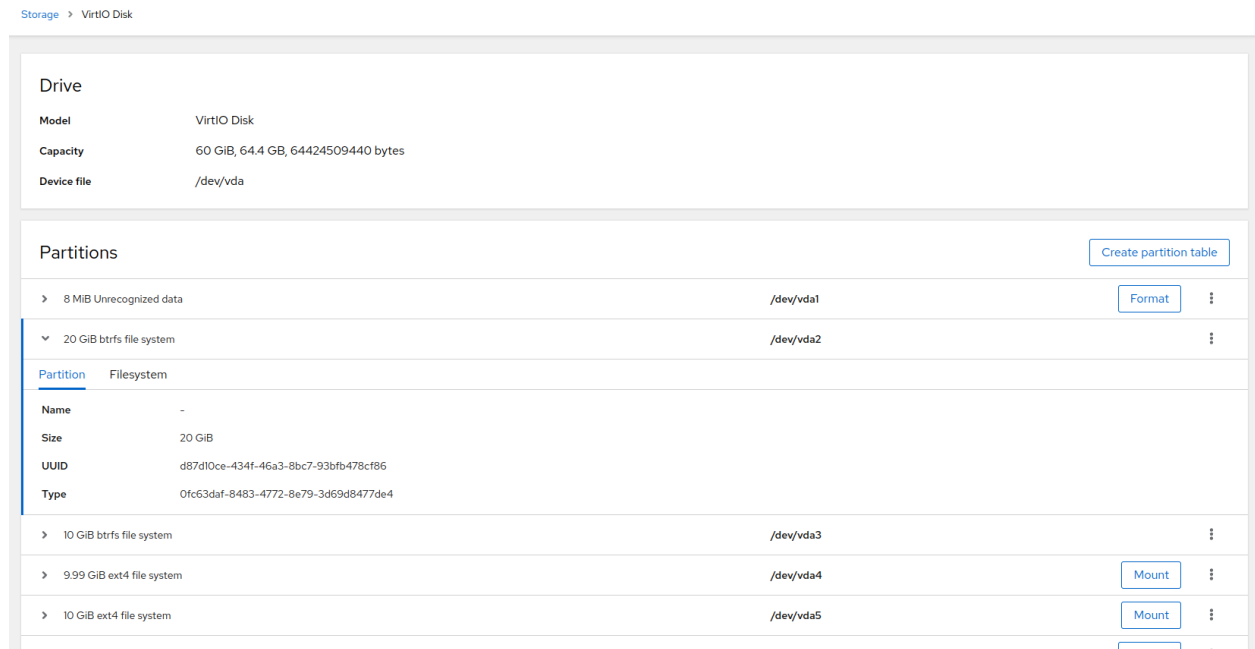


FIGURE 4: FILE SYSTEMS DETAILS

To format the partition, click *Format* next to the particular partition description to open the format window.

Format /dev/vda1

Name

Type BTRFS (recommended) ▼

Erase ☒ Overwrite existing data with zeros

Mount point

Mount options ☒ Mount now
☐ Mount read only
☒ Custom mount options

Encryption LUKS2 (recommended) ▼

Passphrase

Confirm

☐ Store passphrase

⚠ Formatting a storage device will erase all data on it.

Format Cancel

FIGURE 5: **FORMATTING PARTITIONS**

You can specify the following details:

FORMATTING PARTITIONS PARAMETERS

Name

Enter a unique name of the partition.

Type

Select the file system. BTRFS is mandatory for the `/` partition.

Mount point

This is a mandatory option. Specify to which directory the partition will be mounted.

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 \(https://linux.die.net/man/8/mount\)](https://linux.die.net/man/8/mount). Those options are used in the options part of the /etc/fstab file.

Encryption

Select the encryption method or do not encrypt the partition. If you choose any of the encryption methods, the following fields appear:

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.

Custom encryption options

You can pass a list of options described in [supported encrypted options \(https://www.man7.org/linux/man-pages/man5/crypttab.5.html#SUPPORTED_OPTIONS\)](https://www.man7.org/linux/man-pages/man5/crypttab.5.html#SUPPORTED_OPTIONS).

After you specify all required details, click *Format* to proceed.

If a partition is already formatted, you can mount it by clicking *Mount*. Here you can change the mount point and mount options. To proceed, click *Mount*.

6.3 Managing NFS mount points

Use the *Storage* view to add, edit or delete NFS mounts.

To add an NFS mount point, after clicking on the plus icon, specify the following details:

NFS MOUNT DETAILS

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 to.

Mount options

Check any of the options:

- *Mount at boot* – to mount the path automatically after each system start or restart.
- *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.

To edit an NFS mount, click on the particular NFS mount. On the next screen, click *Edit* and specify the details described in [NFS mount details](#).

6.4 Software RAID management

The *Devices* view enables you to add or manage RAIDs.

6.4.1 Creating RAIDs

You can create RAIDs of different levels on your system from the *Devices* part by selecting the *Create RAID device* option in the hamburger menu.

Create RAID device

Name	raid0							
RAID level	RAID 5 (distributed parity) ▼							
Chunk size	512 KiB ▼							
Disks	<table><tr><td><input checked="" type="checkbox"/></td><td>20 GiB LVM PV ovZ5xM-XCGi-kM5r-HQ5y-wbjt-UYxt-dZZYr9 on /dev/vdb</td><td>/dev/vdb</td></tr><tr><td><input checked="" type="checkbox"/></td><td>20 GiB VirtIO Disk</td><td>/dev/vdc</td></tr></table>		<input checked="" type="checkbox"/>	20 GiB LVM PV ovZ5xM-XCGi-kM5r-HQ5y-wbjt-UYxt-dZZYr9 on /dev/vdb	/dev/vdb	<input checked="" type="checkbox"/>	20 GiB VirtIO Disk	/dev/vdc
<input checked="" type="checkbox"/>	20 GiB LVM PV ovZ5xM-XCGi-kM5r-HQ5y-wbjt-UYxt-dZZYr9 on /dev/vdb	/dev/vdb						
<input checked="" type="checkbox"/>	20 GiB VirtIO Disk	/dev/vdc						
<div>Create Cancel</div>								

FIGURE 6: CREATING RAID DEVICE

Enter the following parameters of the RAID:

Name

Enter a unique name of the RAID.

RAID level

Select one of the following RAID levels:

RAID 0

has the data split evenly to two or more disks. Such a solution provides higher performance because you can read data from several disks simultaneously, but does not provide any fault tolerance. You need to have at least 2 disks.

RAID 1

is a mirror. Data are kept in exact copies on several disks. If there is a failure on one disk, you can use data on another disk. Typically, this RAID level uses 2 disks.

RAID 4

provides striping and a dedicated disk for parity. You need to have at least 3 disks.

RAID 5

provides striping, the parity data are distributed on all disks. The minimum number of disks is 3.

RAID 6

provides striping and two parity blocks distributed across disks. The minimum number of disks is 4.

RAID 10

is a combination of striping and mirroring. Each stripe is mirrored to another disk. The minimum number of disks is 4.

Chunk size

The size of chunks. 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.

Confirm the parameters by clicking *Create*. The RAID then appears in the *Devices* part.

6.4.2 Modifying RAIDS

Click the RAID in *Devices* to open the RAID details view. In the detailed view, you can stop or delete the RAID, add or remove disks and format the device.

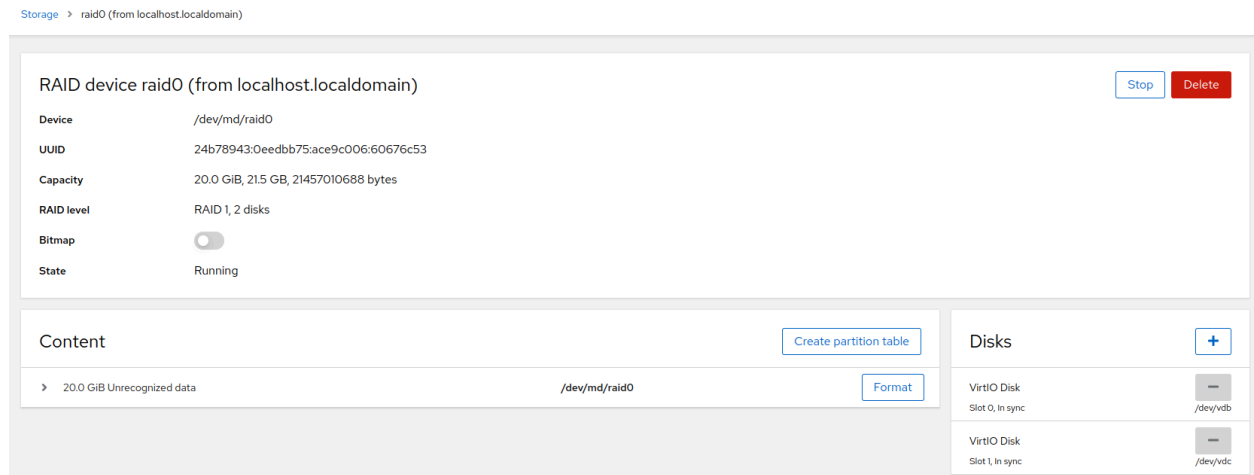


FIGURE 7: RAID DETAILS

With some 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.

You can remove and add disks in the *Disks* part. After any change to the array, the system undergoes resynchronization that might 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.

The *Content* view enables you to create a partition table (either MBR or GPT), format the disks and mount the RAID.

6.5 Logical volume management

Using Cockpit, you can create thinly provisioned logical volumes or logical block devices. To do so, create a volume group of disks by selecting *Create volume group* in the hamburger menu in *Devices*.

Create volume group

Name	<input type="text" value="vgroup0"/>	
Disks	<input type="checkbox"/> 20 GiB VirtIO Disk	/dev/vdb
	<input type="checkbox"/> 20 GiB VirtIO Disk	/dev/vdc
	<input type="checkbox"/> 20 GiB VirtIO Disk	/dev/vdd
<div><div>Create</div><div>Cancel</div></div>		

FIGURE 8: CREATING A VOLUME GROUP

Enter the volume group name and select disks that will be part of the volume group. Confirm the data with *Create*. The volume group appears in the *Devices* view. Now you can create a logical volume (either a block device or a pool of thinly provisioned volumes). To create a logical volume, proceed as follows:

PROCEDURE 1: CREATING A LOGICAL VOLUME

1. Click on the volume group in *Devices*.
2. In the volume group details, click *Create new logical volume*.

Create logical volume

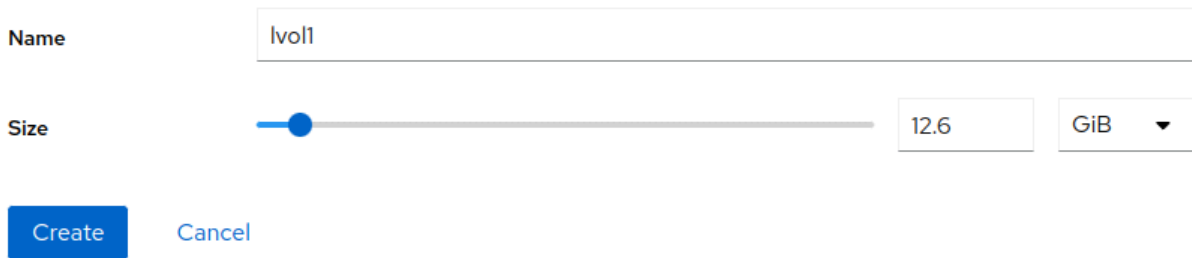
Name	<input type="text" value="lvol0"/>	
Purpose	<div>Block device for filesystems</div>	
Size	<div><div></div><div>60.0</div><div>GiB</div></div>	
<div><div>Create</div><div>Cancel</div></div>		

3. Specify a logical volume name, select either a block device or a pool of thinly provisioned volumes and select the size to use.
4. Click *Create* to confirm the details.

If you created a block logical volume, you need to format the volume by clicking *Format* and filling the details as described in [Formatting partitions parameters](#).

If you created a pool of thinly provisioned volumes, you can create a thin volume by clicking *Create thin volume*.

Create thin volume



The screenshot shows a 'Create thin volume' dialog box. It has a 'Name' field containing 'lvoll'. Below it is a 'Size' section with a slider bar and a numeric input field showing '12.6'. To the right of the slider is a unit dropdown menu set to 'GiB'. At the bottom left are two buttons: 'Create' (highlighted in blue) and 'Cancel'.

FIGURE 9: CREATING A THINLY PROVISIONED VOLUME

Enter a unique name, select the size of the volume, and confirm by clicking *Create*. You can create several volumes of the particular volume group by clicking *Create thin volume* again. Format the volumes by clicking *Format* and filling the details as described in [Formatting partitions parameters](#). Other actions on a particular volume are available in the three dots menu: *Deactivate*, *Create snapshot* or *Delete*.

Shrinking or growing a volume is available by clicking on *Shrink* or *Grow* in the particular volume details. These features are not available for all file systems and volume types.

7 Managing networking

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

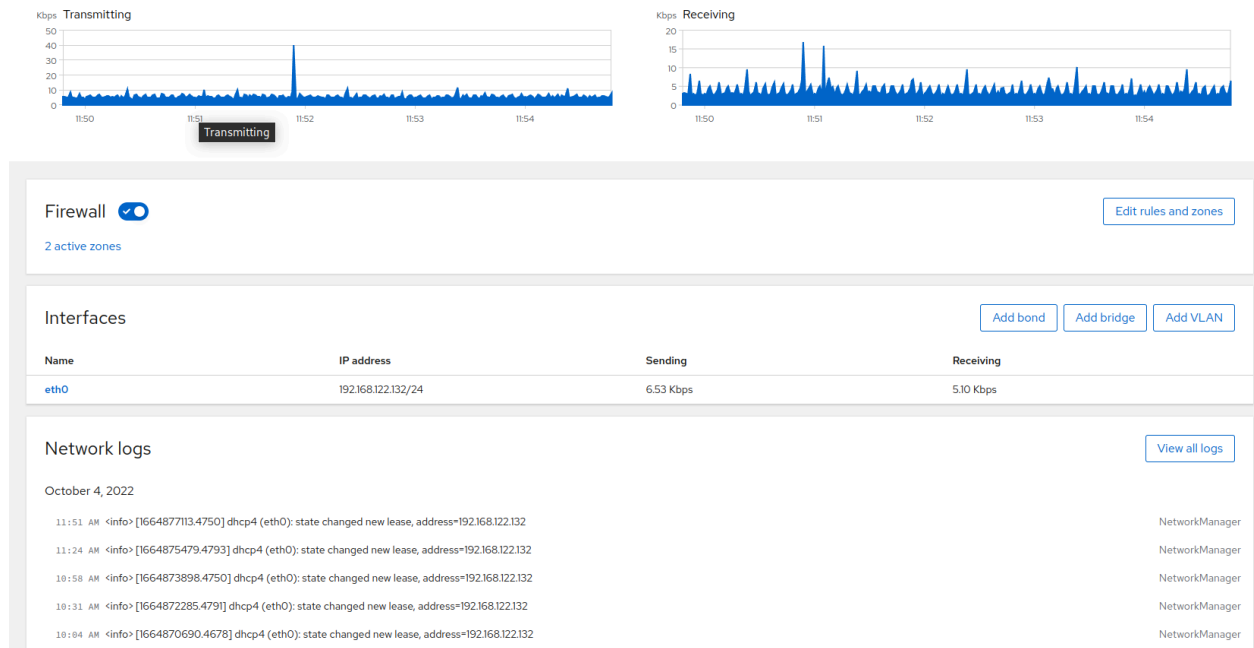


FIGURE 10: NETWORKING

7.1 Firewall



Important: Firewall and Podman

Using firewall along with Podman may result in missing Podman-related firewall rules after reloading the firewalld service. Therefore, it is recommended to keep the firewall in its default setting (disabled) if you intend to use Podman.



Note: The cockpit service is not included in the *public zone*

If you turn on the firewall, add the `cockpit` service to the *public zone*, otherwise the Cockpit web UI might get disconnected.

The *Firewall* view enables you turn on the firewall, which is off by default. After turning it on, you can edit firewall rules and zones by clicking *Edit rules and zones*.

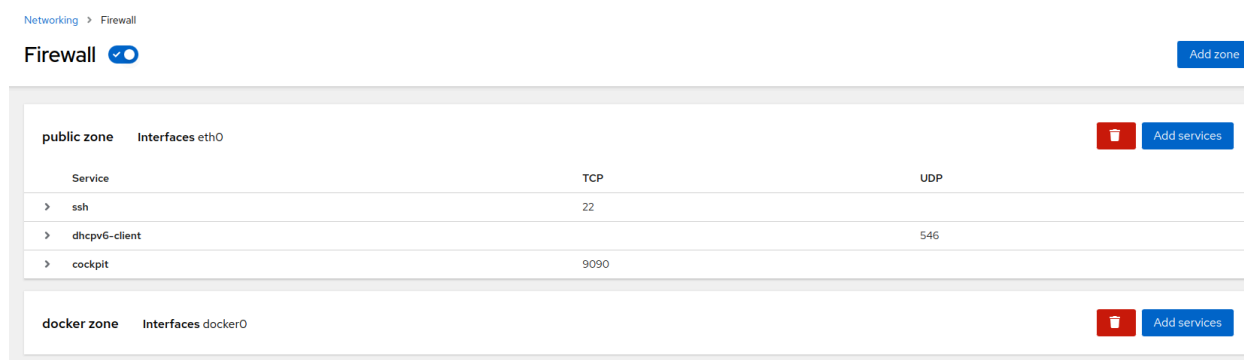


FIGURE 11: FIREWALL SETTINGS

To add a service to an existing zone, click *Add services*. Then select a service from the list and confirm with *Add services*.

If you need to define a zone on top of the currently available zones, click *Add zone*. In the *Add zone* window, 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). You can also allow addresses from the entire subnet within the particular trust level, or you can define a specific range of allowed IP addresses as a comma-separated list.

Add zone

x

Trust level

Sorted from least to most trusted

☐ External

☐ Dmz

☐ Work

☐ Home

☐ Internal

Description

No description available

Included services

None

The cockpit service is automatically included

Interfaces

Allowed addresses

☐ Entire subnet

☒ Range

IP address with routing prefix. Separate multiple values with a comma. Example: 192.0.2.0/24, 2001:db8::/32

Add zone

Cancel

FIGURE 12: ADD FIREWALL ZONE

7.2 Interfaces

In the *Interfaces* part, you can add, modify or delete VLANs, bridges or bond interfaces.

7.2.1 Managing bonds

A bond interface is a combination of several network interfaces into one bond. Depending on the *Mode* (described further), network bonding can increase 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. To add a bond interface, click *Add bond*.

Bond settings

?

×

Name

bond0

Interfaces

☐

eth0

MAC

▼

Mode

Active backup

▼

Primary

-

▼

Link monitoring

MII (recommended)

▼

Monitoring interval

100

Link up delay

0

Link down delay

0

Apply

Cancel

FIGURE 13: **ADDING A BOND**

Specify the following parameters of the bond interface:

Name

Enter a unique name of 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.

Balance 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.

802.3ad Dynamic Link Aggregation

Creates aggregation groups that share the same speed and duplex settings.

Adaptive Transmit Load Balancing

A 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 for 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.

To modify or delete a bond interface, click on the particular bond name to open the details. There you can change the parameters of the bond described below.



FIGURE 14: MODIFYING A BOND

Bond

Select a MAC address from the list.

Connect automatically

The bond connects automatically by default. Uncheck 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

After clicking *edit*, you can specify a particular value of the maximum transmission unit in bytes.

Bond

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

7.2.2 Managing bridges

A network bridge is a device that creates a single aggregated network from multiple networks. To create a bridge, click *Add bridge*.

Bridge settings ✕

Name

bridge0

Ports

☐ eth0

Spanning tree protocol (STP)

☒

STP priority

32768

STP forward delay

15

STP hello time

2

STP maximum message age

20

Apply

Cancel

FIGURE 15: ADDING A BRIDGE

Specify the following:

Name

Specify a unique name of the bridge.

Ports

Select 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 \(https://en.wikipedia.org/wiki/Spanning_Tree_Protocol\)](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 to become 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 (BPDU) 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.

To modify or delete a bridge, click the bridge name to open the details.




FIGURE 16: MODIFYING A BRIDGE

There you can edit the following details:

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.

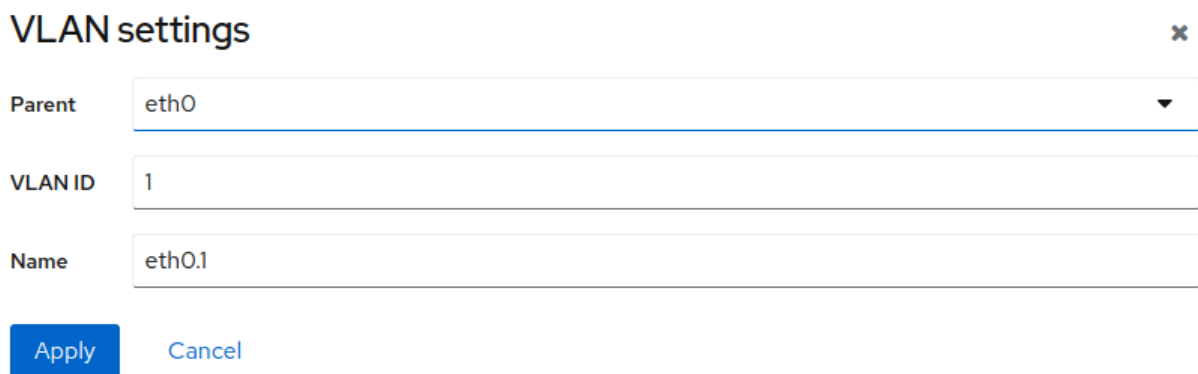
Bridge

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

7.2.3 Managing VLANs

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

To create a VLAN, click *Add VLAN* and select the parent interface, assign an ID number and provide a name for the VLAN.



VLAN settings ✕

Parent

VLAN ID

Name

FIGURE 17: **ADDING VLAN**

To modify or delete a VLAN, click the VLAN name.



FIGURE 18: ADDING VLAN

8 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 container. You can also filter the view by entering a filter criterion into the filter field.

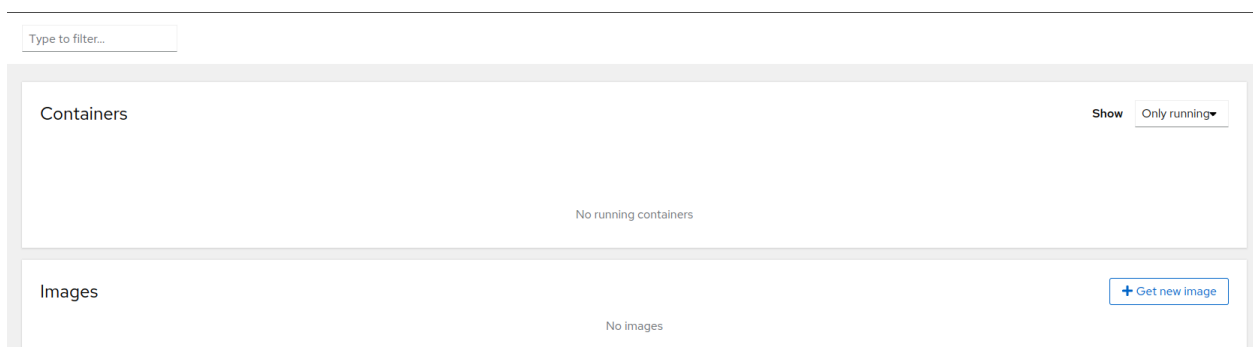


FIGURE 19: CONTAINERS

8.1 Getting container images

To start a container, you need a container image. Click *Get new image* to open a search box for images.

In the search box:

- choose a preferred image registry or proceed with All registries
- define the *Tag*. The default value is latest
- fill in the image name or description

Cockpit suggests possible images according to the entered name, registry and tag. Select the desired image and click *Download*.



Search for an image ×

Search for in

- registry.opensuse.org/devel/bci/sle-15-sp3/containerfile/suse/pcp
- registry.opensuse.org/devel/bci/sle-15-sp4/containerfile/suse/pcp
- registry.opensuse.org/devel/bci/tumbleweed/containerfile/opensuse/pcp
- registry.opensuse.org/home/dancermak/bci/sle-15-sp4/containerfile/suse/pcp

Tag

FIGURE 20: GETTING IMAGES



Note: Adding custom registry

SLE Micro comes with a predefined set of registries. If you need to use a custom registry, configure Cockpit to display them. Add the custom registry to the file `/etc/containers/registries.conf` to the `registries` list. For example, you can add *registry.example* to the list as follows:

```
registries = ["registry.opensuse.org", "docker.io", "registry.example"]
```



Tip: Pulling images using Podman

You can also get container images using Podman by running the `podman pull` command. For details, refer to Article “Podman Guide”.

8.2 Managing containers

After you get an image, you can start a container based on that image by clicking the button next to the image information.

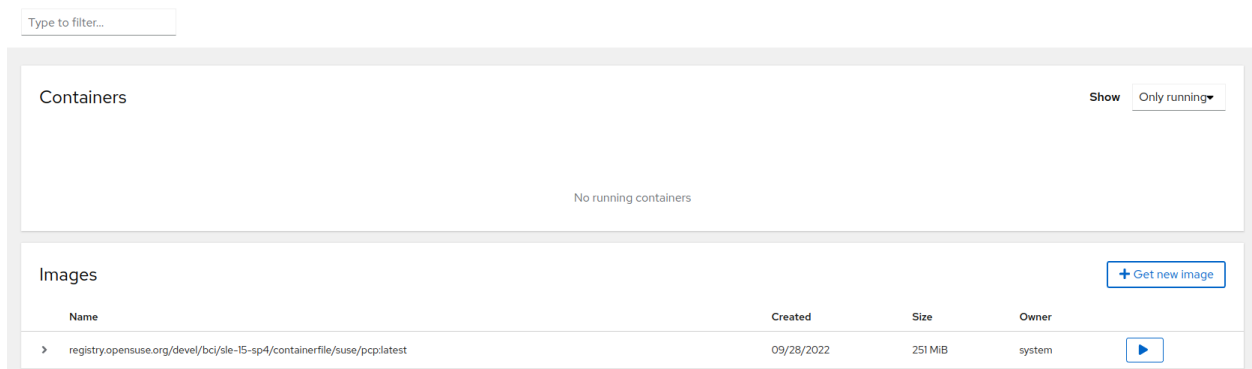


FIGURE 21: **RUNNING A CONTAINER**

In the *Run image* window, enter the container details as described below.

Run image



Image	registry.opensuse.org/devel/bci/sle-15-sp4/containerfile/suse/pcp:latest		
Name	<input type="text" value="stoic_hypatia"/>		
Command	<input type="text" value="/usr/lib/systemd/systemd"/>		
Memory limit	<input type="checkbox"/>	<input type="text" value="512"/>	<div>MiB ▼</div>
CPU shares	<input type="checkbox"/>	<input type="text"/>	
	<input checked="" type="checkbox"/> With terminal		
Ports	<div><input type="text" value="IP (optional)"/> <input type="text" value="Container port"/></div>	<div><input type="text" value="Host port (optional)"/> <div>TCP ▼</div></div>	<div><div>+</div><div>+</div></div>
Volumes	<div><div>Host path ▼</div><div>ReadWrite ▼</div></div>	<div><div>Container path</div><div>No SELinux label ▼</div></div>	<div><div>+</div></div>
Environment	<div><input type="text" value="Key"/></div>	<div><input type="text" value="Value"/></div>	<div><div>+</div></div>
<div><div>Run</div><div>Cancel</div></div>			

FIGURE 22: CONTAINER DETAILS

Name

You can define a custom name for the container. If not specified, the generated value is used.

Command

Defines a command that is run inside the container, for example, `/usr/lib/systemd/systemd` or `/bin/bash`.

Memory limit

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

With terminal

The container will not run in the detached mode.

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 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.

Ports

You can publish ports of the container to the host. If you do not provide the host IP address, the port will be bound to all host IP addresses. If you omit the host port number, a random port is assigned.

To publish a port for both protocols (TCP and UDP), enter the same port number twice, just select different protocols.

Volumes

Maps the volumes on the host system to the paths in the container. The Container path must be an absolute path inside the container. If it does not exist, the directory will be created.

Environment

Enables passing environment variables to the container.

After expanding the container details, you can perform the following tasks:

- access a terminal of the container in the *Console* tab
- view the container logs in the *Logs* tab
- delete the container by clicking the trash bin button
- commit changes performed to the container, for example, installing packages to the container
- start the container by clicking *Start*

- 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 might 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.

9 Users administration



Note: Users administration only for server administrators

Only users with the *Server administrator* role can edit other users.

Cockpit enables you to manage the users of your system. Click *Accounts* to open the user administration page. Here you can create a new account by clicking *Create new account* or manage already existing accounts by clicking on the particular account.




FIGURE 23: THE ACCOUNTS SCREEN

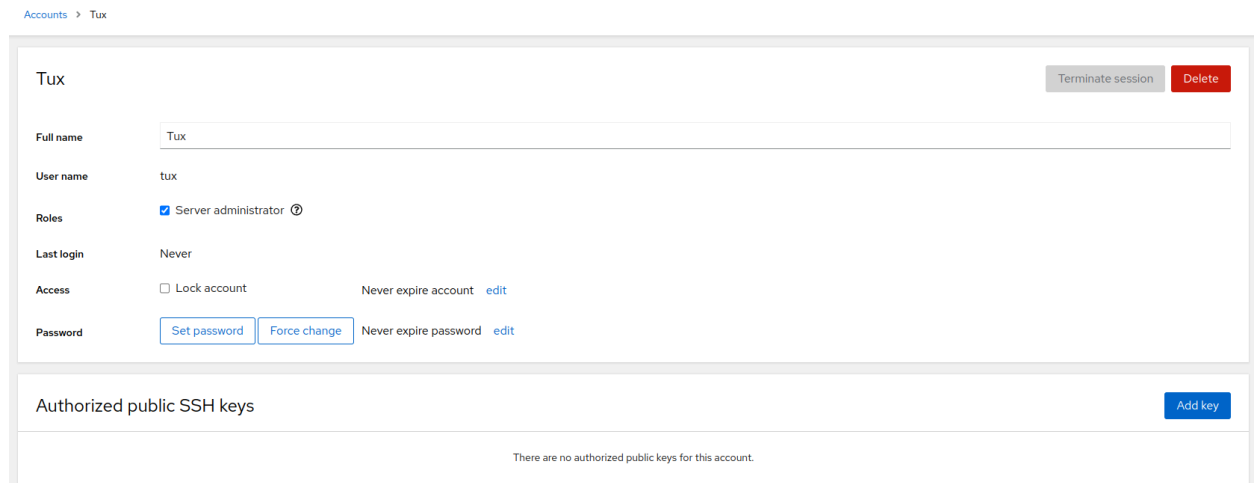
9.1 Creating new accounts

Click *Create new account* to open the window that enables you to add a new user. Fill in the user's login and/or full name and password, then confirm the form by clicking *Create*.

To add authorized SSH keys for the new user or set the *Server administrator* role, edit the already created account by clicking on it. For details, refer to [Section 9.2, “Modifying accounts”](#).

9.2 Modifying accounts

After clicking the user icon in the *Accounts* page, the user details view opens, and you can edit the user.



The screenshot shows the user details view for a user named 'Tux'. At the top left, there is a breadcrumb 'Accounts > Tux'. The user's name 'Tux' is displayed at the top left of the main content area. In the top right corner, there are two buttons: 'Terminate session' (grey) and 'Delete' (red). Below the name, there are several fields and controls: 'Full name' with a text input containing 'Tux'; 'User name' with a text input containing 'tux'; 'Roles' with a checked checkbox for 'Server administrator' and a help icon; 'Last login' with the value 'Never'; 'Access' with a checkbox for 'Lock account', a text input for 'Never expire account', and an 'edit' link; and 'Password' with buttons for 'Set password' and 'Force change', a text input for 'Never expire password', and an 'edit' link. At the bottom, there is a section for 'Authorized public SSH keys' with an 'Add key' button. A message at the very bottom states 'There are no authorized public keys for this account.'

FIGURE 24: USER DETAILS

In the user's 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 out a particular user out of the system.

Change user's role

By checking/unchecking the *Server administrator* check box, you can assign or remove the administrator role from the user.

Manage access to the account

You can lock the account, or you can set a date when the account will expire.

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.

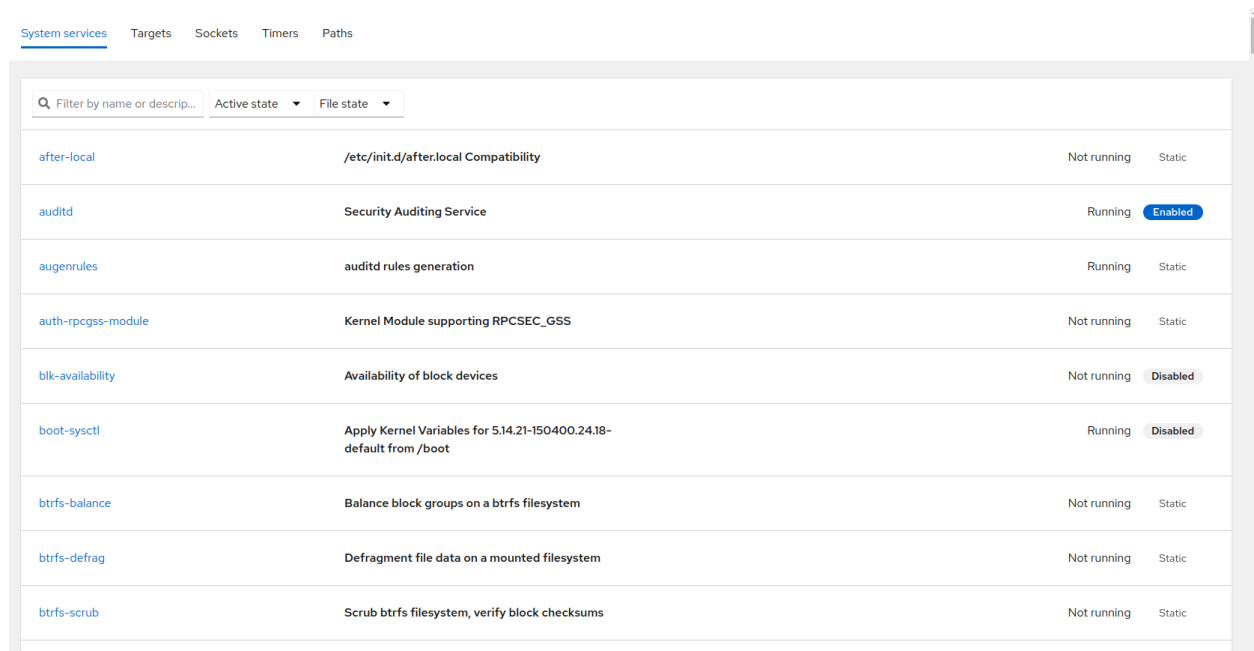
Click *edit* to set whether or when the password expires.

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*.

10 Managing services

In the *Services* view, you can view the status of systemd units, start or stop particular units and create systemd timers.



The screenshot shows a web interface for managing system services. At the top, there are tabs: 'System services' (selected), 'Targets', 'Sockets', 'Timers', and 'Paths'. Below the tabs is a search bar with the text 'Filter by name or descrip...' and two dropdown menus for 'Active state' and 'File state'. The main content is a table listing various systemd units. Each row includes the unit name, a description, the current state, and a button to toggle the state.

Unit Name	Description	Active State	File State
after-local	/etc/init.d/after.local Compatibility	Not running	Static
auditd	Security Auditing Service	Running	Enabled
augenrules	auditd rules generation	Running	Static
auth-rpcgss-module	Kernel Module supporting RPCSEC_GSS	Not running	Static
blk-availability	Availability of block devices	Not running	Disabled
boot-sysctl	Apply Kernel Variables for 5.14.21-150400.24.18-default from /boot	Running	Disabled
btrfs-balance	Balance block groups on a btrfs filesystem	Not running	Static
btrfs-defrag	Defragment file data on a mounted filesystem	Not running	Static
btrfs-scrub	Scrub btrfs filesystem, verify block checksums	Not running	Static

FIGURE 25: SERVICES MANAGEMENT

To manage a systemd unit, select the appropriate tab, then click on the particular unit. 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 dots 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 might cause serious troubles for the system.

10.1 Creating new timers

To create a new timer, click *Create timer* in the *Timers* tab. Fill in the details described below in the opened *Create timer* window.

Create timer

x

Name

Description

Command

Trigger

☐ After system boot
☒ At specific time

Repeat

Don't repeat ▼

Run at

00:00 ⌚

Save

Cancel

FIGURE 26: CREATING TIMERS

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.



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. If you want 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.

11 Managing SELinux mode and policy

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




FIGURE 27: SELINUX TOOL

On SLE Micro, SELinux is in the permissive mode by default. To temporarily switch to the enforcing mode, click the button with the Permissive label. Bear in mind that the change persists only until the next boot. If you need to perform persistent change of the mode, edit the configuration file `/etc/selinux/config` as described in *Book "Security and Hardening Guide", Chapter 1 "SELinux", Section 1.2 "SELinux modes"*.

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 the ansible configuration file that can be applied on other servers.

12 Managing SLE Micro updates and snapshots



Important: No system updates without registering your system

If your system is not registered, the updates are not available and the check for updates fails. Therefore, register your system to view available updates. For details, refer to *Book "Deployment Guide", Chapter 10 "Post-deployment steps"*.



Note: Snapshots and updates management only for system administrators

Only users with the *Server administrator* role can update the system or perform a rollback to another snapshot.

Cockpit enables you to update your SLE Micro instance or perform a rollback from the Software Updates menu.

The screenshot displays the Cockpit Software Updates interface. At the top, there are two main sections: 'Status' and 'Updates'. The 'Status' section shows 'Updates available' with an information icon. The 'Updates' section shows 'Last Checked: Oct 17, 2022, 11:40 AM' and a 'Check for Updates' button. Below these is a 'Snapshots & Updates' section. It features a table of available updates. The first update is 'SUSE-SLE-Micro-5.3-2022-3564', which is marked as 'recommended' and 'critical'. Below the update name, there is a list of snapshots. Each snapshot has a 'Rollback and Reboot' button. There is also an 'Update and Reboot' button at the top right of the snapshots list.

Snapshot	Time	Status	Action
#10 Snapshot Update of #9	less than a minute ago	Active, Default	Rollback and Reboot
#9 Snapshot Update of #8	2 minutes ago		Rollback and Reboot
#8 Snapshot Update of #7	about 1 month ago		Rollback and Reboot
#7 Snapshot Update of #6	about 1 month ago		Rollback and Reboot
#6 Snapshot Update of #5	about 1 month ago		Rollback and Reboot
#5 Snapshot Update of #3	about 1 month ago		Rollback and Reboot
#4 Snapshot Update of #3	about 1 month ago		Rollback and Reboot
#3 Snapshot Update of #1	about 1 month ago		Rollback and Reboot
#2 Snapshot Update of #1	about 1 month ago		Rollback and Reboot
#1 first root filesystem	about 2 months ago		Rollback and Reboot

FIGURE 28: SOFTWARE UPDATES

12.1 Managing updates

Click *Check for Updates* to get a list of new package updates and patches available for your system. You can view each package update or patch information by clicking the blue icon next to the package update/patch. It is recommended to install the patches marked as important as soon as possible.

To apply the updates and patches, click *Update and Reboot*. The equivalent of the **transactional-update up** command is called, and a new snapshot of your system is created. After the update completes successfully, your system will be rebooted and the new snapshot set as default.

You can postpone the reboot of your system after the update process by selecting *Update without Reboot* from the three dots menu. Bear in mind that you need to 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.

You can also update SLE Micro using CLI. For details, refer to Book “*Administration Guide*”, Chapter 3 “*Administration using transactional updates*”, Section 3.1 “**transactional-update** usage”.

12.2 Performing rollbacks

In Snapshots & Updates you can perform a rollback to one of older snapshots by clicking *Rollback and Reboot*, or *Rollback without Reboot* in threedots menu.

After rebooting the system, the snapshot you rolled back to will be set as active. Do not make any changes (install updates, packages, etc.) before rebooting your system, as the snapshot you rolled back to is not active. Any changes performed before you reboot your system will start from the currently active snapshot.

For a system rollback procedure performed using CLI, refer to Book “*Administration Guide*”, Chapter 3 “*Administration using transactional updates*”, Section 3.3 “*System rollback*”.

A GNU licenses

This appendix contains the GNU Free Documentation License version 1.2.

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. Secondly, 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.

- 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.

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".

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.

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.

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.

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.

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/7>.

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.