



Experimental version for  
testing purpose only!

My private, unofficial Version of:

SUSE Linux Enterprise Server 15 SP7

# Quick Start Guides

## Quick Start Guides

SUSE Linux Enterprise Server 15 SP7

---

File generated at 2025-11-17 15:09

This is my own, **experimental version** of a Document from SUSE company. The only purpose of this document is the test of an alternative publishing mechanism. **Errors in the publishing mechanism may lead to wrong content. You can find the original version of this document at [documentation.suse.com](https://documentation.suse.com).**

The books and articles exist as XML sources, conformant to the DocBook standard. SUSE publishes them with the DocBook XSLT 1.0 Stylesheets, which generate XSL-FO, and Apache FOP, which in turn generates PDF.

This version is based on the same DocBook sources, but published with the new [xsITNG Stylesheets](#), which produce XHTML+CSS, and an rendering engine like *Antenna House* or *Weasyprint* to generate PDF. The only purpose of this version is a "*real life test*" of the new publishing mechanism, together with an "*DocBook TNG Framework*" that i wrote. It helps me to use and customize the xsITNG Stylesheets.  
— Frank Steinke, Bremen, Germany



## Contents

- [Installation Quick Start](#)
- [Modules and Extensions Quick Start](#)
- [Virtualization Best Practices](#)
- [Using KubeVirt on SUSE Linux Enterprise](#)
- [AMD Secure Encrypted Virtualization \(AMD-SEV\) Guide](#)
- [NVIDIA Virtual GPU for KVM Guests](#)
- [Raspberry Pi Quick Start](#)

1	Welcome to SUSE Linux Enterprise Server	1
	The Unified Installer	1
	• Installing offline or without registration	1
2	The installation procedure	2
	Language, keyboard and product selection	3
	• License agreement	4
	• Network settings	5
	• IBM Z: disk activation	5
	• Registration	7
	• Extension and module selection	9
	• Add-on product	13
	• System role	14
	• Suggested partitioning	15
	• Clock and time zone	17
	• Local user	18
	• Authentication for the system administrator “root”	19
	• Installation settings	20
	• Start the installation	23
	• The installation process	24
3	Legal notice	24
4	GNU Free Documentation License	25
1	Introduction to modules and extensions	33
	Support	34
2	Modules and extensions for the SUSE Linux Enterprise product family	34
	Basesystem	35
	Certifications	35
	Confidential Computing Technology Preview	35
	Containers	36
	Desktop Applications	37
	Development Tools	37
	High Availability	37
	High Performance Computing (HPC)	38
	Legacy	38
	Live Patching	38
	NVIDIA Compute	39
	Public Cloud	39
	Python 3	40
	Real Time	40
	SAP Applications	40
	SAP Business One Server	41
	Server Applications	41
	SUSE Package Hub	41
	Systems Management	42
	Transactional Server	42
	Web and Scripting	42
	Workstation Extension	42
3	Installing and managing modules and extensions	43
	Installing modules and extensions during system installation	43
	• Installing modules and extensions with YaST	43
	• Installing extensions and third-party add-on products from media	45
	• Deleting modules and extensions with YaST	48
	• Adding or deleting modules and extensions with SUSEConnect	49
4	Module membership of packages	51
5	Legal notice	54
6	GNU Free Documentation License	54
1	Virtualization scenarios	62
2	Before you apply modifications	62
	Back up first	62
	• Test your workloads	62
3	Recommendations	63
	Prefer the <code>libvirt</code> framework	63
	• <code>qemu-system-i386</code> compared to <code>qemu-system-x86_64</code>	63

- 4 VM Host Server configuration and resource allocation 63
  - Memory 63 • Swap 69 • I/O 70 • Storage and file system 72 • CPUs 74 • NUMA tuning 77
- 5 VM Guest images 82
  - VM Guest image formats 82 • Overlay disk images 84 • Opening a VM Guest image 87 • File system sharing 91
- 6 VM Guest configuration 92
  - Virtio driver 92 • Cirrus video driver 94 • Better entropy 95 • Disable unused tools and devices 95 • Updating the guest machine type 96
- 7 VM Guest-specific configurations and settings 97
  - ACPI testing 97 • Keyboard layout 97 • Spice default listen URL 98 • XML to QEMU command line 98 • Change kernel parameters at boot time 99 • Add a device to an XML configuration 99 • Adding and removing CPUs 100 • SCSI persistent reservation on a multipathed device 100
- 8 More information 103
- 9 Legal notice 104
- 10 GNU Free Documentation License 104
- 1 KubeVirt components 112
- 2 Installing KubeVirt on Kubernetes 113
- 3 Updating the KubeVirt deployment 114
- 4 Deleting KubeVirt from a cluster 114
- 5 Containerized Data Importer 114
  - Installing CDI 114 • Updating and deleting CDI: 115
- 6 Running virtual machines 115
- 7 Live migration 117
  - Prerequisites 117 • Initiating live migration 117 • Cancelling live migration 118
- 8 Volume hotplugging 118
- 9 Running Windows VMs with VMDP ISO 118
- 10 Supported features 119
  - VMI virtual hardware 119 • VMI disks and volumes 119 • VMI interfaces and networks 120
- 11 Debugging 120
  - 1 Introducing SEV 122
  - 2 VM host requirements 123

3	VM requirements	123
4	VM configuration	123
5	VM installation	125
6	VM attestation	126
7	SEV with KubeVirt	128
8	Current limitations	129
9	More information	129
10	Legal notice	129
11	GNU Free Documentation License	130
1	Introduction	138
2	Configuring vGPU manager in VM Host Server	138
	Prepare VM Host Server environment	138
	• Remove conflicting and pre-installed NVIDIA drivers	
	139	
	• Install the NVIDIA KVM driver	140
	• Verify the driver installation	141
3	Creating a vGPU device	142
	Create a legacy vGPU device without support for SR-IOV	142
	• Create a vGPU device with support	
	for SR-IOV	144
	• Creating a MIG-backed vGPU	146
4	Assign the vGPU device to a VM Guest	149
	Assign by <code>libvirt</code>	149
	• Assign by QEMU	149
5	Configuring vGPU in VM Guest	149
	Prepare the VM Guest	149
	• Install the vGPU driver in the VM Guest	150
6	Licensing vGPU in the VM Guest	151
7	Configuring a graphics mode	152
	Create or update the <code>/etc/X11/xorg.conf</code> file	152
	• Verify the graphics mode	153
	• Remote display	153
8	Configuring compute mode	154
9	Additional tasks	154
	Disabling Frame Rate Limiter	154
	• Enabling/Disabling Error Correcting Code (ECC)	154
	• Black screen in Virt-manager	155
	• Black screen in VNC client when using a non-QEMU VNC server	
	155	
	• Kernel panic occurs because the Nouveau and NVIDIA drivers compete on GPU resources	
	155	
	• Filing an NVIDIA vGPU bug	156
	• Configuring a License Server	157
10	For more information	157

11	NVIDIA virtual GPU background 157	
	NVIDIA GPU architectures 157 • vGPU types 158 • Valid vGPU configurations on a single GPU 159	
12	GNU Free Documentation License 159	
1	Platform overview 168	
	Technical details of the Raspberry Pi 3 model B 168 • SUSE Linux Enterprise Server for Arm 15	
	SP7 170	
2	Installation 173	
	Preparing the card on Linux 173 • Preparing the card on Microsoft Windows operating systems	
	174 • Preparing the card on macOS 175	
3	First boot 176	
	Initial system setup with JeOS First Boot Assistant 176 • Registration process 182	
4	Initial system configuration 185	
	Changing the language 186 • Changing the host name 186 • Setting up networking 186	
5	General system usage 187	
	Desktop 187 • Bluetooth 187	
6	Product documentation 188	
	Product documentation 188 • SUSE forums 188	
7	Trademarks 189	
8	Legal notice 189	
9	GNU Free Documentation License 189	
A	<b>GNU licenses 197</b>	
A . 1	GNU Free Documentation License 197	





Koordinierungsstelle  
für IT-Standards

## Installation Quick Start

This Quick Start guides you step by step through the installation of SUSE® Linux Enterprise Server15 SP7.

### 1. Welcome to SUSE Linux Enterprise Server

Use the following procedures to install the product on all supported hardware platforms. They assume you have successfully booted/IPLed into the installation system. For more detailed installation instructions and deployment strategies, refer to *Deployment Guide*. For platform-specific preparations required to boot or IPL into the installation system, refer to the following parts of the aforementioned guide:

- *Chapter 3, Installation on AMD64 and Intel 64*
- *Chapter 4, Installation on AArch64*
- *Chapter 5, Installation on IBM POWER*
- *Chapter 6, Installation on IBM Z and LinuxONE*

#### 1.1. The Unified Installer

Starting with SUSE Linux Enterprise Server 15, the installation medium only consists of the Unified Installer, a minimal system for installing, updating and registering all SUSE Linux Enterprise base products. During the installation you can add functionality by selecting modules and extensions to be installed on top of the Unified Installer.

#### 1.2. Installing offline or without registration

The default installation medium SLE-15-SP7-Online-ARCH-GM-media1.iso is optimized for size and does not contain any modules and extensions. Therefore, the installation requires network access to register your product and retrieve repository data for the modules and extensions.

For installation without registering the system, use the SLE-15-SP7-Full-ARCH-GM-media1.iso image from <https://www.suse.com/download/sles/> and refer to the section called “*Installing without registration*” for instructions.

## Copying the installation media image to a removable flash disk



Use the following command to copy the contents of the installation image to a removable flash disk.

```
>sudo dd if=IMAGE of=FLASH_DISK bs=4M && sync
```

*IMAGE* needs to be replaced with the path to the SLE-15-SP7-Online-ARCH-GM-media1.iso or SLE-15-SP7-Full-ARCH-GM-media1.iso image file. *FLASH\_DISK* needs to be replaced with the flash device. To identify the device, insert it and run:

```
#grep -Ff <(hwinfo --disk --short) <(hwinfo --usb --short)
disk:
  /dev/sdc           General USB Flash Disk
```

Make sure the size of the device is sufficient for the desired image. You can check the size of the device with:

```
#fdisk -l /dev/sdc | grep -e "^\ /dev"
  /dev/sdc1  *  2048 31490047 31488000  15G 83 Linux
```

In this example, the device has a capacity of 15 GB. The command to use for the SLE-15-SP7-Full-ARCH-GM-media1.iso would be:


```
dd if=SLE-15-SP7-Full-ARCH-GM-media1.iso of=/dev/sdc bs=4M && sync
```

The device must not be mounted when running the **dd** command. Note that all data on the partition will be erased!

## 2. The installation procedure

To install SUSE Linux Enterprise Server, boot or IPL into the installer from the Unified Installer medium and start the installation.

## 2.1. Language, keyboard and product selection



The *Language* and *Keyboard Layout* settings are initialized with the language you chose on the boot screen. If you did not change the default, it will be English (US). Change the settings here, if necessary. Use the *Keyboard Test* text box to test the layout.

With the Unified Installer, you can install all SUSE Linux Enterprise base products:

- SUSE Linux Enterprise Server 15 SP7 (covered here)
- SUSE Linux Enterprise Desktop 15 SP7 (for installation instructions, refer to <https://documentation.suse.com/sled/>)
- SUSE Linux Enterprise Real Time 15 SP7 (for installation instructions, refer to <https://documentation.suse.com/sle-rt/>)
- SUSE Linux Enterprise Server for SAP applications 15 SP7 (for installation instructions, refer to <https://documentation.suse.com/sles-sap>)

Select a product for installation. You need to have a registration code for the respective product. In this document, it is assumed you have chosen SUSE Linux Enterprise Server. Proceed with *Next*.

## Light and high-contrast themes



If you have difficulties reading the labels in the installer, you can change the widget colors and theme.

Click the button or press **Shift-F3** to open a theme selection dialog. Select a theme from the list and *Close* the dialog.


**Shift-F4** switches to the color scheme for vision-impaired users. Press the buttons again to switch back to the default scheme.

## 2.2. License agreement



Read the License Agreement. It is presented in the language you have chosen on the boot screen. Translations are available via the *License Language* drop-down list. You need to accept the agreement by checking *I Agree to the License Terms* to install SUSE Linux Enterprise Server. Proceed with *Next*.

## 2.3. Network settings



A system analysis is performed, where the installer probes for storage devices and tries to find other installed systems. If the network was automatically configured via DHCP during the start of the installation, you will be taken to the registration step.

If the network is not yet configured, the *Network Settings* dialog opens. Choose a network interface from the list and configure it with *Edit*. Alternatively, *Add* an interface manually. See *the section called “Network settings”* and *the section called “Configuring a network connection with YaST”* for more information. If you prefer to do an installation without network access, skip this step without making any changes and proceed with *Next*.

## 2.4. IBM Z: disk activation

Skip this step if you are not installing on IBM Z hardware.




IBM z13

## Disk Activation

[Network Configuration...](#)[Configure DASD Disks](#)[Configure zFCP Disks](#)[Configure iSCSI Disks](#)[Help](#)[Abort](#)[Back](#)[Next](#)

Configure the attached hard disks. Select DASD, Fibre Channel Attached SCSI Disks (zFCP), or iSCSI for the installation of SUSE Linux Enterprise Server. The DASD and zFCP configuration buttons are only available if the corresponding devices are attached. Proceed with *Next* to set up the selected disk type. For more information, refer to *the section called “IBM Z: disk activation”*.

## 2.5. Registration



To get technical support and product updates, you need to register and activate SUSE Linux Enterprise Server with the SUSE Customer Center or a local registration server. Registering your product at this stage also grants you immediate access to the update repository. This enables you to install the system with the latest updates and patches available.

When registering, repositories and dependencies for modules and extensions are loaded from the registration server.

### ***Register system via scc.suse.com***

To register with the SUSE Customer Center, enter the *E-mail Address* associated with your SUSE Customer Center account and the *Registration Code* for SUSE Linux Enterprise Server. Proceed with *Next*.

### ***Register system via local RMT server***

If your organization provides a local registration server, you may alternatively register there. Activate *Register System via local RMT Server* and either choose a URL from the drop-down list or type in an address. Proceed with *Next*.

## **Skip registration**

If you are offline or want to skip registration, activate *Skip Registration*. Accept the warning with *OK* and proceed with *Next*.

### **Skipping the registration**



Your system and extensions need to be registered in order to retrieve updates and to be eligible for support. Skipping the registration is only possible when installing from the `SLE-15-SP7-Full-ARCH-GM-media1.iso` image.

If you do not register during the installation, you can do so at any time later from the running system. To do so, run `YaST > Product Registration` or the command-line tool **SUSEConnect**.

### **Installing product patches at installation time**




After SUSE Linux Enterprise Server has been successfully registered, you are asked whether to install the latest available online updates during the installation. If choosing **Yes**, the system will be installed with the most current packages without having to apply the updates after installation. Activating this option is recommended.

### **Firewall settings for receiving updates**



If your system is behind a firewall that blocks outgoing traffic, make sure to allow connections to `https://scc.suse.com/` and `https://updates.suse.com` on ports 80 and 443 in order to receive updates. For more information, such as IP addresses and proxy server configuration, refer to <https://www.suse.com/support/kb/doc/?id=000021034>.

## 2.6. Extension and module selection



After the system is successfully registered, the installer lists modules and extensions that are available for SUSE Linux Enterprise Server. Modules are components that allow you to customize the product according to your needs. They are included in your SUSE Linux Enterprise Server subscription. Extensions add functionality to your product. They must be purchased separately.

The availability of certain modules or extensions depends on the product selected in the first step of the installation. For a description of the modules and their lifecycles, select a module to see the accompanying text. More detailed information is available in the [Modules and Extensions Quick Start](#).

The selection of modules indirectly affects the scope of the installation, because it defines which software sources (repositories) are available for installation and in the running system.

The following modules and extensions are available for SUSE Linux Enterprise Server:

### Basesystem Module

This module adds a basic system on top of the Unified Installer. It is required by all other modules and extensions. The scope of an installation that only contains the base system is comparable to the installation pattern *minimal system* of previous SUSE Linux Enterprise

Server versions. This module is selected for installation by default and should not be deselected.

*Dependencies:* None

### **Certifications Module**

Contains the FIPS certification packages.

*Dependencies:* Server Applications

### **Confidential Computing Technical Preview**

Contains packages related to confidential computing.

*Dependencies:* Basesystem

### **Containers Module**

Contains support and tools for containers.

*Dependencies:* Basesystem

### **Desktop Applications Module**

Adds a graphical user interface and essential desktop applications to the system.

*Dependencies:* Basesystem

### **Development Tools Module**

Contains compilers (including gcc) and libraries required for compiling and debugging applications. Replaces the former Software Development Kit (SDK).

*Dependencies:* Basesystem, Desktop Applications

### **High Performance Computing (HPC) Module**

Provides specific tools commonly used for high performance, numerically intensive workloads.

*Dependencies:* Basesystem

### **Legacy Module**

Helps you with migrating applications from earlier versions of SUSE Linux Enterprise Server and other systems to SLES 15 SP7, by providing packages which are discontinued on SUSE Linux Enterprise. Packages in this module are selected based on the requirement for migration and the level of complexity of configuration.

This module is recommended when migrating from a previous product version.

*Dependencies:* Basesystem, Server Applications

### **NVIDIA Compute Module**

Contains the NVIDIA CUDA (Compute Unified Device Architecture) drivers.

The software in this module is provided by NVIDIA under the [CUDA End User License Agreement](#) and is not supported by SUSE.

*Dependencies:* Basesystem

### **Public Cloud Module**

Contains all tools required to create images for deploying SUSE Linux Enterprise Server in cloud environments such as Amazon Web Services (AWS), Microsoft Azure, Google Compute Platform, or OpenStack.

*Dependencies:* Basesystem, Server Applications

### **Python 3 Module**

This module contains the most recent version of the selected Python 3 packages.

*Dependencies:* Basesystem

### **SAP Business One Server**

This module contains packages and system configuration specific to SAP Business One Server. It is maintained and supported by the SUSE Linux Enterprise Server product subscription.

*Dependencies:* Basesystem, Server Applications, Desktop Applications, Development Tools

### **Server Applications Module**

Adds server functionality by providing network services such as DHCP server, name server, or Web server. This module is selected for installation by default; deselecting it is not recommended.

*Dependencies:* Basesystem

### **SUSE Linux Enterprise High Availability**

Adds clustering support for mission critical setups to SUSE Linux Enterprise Server. This extension requires a separate license key.

*Dependencies:* Basesystem, Server Applications

### **SUSE Linux Enterprise Live Patching**

Adds support for performing critical patching without having to shut down the system. This extension requires a separate license key.

*Dependencies:* Basesystem, Server Applications

### **SUSE Linux Enterprise Workstation Extension**

Extends the functionality of SUSE Linux Enterprise Server with packages from SUSE Linux Enterprise Desktop, like additional desktop applications (office suite, e-mail client, graphical editor, etc.) and libraries. It allows to combine both products to create a fully featured workstation. This extension requires a separate license key.

*Dependencies:* Basesystem, Desktop Applications

### **SUSE Package Hub**

Provides access to packages for SUSE Linux Enterprise Server maintained by the open-SUSE community. These packages are delivered without L3 support and do not interfere with the supportability of SUSE Linux Enterprise Server. For more information, refer to <https://packagehub.suse.com/>.

*Dependencies:* Basesystem

### **Transactional Server Module**

Adds support for transactional updates. Updates are either applied to the system all together in a single transaction, or not. This happens without influencing the running system. If an update fails, or if the successful update is deemed to be incompatible or otherwise incorrect, it can be discarded to immediately return the system to its previous functioning state.

*Dependencies:* Basesystem

### **Web and Scripting Module**

Contains packages intended for a running Web server.

*Dependencies:* Basesystem, Server Applications

Some modules depend on the installation of other modules. Therefore, when selecting a module, other modules may be selected automatically to fulfill dependencies.

Depending on the product, the registration server can mark modules and extensions as recommended. Recommended modules and extensions are preselected for registration and installation. To avoid installing these recommendations, deselect them manually.


Select the modules and extensions you want to install and proceed with *Next*. In case you have chosen one or more extensions, you will be prompted to provide the respective registration codes. Depending on your choice, it may also be necessary to accept additional license agreements.

### Default modules for offline installation



When performing an offline installation from the SLE-15-SP7-Full-ARCH-GM-medi-a1.iso, only the *Basesystem Module* is selected by default. To install the complete default package set of SUSE Linux Enterprise Server, additionally select the *Server Applications Module* and the *Python 3 Module*.

## 2.7. Add-on product



The *Add-On Product* dialog allows you to add additional software sources (called “repositories”) to SUSE Linux Enterprise Server that are not provided by the SUSE Customer Center. Add-on products may include third-party products and drivers as well as additional software for your system.

## Adding drivers during the installation




You can also add driver update repositories via the *Add-On Product* dialog. Driver updates for SUSE Linux Enterprise are provided at <https://drivers.suse.com/>. These drivers have been created via the SUSE SolidDriver Program.

If you want to skip this step, proceed with *Next*. Otherwise, activate *I would like to install an additional Add On Product*. Specify a media type, a local path, or a network resource hosting the repository and follow the on-screen instructions.

Check *Download Repository Description Files* to download the files describing the repository now. If deactivated, they will be downloaded after the installation has started. Proceed with *Next* and insert a medium if required. Depending on the content of the product, it may be necessary to accept additional license agreements. Proceed with *Next*. If you have chosen an add-on product requiring a registration key, you will be asked to enter it before proceeding to the next step.

## 2.8. System role



The availability of system roles depends on your selection of modules and extensions. System roles define, for example, the set of software patterns that is going to be preselected for the installation. Refer to the description on screen to make your choice. Select a role and proceed with *Next*.


If from the enabled modules only one role or no role is suitable for the respective base product, the *System Role* dialog is omitted.

### Release notes



From this point on, the Release Notes can be viewed from any screen during the installation process by selecting *Release Notes*.

## 2.9. Suggested partitioning



Review the partition setup proposed by the system. If necessary, change it. You have the following options:

#### **Guided setup**

Starts a wizard which lets you refine the partitioning proposal. The options available here depend on your system setup. If it contains more than a single hard disk, you may choose which disk(s) to use and where to place the root partition. If the disk(s) already contain partitions, decide whether to remove or resize them.

In subsequent steps, you may also add LVM support and disk encryption. You can change the file system for the root partition and decide whether or not to have a separate home partition.

### ***Expert partitioner***

Opens the *Expert Partitioner* described in the section called “*Using the Expert Partitioner*”. This gives you full control over the partitioning setup and lets you create a custom setup. This option is intended for experts.

### **Disk space units**



Note that for partitioning purposes, disk space is measured in binary units, rather than in decimal units. For example, if you enter sizes of 1GB, 1GiB or 1G, they all signify 1 GiB (Gibibyte), as opposed to 1 GB (Gigabyte).

#### **Binary**

$1 \text{ GiB} = 1\,073\,741\,824 \text{ bytes.}$

#### **Decimal**


$1 \text{ GB} = 1\,000\,000\,000 \text{ bytes.}$

#### **Difference**

$1 \text{ GiB} \approx 1.07 \text{ GB.}$


To accept the proposed setup without any changes, choose *Next* to proceed.

## 2.10. Clock and time zone



Select the clock and time zone to use in your system. To manually adjust the time or to configure an NTP server for time synchronization, choose *Other Settings*. See the section called “Clock and time zone” for detailed information. Proceed with *Next*.

## 2.11. Local user



To create a local user, type the first and last name in the *User's Full Name* field, the login name in the *Username* field, and the password in the *Password* field.

The password should be at least eight characters long and should contain both uppercase and lowercase letters and numbers. The maximum length for passwords is 72 characters, and passwords are case-sensitive.


For security reasons, it is also strongly recommended *not* to enable the *Automatic Login*. You should also *notUse this Password for the System Administrator* but rather provide a separate root password in the next installation step.

If you install on a system where a previous Linux installation was found, you may *Import User Data from a Previous Installation*. Click *Choose User* for a list of available user accounts. Select one or more users.

In an environment where users are centrally managed (for example, by NIS or LDAP), you can skip the creation of local users. Select *Skip User Creation* in this case.

Proceed with *Next*.

## 2.12. Authentication for the system administrator “root”



Type a password for the system administrator (called the `root` user) or provide a public SSH key. If you want, you can use both.

Because the `root` user is equipped with extensive permissions, the password should be chosen carefully. You should never forget the `root` password! After you entered it here, the password cannot be retrieved.

### Passwords and keyboard layout




It is recommended to only use US ASCII characters. In case of a system error or when you need to start your system in rescue mode, the keyboard may not be localized.

If you want to access the system remotely via SSH using a public key, import a key from a removable media or an existing partition. See *the section called “Authentication for the system administrator `root`” for more information*.

Proceed with `Next`.

## 2.13. Installation settings



Use the *Installation Settings* screen to review and—if necessary—change several proposed installation settings. The current configuration is listed for each setting. To change it, click the headline. Some settings, such as firewall or SSH, can be changed directly by clicking the respective links.

### Remote access



Changes you can make here can also be made later at any time from the installed system. However, if you need remote access right after the installation, you may need to open the SSH port in the *Security* settings.

### Software

The scope of the installation is defined by the modules and extensions you have chosen for this installation. However, depending on your selection, not all packages available in a module are selected for installation.

Clicking *Software* opens the *Software Selection and System Tasks* screen, where you can change the software selection by selecting or deselecting patterns. Each pattern contains several software packages needed for specific functions (for example, *KVM Host Server*).

For a more detailed selection based on software packages to install, select *Details* to switch to the YaST *Software Manager*. See *Chapter 8, Installing or removing software* for more information.

## Booting

This section shows the boot loader configuration. Changing the defaults is only recommended if really needed. Refer to *Chapter 18, The boot loader GRUB 2* for details.

## Security

The *CPU Mitigations* refer to kernel boot command-line parameters for software mitigations that have been deployed to prevent CPU side-channel attacks. Click the selected entry to choose a different option. For details, see *CPU Mitigations* .

By default, the *Firewall* is enabled on all configured network interfaces. To completely disable `firewalld`, click *disable* (not recommended). Refer to *Chapter 23, Masquerading and firewalls* for configuration details.

### Firewall settings for receiving updates



If your system is behind a firewall that blocks outgoing traffic, make sure to allow connections to `https://scc.suse.com/` and `https://updates.suse.com` on ports 80 and 443 in order to receive updates. For more information, such as IP addresses and proxy server configuration, refer to <https://www.suse.com/support/kb/doc/?id=000021034>.

The *SSH service* is enabled by default, but its port (22) is closed in the firewall. Click *open* to open the port or *disable* to disable the service. Note that if SSH is disabled, remote logins will not be possible. Refer to *Chapter 22, Securing network operations with OpenSSH* for more information.

The default *Major Linux Security Module* is *AppAmpor*. To disable it, select *None* as module in the *Security* settings.

## Security Policies

Click to *enable* the Defense Information Systems Agency STIG security policy. If any installation settings are incompatible with the policy, you will be prompted to modify them accordingly. Some settings can be adjusted automatically, others require user input.

Enabling a security profile enables a full SCAP remediation on first boot. You can also perform a *scan only* or *do nothing* and manually remediate the system later with OpenSCAP. For more information, refer to the section called “*Security Profiles* ”.

## **Network configuration**

Displays the current network configuration. By default, **wicked** is used for server installations and NetworkManager for desktop workloads. Click *Network Configuration* to change the settings. For details, see *the section called “Configuring a network connection with YaST”*.

## **Support for NetworkManager**



SUSE only supports NetworkManager for desktop workloads with SLED or the Workstation extension. All server certifications are done with **wicked** as the network configuration tool, and using NetworkManager may invalidate them. NetworkManager is not supported by SUSE for server workloads.

## **Kdump**

Kdump saves the memory image (“core dump”) to the file system in case the kernel crashes. This enables you to find the cause of the crash by debugging the dump file. Kdump is pre-configured and enabled by default. See *the section called “Basic Kdump configuration”* for more information.


## **Default systemd target**

If you have installed the desktop applications module, the system boots into the *graphical* target, with network, multiuser and display manager support. Switch to *multi-user* if you do not need to log in via display manager.

## **System**


View detailed hardware information by clicking *System*. In the resulting screen you can also change *Kernel Settings*—see *the section called “System”* for more information.

## 2.14. Start the installation



After you have finalized the system configuration on the *Installation Settings* screen, click *Install*. Depending on your software selection, you may need to agree to license agreements before the installation confirmation screen pops up. Up to this point, no changes have been made to your system. After you click *Install* a second time, the installation process starts.

## 2.15. The installation process



During the installation, the progress is shown. After the installation routine has finished, the computer is rebooted into the installed system.

### 3. Legal notice

Copyright© 2006– 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.

## 4. 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:

1. 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.
2. 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.
3. State on the Title page the name of the publisher of the Modified Version, as the publisher.
4. Preserve all the copyright notices of the Document.
5. Add an appropriate copyright notice for your modifications adjacent to the other copyright notices.
6. 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.
7. Preserve in that license notice the full lists of Invariant Sections and required Cover Texts given in the Document's license notice.
8. Include an unaltered copy of this License.
9. 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.
10. 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.

11. 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.
12. 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.
13. Delete any section Entitled "Endorsements". Such a section may not be included in the Modified Version.
14. Do not retitle any existing section to be Entitled "Endorsements" or to conflict in title with any Invariant Section.
15. 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.



# Modules and Extensions Quick Start

Koordinierungsstelle  
für IT-Standards

This Quick Start gives you an overview of the modules and extensions available for the SUSE Linux Enterprise family. Learn about their purpose and features, and how to manage them.

## 1. Introduction to modules and extensions

Starting with SUSE Linux Enterprise Server 15, the installation medium consists of only the Unified Installer: a minimal system for deploying, updating and registering SUSE Linux Enterprise Server. During deployment, you can add functionality by selecting modules and extensions to be installed on top of the Unified Installer.

Technically, modules and extensions are very similar. Both can be managed with YaST or the command-line tool **SUSEConnect**. Both modules and extensions add a set of online repositories (main, update, source, debug) containing all necessary packages. Additionally, both may add partitioning proposals and system roles that can be chosen during deployment.

However, there are significant differences between modules and extensions:

### Modules

Modules allow you to shape the product according to your needs. Each module has a clearly defined scope ranging from tools for Web Development and Scripting, through a Public Cloud, all the way to the SUSE Package Hub, offering community-managed packages.

Modules enable you to provide only the set of packages required for the machine's purpose, making the system lean, fast, and more secure. This modular approach also makes it easy to provide tailor-made images for container and cloud environments.

Modules can be added or removed at any time during the lifecycle of the system. This allows you to easily adjust the system to changing requirements.

Modules are part of your SUSE Linux Enterprise Server subscription and therefore free of charge. They differ from the base of SUSE Linux Enterprise Server only by their lifecycle.

### Extensions

Extensions provide advanced capabilities for tasks such as live patching and high availability clustering to SUSE Linux Enterprise Server. They are offered as subscriptions and require a

registration key that is liable for costs. Usually, extensions have their own release notes that are available from <https://www.suse.com/releasenotes>.

## 1.1. Support

L3 support is generally provided for packages except for SUSE Package Hub and third-party modules. To check the support level for a package, run **zypper info PACKAGE**.

Long-Term Service Pack Support(LTSS) is available for most modules and extensions. For more information, refer to <https://www.suse.com/products/long-term-service-pack-support/>.

Some modules and extensions also include Extended Service Pack Overlay Support (ESPOS), which gives you significantly more time to upgrade and/or skip a service pack completely.

For more information on support and lifecycles, refer to the [SUSE Technical Support Policy](#) and the [Product Lifecycle Support Policies](#). For a complete list of lifecycle dates by product, refer to the [Product Support Lifecycle](#) page.

## 2. Modules and extensions for the SUSE Linux Enterprise product family

The following modules and extensions are available within the SUSE Linux Enterprise product family:

- *Basesystem*
- *Certifications*
- *Confidential Computing Technology Preview*
- *Containers*
- *Desktop Applications*
- *Development Tools*
- *High Availability*
- *High Performance Computing (HPC)*
- *Legacy*
- *Live Patching*
- *NVIDIA Compute*
- *Public Cloud*
- *Python 3*
- *Real Time*
- *SAP Applications*
- *Server Applications*
- *SUSE Package Hub*

- *Systems Management*
- *Transactional Server*
- *Web and Scripting*
- *Workstation Extension*

Please note that the availability depends on the product. Not all modules and extensions are available for all products. Some modules are included in one product but also available as extension for another.

## 2.1. Basesystem

This module adds a basic system on top of the Installer. It is required by all other modules and extensions. The scope of an installation that only contains the base system is comparable to the minimal system installation pattern of previous SUSE Linux Enterprise Server versions. This module is selected for installation by default and should not be deselected.

- **Dependencies:** None
- **Availability:** Default on all products
- **Lifecycle:** 10 years
- **Extended Support:** 3 years LTSS
- **Support level:** L3

## 2.2. Certifications

This module contains the FIPS certification packages.

- **Depends on:** Basesystem
- **Available for:** SLES, SLED, SLES for SAP, SUSE Linux Enterprise Real Time
- **Lifecycle:** 10 years
- **Extended support:** 3 years LTSS
- **Support level:** L3

## 2.3. Confidential Computing Technology Preview

This module contains a Secure Virtual Machine Service Module (SVSM), which aims to provide secure services and device emulations to guest operating systems on confidential virtual machines (CVMs).

### Technology preview



The module is a technology preview and is not supported. Its lifecycle and support options will be defined at a later stage.

## AMD SEV required



The SVSM requires AMD Secure Encrypted Virtualization (AMD SEV) with Secure Nested Paging. It will only work on AMD systems.

- **Depends on:** Server Applications
- **Available for:** SLES on AMD64/Intel 64
- **Lifecycle:** Until the release of SLE 15 SP7
- **Extended support:** None
- **Support level:** This module is a technology preview and is not supported.

## 2.4. Containers

This module contains packages relating to containers, including the container engine and core container-related tools such as on-premise registry.

### docker-stable



SUSE ships a package named `docker-stable`, based on selected versions of the Moby project, which is the upstream for Docker. When this version is no longer supported by the Moby project, SUSE will provide reactive support for an additional 3 years, which may require a LTSS subscription depending on the lifecycle of the underlying product.

### Migration from docker to docker-stable



In alignment with Moby project practices, SUSE neither tests nor supports downgrading `docker` to older versions. This policy is particularly relevant when considering a migration from the latest `docker` package to the `docker-stable` package, which in most cases would be an older version. Such a downgrade can result in unexpected side effects, compatibility issues and even disruptions to container production environments, and is not recommended for production systems.

### Docker Swarm is not supported



SUSE does not provide any support related to use of Docker Swarm orchestration. Related packages are shipped for customer convenience only.

- **Dependencies:** Basesystem

- **Availability:** SLES, SLES for SAP
- **Lifecycle:** By package; follows the upstream community of the respective package.
- **Extended support:** 3 years LTSS
- **Support level:** L3

## 2.5. Desktop Applications

This module adds a graphical user interface and essential desktop applications to the system.

- **Dependencies:** Basesystem
- **Availability:** SLES, SLED (default), SLE WE (default), SLES for SAP, SLE RT (default)
- **Lifecycle:** 10 years
- **Extended support:** 3 years LTSS
- **Level:** L3

## 2.6. Development Tools

This module contains compilers (including gcc) and libraries required for compiling and debugging applications. Replaces the former Software Development Kit (SDK).

**Dependencies:** Basesystem, Desktop Applications

- **Availability:** SLES, SLES for SAP, SLED, SLE RT (default)
- **Lifecycle:** 10 years
- **Extended support:** 3 years LTSS
- **Level:** L3

## 2.7. High Availability

This extension provides high-availability clustering technologies that can be deployed in physical and/or virtual environments.

High Availability is included in the subscription for SLES for SAP as a module. It is also available as an extension for SLES. For more information about SUSE Linux Enterprise High Availability, refer to <https://www.suse.com/products/highavailability>.

- **Dependencies:** Basesystem, Server Applications
- **Availability:** SLES, SLES for SAP (included), SUSE Linux Enterprise Real Time
- **Lifecycle:** 10 years
- **Extended support:** 3 years LTSS
- **Level:** L3

## 2.8. High Performance Computing (HPC)

This module provides specific tools commonly used for high performance, numerically intensive workloads.

SUSE packages these tools in the HPC module to provide greater flexibility to deliver new versions or new tools more rapidly than the standard SUSE Linux Enterprise lifecycle would permit. Packages in this module are generally supported until a newer version of the package is released or a package is dropped from the module.

- **Dependencies:** Basesystem
- **Availability:** SLES on AMD64/Intel 64 and AArch64
- **Lifecycle:** 10 years
- **Extended support:** 1 year ESPOS, 3 years LTSS
- **Level:** L3

## 2.9. Legacy

This module helps you with migrating applications from earlier versions of SUSE Linux Enterprise Server and other systems to SLES 15 SP7, by providing packages which are discontinued on SUSE Linux Enterprise. Packages in this module are selected based on the requirement for migration and the level of complexity of configuration.

This module is recommended when migrating from a previous product version.

Please note that this module has a different lifecycle than some of its packages. Not all packages are supported for the complete lifecycle of the module, but depend on migration requirements and upstream lifecycles.

- **Dependencies:** Basesystem, Server Applications
- **Availability:** SLES, SLES for SAP
- **Lifecycle:** 10 years
- **Extended support:** 3 years LTSS
- **Level:** L3

## 2.10. Live Patching

This extension provides packages to update critical components in SUSE Linux Enterprise without shutting down your system, which reduces planned downtime and increases service availability.

- **Dependencies:** Basesystem
- **Availability:** SLES, SLES for SAP on AMD64/Intel 64, s390x, and ppc64le

- **Lifecycle:** 10 years
- **Extended support:** 3 years LTSS
- **Level:** L3

## 2.11. NVIDIA Compute

This module contains the NVIDIA CUDA (Compute Unified Device Architecture) drivers. The software in this module is provided by NVIDIA under the [CUDA End User License Agreement](#) and is not supported by SUSE.

### Importing third-party GPG keys



As this module is provided by a third party, the repository's GPG key is unknown to the system. To activate the module and import the key with SUSEConnect, use the `--gpg-auto-import-keys` option:

```
suseconnect -p sle-module-NVIDIA-compute/15/x86_64 --gpg-auto-import-keys
```

Alternatively, you can activate the module with `yast registration` or during installation.

### Combining Workstation Extension and NVIDIA Compute module is unsupported



The Workstation Extension provides drivers for NVIDIA graphics cards. These drivers share some packages with CUDA, however, versions may differ. It is therefore neither recommended nor supported to enable both the NVIDIA Compute Module and the Workstation Extension at the same time.

- **Dependencies:** Basesystem
- **Availability:** SLES on AMD64/Intel 64 and AArch64
- **Support:** Software in this module is provided by NVIDIA and is not supported by SUSE

## 2.12. Public Cloud

This module contains all tools required to create images for deploying SUSE Linux Enterprise Server in cloud environments. For example: Amazon Web Services (AWS), Microsoft Azure, Google Compute Platform, or OpenStack.

- **Dependencies:** Basesystem, Server Applications
- **Availability:** SLES, SLES for SAP
- **Lifecycle:** 10 years

- **Extended support:** 3 years LTSS
- **Level:** L3

## 2.13. Python 3

This module provides two additional Python interpreters:

- Python 3.11 and many additional maintained Python 3.11 modules. The Python 3.11 interpreter and modules are supported until the end of December 2029 and might be updated with the next service pack (if compatible and when needed).
- Python 3.13 with all necessary modules for getting Python modules from PyPI. The Python 3.13 interpreter and modules are supported for 1.5 years until the end of December 2026.

Packages from the Python 3 module can be installed alongside existing Python packages, and they can coexist in the same system without impacting Python 3.6 workloads.

- **Dependencies:** Basesystem
- **Availability:** SLES (default), SLES for SAP, SLED (default)
- **Lifecycle:** Python 3.11: 31 Dec 2029, Python 3.13: 31 Dec 2026
- **Extended support:** None
- **Level:** L3

## 2.14. Real Time

This extension aims to reduce the latency and increase the predictability and reliability of time-sensitive mission-critical applications.

Packages in this module are generally supported until a newer version of the package is released or the package is dropped from the module.

- **Dependencies:** Basesystem, Desktop Applications, Development Tools, Server Applications
- **Availability:** SLE RT (included)
- **Lifecycle:** 10 years
- **Extended support:** None
- **Level:** L3

## 2.15. SAP Applications

This module contains packages specific to SLES for SAP. It has a different lifecycle than SUSE Linux Enterprise itself. Subscriptions for SUSE Linux Enterprise Server for SAP applications include 1.5 years of General Support plus 3 years of Extended Service Pack Overlap Support (ES-

POS) for each service pack. This eliminates the need for customers to purchase Long Term Service Pack Support (LTSS) until the final service pack of a release.

- **Dependencies:** Basesystem
- **Availability:** SLES for SAP (included)
- **Lifecycle:** 10 years
- **Extended support:** 3 years ESPOS, 3 years LTSS for last SP
- **Level:** L3

## 2.16. SAP Business One Server

This module contains packages and system configuration specific to SAP Business One Server. It is maintained and supported by the SUSE Linux Enterprise Server product subscription.

- **Dependencies:** Basesystem, Server Applications, Desktop Applications, Development Tools
- **Availability:** SLES (AMD64/Intel 64)
- **Lifecycle:** 10 years
- **Extended support:** 3 LTSS
- **Level:** L3

## 2.17. Server Applications

This module adds server functionality by providing network services such as DHCP server, name server or Web server. This module is selected for installation by default; deselecting it is not recommended.

- **Dependencies:** Basesystem
- **Availability:** SLES (default), SLES for SAP, SLE RT (default)
- **Lifecycle:** 10 years
- **Extended support:** 3 years LTSS
- **Level:** L3

## 2.18. SUSE Package Hub

This module provides access to packages for SUSE Linux Enterprise Server maintained by the openSUSE community. These packages are delivered without L3 support but do not interfere with the supportability of SUSE Linux Enterprise Server. For more information, refer to <https://package-hub.suse.com/>.

- **Dependencies:** Basesystem
- **Availability:** SLES, SLES for SAP, SLE RT, SLED
- **Lifecycle:** None

- **Extended support:** None
- **Level:** None

## 2.19. Systems Management

This module provides recent and more frequently updated versions of systems management packages such as Salt and Ansible.

- **Dependencies:** Python 3
- **Availability:** SLES, SLES for SAP, SLED
- **Lifecycle:** 6 years
- **Extended support:** 3 years LTSS
- **Level:** L3

## 2.20. Transactional Server

This module provides SUSE Linux Enterprise systems with a method of updating the operating system and its packages in an entirely “atomic” way. Updates are either applied to the system all together in a single transaction, or not at all. This happens without influencing the running system. If an update fails, or if a successful update is deemed to be incompatible or incorrect, you can discard it immediately and return to its previous functioning state.

- **Dependencies:** Basesystem
- **Availability:** SLES
- **Lifecycle:** 10 years
- **Extended support:** 3 years LTSS
- **Level:** L3

## 2.21. Web and Scripting

This module contains packages intended for a running Web server.

- **Dependencies:** Basesystem, Server Applications
- **Availability:** SLES, SLES for SAP
- **Lifecycle:** 10 years
- **Extended support:** 3 years LTSS
- **Level:** L3

## 2.22. Workstation Extension

This extension offers additional desktop applications and libraries. It is installed by default on SUSE Linux Enterprise Desktop. Adding the Workstation Extension to a SUSE Linux Enterprise

Server installation allows you to seamlessly combine both products to create a fully featured server workstation. For more information, refer to <https://www.suse.com/products/workstation-extension>.

### Unsupported kernel modules



The Workstation Extension installs the `kernel-default-extra` package, which includes additional kernel modules for desktop usage and enables loading of unsupported modules. If actually loaded, these modules may compromise your SLE support status.

Please evaluate if your hardware really requires additional kernel modules, and if so, test this expanded driver set in addition to the default set from the `kernel-default` package for possible configuration conflicts.

To prevent installation of the `kernel-default-extra` package, run `zypper addlock kernel-default-extra`.

- **Dependencies:** Basesystem, Desktop
- **Availability:** SLES, SLED (default), SLES for SAP
- **Lifecycle:** 10 years
- **Extended support:** None
- **Level:** mixed L2/L3 (depending on package)

## 3. Installing and managing modules and extensions

Modules and extensions can be installed when initially setting up the system and on an existing SUSE Linux Enterprise Server installation.

### 3.1. Installing modules and extensions during system installation

Modules and extensions can easily be installed when setting up the system. The respective steps are part of the installation routine. Following the default installation path, this requires network access to contact the SUSE Customer Center or a local registration server ([SUSE Multi-Linux Manager](#) or Repository Mirroring Tool), which provide the respective repositories. Offline installation is supported, too. In that case, an additional installation media is required. For detailed information, refer to *the section called “Extension and module selection”*.

### 3.2. Installing modules and extensions with YaST

To install modules & extensions from the running system, you need to make sure your system is registered with the SUSE Customer Center or a local registration server. If you have not done so during the installation, refer to *the section called “Registering from the installed system”* for instructions.


## Viewing already installed add-ons



To view already installed add-ons, start YaST and select *Software > Add-Ons*.

### Procedure 1. Installing add-ons and extensions from online channels with YaST

1. Start YaST and select *Software > System Extensions*. YaST connects to the registration server and displays a list of *Available Extensions and Modules*.



### Available extensions and modules



The number of available extensions and modules depends on the registration server. A local registration server may only offer update repositories and no additional extensions.

2. Click an entry to see its description.
3. To add modules or extensions, select all components you want to install. Note that all extensions require additional registration codes which are liable for cost.
4. Click *Next* to proceed.
5. Depending on the repositories to be added for the extension or module, you may be prompted to import the repository's GPG key or asked to agree to a license.

6. The YaST package manager opens to install release-packages for each module and, depending on your choice of modules and extensions, additional packages. It is strongly recommended *not to deselect* any of the preselected packages; you may, however, add additional packages.

Choose *Accept* and *Finish* to conclude the process.

7. If needed, adjust the repository *Properties* as described in *the section called “Managing repository properties”*.

### Module dependencies




Similar to software packages, which may depend on other packages to function, a module may have dependencies on other modules. If this is the case, the modules on which it depends are automatically selected for installation.

## 3.3. Installing extensions and third-party add-on products from media


When installing an extension or add-on product from media, you can select various types of product media, like DVD/CD, removable mass storage devices (such as flash disks), or a local directory or ISO image. The media can also be provided by a network server, for example, via HTTP, FTP, NFS, or Samba.

1. Start YaST and select *Software > Add-On Products*. Alternatively, start the YaST *Add-On Products* module from the command line with **sudo yast2 add-on**.

The dialog will show an overview of already installed add-on products, modules and extensions.



2. Click *Add* to install a new add-on product.
3. In the *Add-On Product* dialog, select the option that matches the type of medium from which you want to install:



- To scan your network for installation servers announcing their services via SLP, select *Scan Using SLP* and click *Next*.
- To add a repository from a removable medium, choose the relevant option and insert the medium or connect the USB device to the machine, respectively. Click *Next* to start the installation.
- For most media types, you will be prompted to specify the path (or URL) to the media after selecting the respective option and clicking *Next*. Specifying a *Repository Name* is optional. If none is specified, YaST will use the product name or the URL as the repository name.

The option *Download Repository Description Files* is activated by default. If you deactivate the option, YaST will automatically download the files later, if needed.


4. Depending on the repository you have added, you may be prompted to import the repository's GPG key or asked to agree to a license.

After confirming these messages, YaST will download and parse the metadata. It will add the repository to the list of *Configured Repositories*.

5. If needed, adjust the repository *Properties* as described in the section called “*Managing repository properties*”.
6. Confirm your changes with *OK* to close the configuration dialog.
7. After having successfully added the repository for the add-on media, the software manager starts, and you can install packages. For details, refer to *Chapter 8, Installing or removing software*.

### 3.4. Deleting modules and extensions with YaST

1. Start YaST > Software > Add-On Products.



2. Choose the module or extension that should be removed and click *Delete*. Confirm the warning saying that all packages from the selected component will be removed.
3. The YaST Software Manager opens and lists all installed packages from the deleted module or extension. Click *Accept* to remove all of them. It is strongly recommended to do so, because you will no longer get updates for packages from deleted modules or extensions. If you keep packages, make sure to at least remove the \*-release package for each module or extension that gets deleted.

Proceed with *Accept* and then *OK*.

#### Deleting modules



Note that you should never delete the *Basesystem Module*. It is also not recommended to delete the *Server Applications Module*.

#### No updates for packages from deleted modules and extensions



If you choose to keep packages from deleted modules or extensions, you will no longer receive updates for these packages. Because this includes security fixes, keeping such packages may introduce a security risk to your system.

### 3.5. Adding or deleting modules and extensions with SUSEConnect

1. Run **SUSEConnect --list-extensions** to get an overview of available extensions:

```

>sudo SUSEConnect --list-extensions
AVAILABLE EXTENSIONS AND MODULES

  Basesystem Module 15 SP7 x86_64 (Activated)
  Deactivate with: SUSEConnect -d -p sle-module-basesystem/15.7/x86_64

  Containers Module 15 SP7 x86_64
  Activate with: SUSEConnect -p sle-module-containers/15.7/x86_64

  Desktop Applications Module 15 SP7 x86_64
  Activate with: SUSEConnect -p sle-module-desktop-applications/15.7/
x86_64

  Development Tools Module 15 SP7 x86_64
  Activate with: SUSEConnect -p sle-module-development-tools/
15.7/x86_64

  NVIDIA Compute Module 15 x86_64
  Activate with: SUSEConnect -p sle-module-NVIDIA-compute/15/
x86_64

  SUSE Linux Enterprise Workstation Extension 15 SP7 x86_64
  Activate with: SUSEConnect -p sle-we/15.7/x86_64 -r ADDITIONAL
REGCODE

  Python 3 Module 15 SP7 x86_64
  Activate with: SUSEConnect -p sle-module-python3/15.7/x86_64

  SUSE Linux Enterprise Live Patching 15 SP7 x86_64
  Activate with: SUSEConnect -p sle-module-live-patching/15.7/x86_64
-r ADDITIONAL REGCODE

  SUSE Package Hub 15 SP7 x86_64
  Activate with: SUSEConnect -p PackageHub/15.7/x86_64

  Server Applications Module 15 SP7 x86_64 (Activated)
  Deactivate with: SUSEConnect -d -p sle-module-server-applications/
15.7/x86_64

  Legacy Module 15 SP7 x86_64
  Activate with: SUSEConnect -p sle-module-legacy/15.7/x86_64

  Public Cloud Module 15 SP7 x86_64
  Activate with: SUSEConnect -p sle-module-public-cloud/15.7/
x86_64

  SUSE Linux Enterprise High Availability Extension 15 SP7 x86_64
  Activate with: SUSEConnect -p sle-ha/15.7/x86_64 -r ADDITIONAL
REGCODE

  Web and Scripting Module 15 SP7 x86_64
  Activate with: SUSEConnect -p sle-module-web-scripting/15.7/
x86_64

  Transactional Server Module 15 SP7 x86_64
  Activate with: SUSEConnect -p sle-module-transactional-server/15.7/
x86_64

  Certifications Module 15 SP7 x86_64
  Activate with: SUSEConnect -p sle-module-certifications/15.7/x86_64

```

#### REMARKS

(Not available) The module/extension is not enabled on your RMT/SMT  
(Activated) The module/extension is activated on your system

#### MORE INFORMATION

You can find more information about available modules here:  
[https://www.suse.com/documentation/sles-15/singlehtml/art\\_modules/art\\_modules.html](https://www.suse.com/documentation/sles-15/singlehtml/art_modules/art_modules.html)

2. Run the commands in the listing for activating/deactivating a module or extension to add or delete a component. Note that adding an extension requires additional registration codes which are liable for cost.

### Deleting modules



Note that you should never delete the *Basesystem Module*. It is also not recommended to delete the *Server Applications Module*.

### No automatic installation or removal of packages



**SUSEConnect** only adds or removes modules and extensions. It registers or deregisters the components and enables or disables their repositories, but it does not install or remove any packages. If you want this to be done automatically, use YaST to manage modules and extensions.

When adding a module or extension, **SUSEConnect** does not install default packages or patterns. To do this manually, use Zypper or YaST > *Software Management*.

When deleting a module or extension, **SUSEConnect** does not perform a cleanup. Packages from the module or extension remain installed on the system, but are no longer updated from a repository. To list these “orphaned” packages, run **zypper packages --orphaned**. To remove one or more packages, run **zypper remove PACKAGE [ANOTHER\_PACKAGE]**. Alternatively, use YaST > *Software Management* and then *View > Package Classification > Orphaned Packages* to list and delete orphaned packages.

### No updates for packages from deleted modules and extensions



If you choose to keep packages from deleted modules or extensions, you will no longer receive updates for these packages. Because this includes security fixes, keeping such packages may introduce a security risk to your system.

## 4. Module membership of packages

By default, tools for searching software packages (YaST Software Management or **zypper search**) are restricted to searching within active repositories. For an extended search including not yet activated remote repositories, you can run

```
zypper search-packages PACKAGE
```

## Searching for packages in the SUSE Customer Center



You can also search for packages and their respective modules in the SUSE Customer Center Web interface at <https://scc.suse.com/packages>. A SUSE login is *not* required for this service.

To refine the search results, **zypper search-packages** offers the following options:

### **-x, --match-exact**

Searches for exact matches.

### **-g, --group-by-module**

Groups the search results by modules rather than by package name.

### **-d, --details**

Additionally lists version numbers and repositories containing the package.

### **--xmlout**

Generates XML output.

Below are a few examples of how to use **zypper search-packages**.

### Example 1. Simple search

The following command searches for the packages nodejs and hawk. The first one is available in the Web and Scripting module, while the second one is part of SUSE Linux Enterprise High Availability. Both packages are not yet installed:

```
>zypper search-packages openjdk nodejs
Following packages were found in following modules:
Package           Module or Repository
                  SUSEConnect Activation Command
-----
java-11-openjdk  Basesystem Module (sle-module-basesystem/15.4/x86_64)
                  SUSEConnect --product sle-module-basesystem/15.4/x86_64
java-17-openjdk  Basesystem Module (sle-module-basesystem/15.4/x86_64)
                  SUSEConnect --product sle-module-basesystem/15.4/x86_64
nodejs14         Web and Scripting Module (sle-module-web-scripting/15.4/x86_64)
                  SUSEConnect --product sle-module-web-scripting/15.4/x86_64
nodejs16         Web and Scripting Module (sle-module-web-scripting/15.4/x86_64)
                  SUSEConnect --product sle-module-web-scripting/15.4/x86_64
[...]
```

## Example 2. Searching for exact matches

If you know the exact package names, you can narrow the results down to exact matches by using the option `-x`:

```
>zypper search-packages -x java-11-openjdk nodejs14
Following packages were found in following modules:

Package           Module or Repository
                  SUSEConnect Activation Command
-----
-----
```

Package	Module or Repository
java-11-openjdk	Basesystem Module (sle-module-basesystem/15.4/x86_64) SUSEConnect --product sle-module-basesystem/15.4/x86_64
nodejs14	Web and Scripting Module (sle-module-web-scripting/15.4/x86_64) SUSEConnect --product sle-module-web-scripting/15.4/x86_64
[...]	

## Example 3. Searching for packages that are already installed or available

If you search for packages that are already installed or are available in one of the modules that are available on your system, the search results differ. In the following example, the package `vim` is already installed. The package `emacs` is not installed, but is available for installation without having to add a module.

```
>zypper search-packages -x vim
Following packages were found in following modules:

Package           Module or Repository
                  SUSEConnect Activation Command
-----
-----
```

Package	Module or Repository
vim	Basesystem Module (sle-module-basesystem/15.4/x86_64) SUSEConnect --product sle-module-basesystem/15.4/x86_64
vim	Installed
vim	Available in repo Basesystem_Module_15_SP4_x86_64:SLE-Module-Basesystem15-SP4-Pool
vim	Available in repo Basesystem_Module_15_SP4_x86_64:SLE-Module-Basesystem15-SP4-Updates

```
>zypper search-packages -x emacs
Following packages were found in following modules:

Package           Module or Repository
                  SUSEConnect Activation Command
-----
-----
```

Package	Module or Repository
emacs	Basesystem Module (sle-module-basesystem/15.4/x86_64) SUSEConnect --product sle-module-basesystem/15.4/x86_64
emacs	Available in repo Basesystem_Module_15_SP4_x86_64:SLE-Module-Basesystem15-SP4-Pool

Instead of searching for packages, you can also view the list of packages that belong to a module or extension. Start the YaST *Software Management* module and click *View > Services*. Select an entry to display the list of packages that belong to this module or extension.

## 5. Legal notice

Copyright© 2006– 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.

## 6. 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:

1. 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.
2. 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.
3. State on the Title page the name of the publisher of the Modified Version, as the publisher.
4. Preserve all the copyright notices of the Document.
5. Add an appropriate copyright notice for your modifications adjacent to the other copyright notices.

6. 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.
7. Preserve in that license notice the full lists of Invariant Sections and required Cover Texts given in the Document's license notice.
8. Include an unaltered copy of this License.
9. 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.
10. 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.
11. 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.
12. 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.
13. Delete any section Entitled "Endorsements". Such a section may not be included in the Modified Version.
14. Do not retitle any existing section to be Entitled "Endorsements" or to conflict in title with any Invariant Section.
15. 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.



### 1. Virtualization scenarios

Virtualization offers a lot of capabilities for your environment. It can be used in multiple scenarios. To get more details about it, refer to the *Virtualization Guide*, and in particular to the following sections:

- *the section called “Server consolidation”*
- *the section called “Virtualization benefits”*

This best practice guide provides advice for making the right choice in your environment. It recommends or discourages the use of options depending on your workload. Fixing configuration issues and performing tuning tasks increases the performance of VM Guests near to bare metal.

### 2. Before you apply modifications

#### 2.1. Back up first

Changing the configuration of the VM Guest or the VM Host Server can lead to data loss or an unstable state. It is really important that you do backups of files, data, images, etc. before making any changes. Without backups, you cannot restore the original state after a data loss or a misconfiguration. Do not perform tests or experiments on production systems.

#### 2.2. Test your workloads

The efficiency of a virtualization environment depends on many factors. This guide helps to make good choices when configuring virtualization in a production environment. Nothing is *carved in stone*. Hardware, workloads, resource capacity, etc. should all be considered when planning, testing and deploying your virtualization infrastructure. Testing your virtualized workloads is vital to a successful virtualization implementation.

## 3. Recommendations

### 3.1. Prefer the `libvirt` framework

SUSE strongly recommends using the `libvirt` framework to configure, manage and operate VM Host Servers and VM Guest. It offers a single interface (GUI and shell) for all supported virtualization technologies and therefore is easier to use than the hypervisor-specific tools.

We do not recommend using `libvirt` and hypervisor-specific tools at the same time, because changes done with the hypervisor-specific tools may not be recognized by the `libvirt` tool set. See *Chapter 8, libvirt daemons* for more information on `libvirt`.

### 3.2. `qemu-system-i386` compared to `qemu-system-x86_64`

Similar to real 64-bit PC hardware, `qemu-system-x86_64` supports VM Guests running a 32-bit or a 64-bit operating system. Because `qemu-system-x86_64` usually also provides better performance for 32-bit guests, SUSE recommends using `qemu-system-x86_64` for both 32-bit and 64-bit VM Guests on KVM. Scenarios where `qemu-system-i386` performs better are not supported by SUSE.

Xen also uses binaries from the `qemu` package but prefers `qemu-system-i386`, which can be used for both 32-bit and 64-bit Xen VM Guests. To maintain compatibility with the upstream Xen Community, SUSE encourages using `qemu-system-i386` for Xen VM Guests.

## 4. VM Host Server configuration and resource allocation

Allocation of resources for VM Guests is a crucial point when administrating virtual machines. When assigning resources to VM Guests, be aware that overcommitting resources may affect the performance of the VM Host Server and the VM Guests. If all VM Guests request all their resources simultaneously, the host needs to provide them all. If not, the host's performance is negatively affected and this in turn also has negative effects on the VM Guest's performance.

### 4.1. Memory

Linux manages memory in units called pages. On most systems the default page size is 4 KB. Linux and the CPU need to know which pages belong to which process. That information is stored in a page table. If a lot of processes are running, it takes more time to find where the memory is mapped, because of the time required to search the page table. To speed up the search, the TLB (Translation Lookaside Buffer) was invented. But on a system with a lot of memory, the TLB is not enough. To avoid any fallback to normal page table (resulting in a cache miss, which is time consuming), huge pages can be used. Using huge pages will reduce TLB overhead and TLB misses (pagewalk). A host with 32 GB ( $32*1024*1024 = 33,554,432$  KB) of memory and a 4 KB page size has a TLB with  $33,554,432/4 = 8,388,608$  entries. Using a 2 MB (2048 KB) page size, the TLB only has  $33554432/2048 = 16384$  entries, considerably reducing the TLB misses.

#### 4.1.1. Configuring the VM Host Server and the VM Guest to use huge pages

The AMD64/Intel 64 CPU architecture supports larger pages than 4 KB: huge pages. To determine the size of huge pages available on your system (could be 2 MB or 1 GB), check the flags line in the output of `/proc/cpuinfo` for occurrences of `pse` and/or `pdpe1gb`.

**Table 1. Determine the available huge pages size**

CPU flag	Huge pages size available
Empty string	No huge pages available
<code>pse</code>	2 MB
<code>pdpe1gb</code>	1 GB

Using huge pages improves the performance of VM Guests and reduces host memory consumption.

By default, the system uses THP. To make huge pages available on your system, activate it at boot time with `hugepages=1`, and—optionally—add the huge pages size with, for example, `hugepagesz=2MB`.



##### 1 GB huge pages

1 GB pages can only be allocated at boot time and cannot be freed afterward.

To allocate and use the huge page table (HugeTlbPage), you need to mount `hugetlbfs` with correct permissions.



##### Restrictions of huge pages

Even if huge pages provide the best performance, they do come with some drawbacks. You lose features such as Memory ballooning (see the section called “*virtio balloon*”), KSM (see the section called “*KSM and page sharing*”), and huge pages cannot be swapped.

#### Procedure 1. Configuring the use of huge pages

1. Mount `hugetlbfs` to `/dev/hugepages`:

```
>sudo mount -t hugetlbfs hugetlbfs /dev/hugepages
```

2. To reserve memory for huge pages, use the **sysctl** command. If your system has a huge page size of 2 MB (2048 KB), and you want to reserve 1 GB (1,048,576 KB) for your VM Guest, you need  $1,048,576/2048=512$  pages in the pool:

```
>sudo sysctl vm.nr_hugepages=512
```

The value is written to `/proc/sys/vm/nr_hugepages` and represents the current number of *persistent* huge pages in the kernel's huge page pool. *Persistent* huge pages are returned to the huge page pool when freed by a task.

3. Add the `memoryBacking` element in the VM Guest configuration file (by running **virsh edit** *CONFIGURATION*).

```
<memoryBacking>
  <hugepages/>
</memoryBacking>
```

4. Start your VM Guest and check on the host whether it uses hugepages:

```
>cat /proc/meminfo | grep HugePages_
HugePages_Total:① 512
HugePages_Free:② 92
HugePages_Rsvd:③ 0
HugePages_Surp:④ 0
```

- ① Size of the pool of huge pages
- ② Number of huge pages in the pool that are not yet allocated
- ③ Number of huge pages for which a commitment to allocate from the pool has been made, but no allocation has yet been made
- ④ Number of huge pages in the pool above the value in `/proc/sys/vm/nr_hugepages`. The maximum number of surplus huge pages is controlled by `/proc/sys/vm/nr_overcommit_hugepages`

#### 4.1.2. Transparent huge pages

Transparent huge pages (THP) provide a way to dynamically allocate huge pages with the **khugepaged** kernel thread, rather than manually managing their allocation and use. Workloads with contiguous memory access patterns can benefit greatly from THP. A 1000 fold decrease in page faults can be observed when running synthetic workloads with contiguous memory access patterns. Conversely, workloads with sparse memory access patterns (like databases) may perform poorly with THP. In such cases, it may be preferable to disable THP by adding the kernel parameter `transparent_hugepage=never`, rebuild your grub2 configuration, and reboot. Verify if THP is disabled with:

```
>cat /sys/kernel/mm/transparent_hugepage/enabled
always madvise [never]
```

If disabled, the value `never` is shown in square brackets like in the example above.



## Xen

THP is not available under Xen.

### 4.1.3. Xen-specific memory notes

#### 4.1.3.1. Managing domain-0 memory

In previous versions of SUSE Linux Enterprise Server, the default memory allocation scheme of a Xen host was to allocate all host physical memory to Dom0 and enable auto-ballooning. Memory was automatically ballooned from Dom0 when additional domains were started. This behavior has always been error prone and disabling it was strongly encouraged. Starting in SUSE Linux Enterprise Server 15 SP1, auto-ballooning has been disabled by default and Dom0 is given 10% of host physical memory + 1 GB. For example, on a host with 32 GB of physical memory, 4.2 GB of memory is allocated to Dom0.

The use of `dom0_mem` Xen command-line option in `/etc/default/grub` is still supported and encouraged (see *the section called “Change kernel parameters at boot time”* for more information). You can restore the old behavior by setting `dom0_mem` to the host physical memory size and enabling the `autoballoon` setting in `/etc/xen/xl.conf`.

### 4.1.4. KSM and page sharing

Kernel Samepage Merging is a kernel feature that reduces memory consumption on the VM Host Server by sharing blocks of memory that VM Guests have in common. The KSM daemon `ksmd` periodically scans user memory, looking for pages with identical contents, which can be replaced by a single write-protected page. To enable the KSM service, first make sure that the package `qemu-ksm` is installed, then run the command:

```
>sudo systemctl enable --now ksm.service
```

Alternatively, it can also be started by running the command:

```
# echo 1 > /sys/kernel/mm/ksm/run
```

One advantage of using KSM from a VM Guest's perspective is that all guest memory is backed by host anonymous memory. You can share `pagecache`, `tmpfs` or any kind of memory allocated in the guest.

KSM is controlled by `sysfs`. You can check KSM's values in `/sys/kernel/mm/ksm/`:

- `pages_shared`: the number of shared pages that are being used (read-only).
- `pages_sharing`: the number of sites sharing the pages (read-only).

- `pages_unshared`: the number of pages that are unique and repeatedly checked for merging (read-only).
- `pages_volatile`: the number of pages that are changing too fast to be considered for merging (read-only).
- `full_scans`: the number of times all mergeable areas have been scanned (read-only).
- `sleep_millisecs`: the number of milliseconds `ksmd` should sleep before the next scan. A low value will overuse the CPU, consuming CPU time that could be used for other tasks. We recommend a value greater than 1000.
- `pages_to_scan`: the number of present pages to scan before `ksmd` goes to sleep. A high value will overuse the CPU. We recommend starting with a value of 1000 and then adjusting based on the KSM results observed while testing your deployment.
- `merge_across_nodes`: by default, the system merges pages across NUMA nodes. Set this option to 0 to disable this behavior.

### Use cases



KSM is a good technique to over-commit host memory when running multiple instances of the same application or VM Guest. When applications and VM Guest are heterogeneous and do not share any common data, it is preferable to disable KSM. To do that, run:

```
>sudo systemctl disable --now ksm.service
```

Alternatively, it can also be disabled by running the command:

```
# echo 0 > /sys/kernel/mm/ksm/run
```

In a mixed heterogeneous and homogeneous environment, KSM can be enabled on the host but disabled on a per VM Guest basis. Use `virsh edit` to disable page sharing of a VM Guest by adding the following to the guest's XML configuration:

```
<memoryBacking>
  <nosharepages/>
</memoryBacking>
```

### Avoid out-of-memory conditions



KSM can free up certain memory on the host system, but the administrator should reserve enough swap to avoid out-of-memory conditions if that shareable memory decreases. If the amount of shareable memory decreases, the use of physical memory is increased.

## KSM as a side channel



Because of its nature, KSM can form a side channel between otherwise isolated guests. It is discouraged to enable KSM in environments where guests from different security domains are executed.

## Memory access latencies



By default, KSM will merge common pages across NUMA nodes. If the merged, common page is now located on a distant NUMA node (relative to the node running the VM Guest vCPUs), this may degrade VM Guest performance. If increased memory access latencies are noticed in the VM Guest, disable cross-node merging with the `merge_across_nodes` sysfs control:

```
# echo 0 > /sys/kernel/mm/ksm/merge_across_nodes
```

### 4.1.5. VM Guest: memory hotplug

To optimize the usage of your host memory, it may be useful to hotplug more memory for a running VM Guest when required. To support memory hotplugging, you must first configure the `<maxMemory>` tag in the VM Guest's configuration file:

```
<maxMemory① slots='16'② unit='KiB'>20971520③</maxMemory>
  <memory④ unit='KiB'>1048576</memory>
  <currentMemory⑤ unit='KiB'>1048576</currentMemory>
```

① Runtime maximum memory allocation of the guest

② Number of slots available for adding memory to the guest

③ Valid units are:

- “KB” for kilobytes (1,000 bytes)
- “k” or “KiB” for kibibytes (1,024 bytes)
- “MB” for megabytes (1,000,000 bytes)
- “M” or “MiB” for mebibytes (1,048,576 bytes)
- “GB” for gigabytes (1,000,000,000 bytes)
- “G” or “GiB” for gibibytes (1,073,741,824 bytes)
- “TB” for terabytes (1,000,000,000,000 bytes)
- “T” or “TiB” for tebibytes (1,099,511,627,776 bytes)

④ Maximum allocation of memory for the guest at boot time

## 5 Actual allocation of memory for the guest

To hotplug memory devices into the slots, create a file `mem-dev.xml` like the following:

```
<memory model='dimm'>
  <target>
    <size unit='KiB'>524287</size>
    <node>0</node>
  </target>
</memory>
```

And attach it with the following command:

```
>virsh attach-device vm-name mem-dev.xml
```

For memory device hotplug, the guest must have at least 1 NUMA cell defined (see *the section called “VM Guest virtual NUMA topology”*).

## 4.2. Swap

**Swap** is used by the system to store underused physical memory (low usage, or not accessed for a long time). To prevent the system from running out of memory, setting up a minimum swap is highly recommended.

### 4.2.1. swappiness

The **swappiness** setting controls your system's swap behavior. It defines how memory pages are swapped to disk. A high value of **swappiness** results in a system that swaps more often. Available values range from 0 to 200. A value of 200 tells the system to find inactive pages and put them in swap. A value of 0 disables swapping.

To test on a live system, change the value of `/proc/sys/vm/swappiness` on the fly and check the memory usage afterward:

```
# echo 35 > /proc/sys/vm/swappiness
```

```
>free -h
total        used         free        shared       buffers       cached
Mem:      24616680      4991492      19625188      167056      144340      2152408
-/+ buffers/cache:  2694744  21921936
Swap:      6171644          0      6171644
```

To permanently set a **swappiness** value, add a line in `/etc/systcl.conf`, for example:

```
vm.swappiness = 35
```

You can also control the swap by using the `swap_hard_limit` element in the XML configuration of your VM Guest. Before setting this parameter and using it in a production environment, test it because the host can stop the domain if the value is too low.

```

<memtune>❶
  <hard_limit unit='G'>1</hard_limit>❷
  <soft_limit unit='M'>128</soft_limit>❸
  <swap_hard_limit unit='G'>2</swap_hard_limit>❹
</memtune>

```

- ❶ This element provides memory tunable parameters for the domain. If this is omitted, it defaults to the defaults provided by the operating system.
- ❷ Maximum memory the guest can use. To avoid any problems on the VM Guest, we strongly recommend not to use this parameter.
- ❸ The memory limit to enforce during memory contention.
- ❹ The maximum memory plus swap the VM Guest can use.

## 4.3. I/O

### 4.3.1. I/O scheduler

The I/O scheduler for SUSE Linux Enterprise 15 SP2 and up is Budget Fair Queueing (BFQ). The main aim of the BFQ scheduler is to provide a fair allocation of the disk I/O bandwidth for all processes that request an I/O operation. You can have different I/O schedulers for different devices.

To get better performance in host and VM Guest, use `none` in the VM Guest (disable the I/O scheduler) and the `mq-deadline` scheduler for a virtualization host.

#### Procedure 2. Checking and changing the I/O scheduler at runtime

1. To check your current I/O scheduler for your disk (replace `sdX` by the disk you want to check), run:

```

>cat /sys/block/sdX/queue/scheduler
mq-deadline kyber [bfq] none

```

The value in square brackets is the one currently selected (bfq in the example above).

2. You can change the scheduler at runtime by running the following command as root:

```
# echo mq-deadline > /sys/block/sdX/queue/scheduler
```

If you need to specify different I/O schedulers for each disk, create the file `/usr/lib/tmpfiles.d/I0_ioscheduler.conf` with content similar to the following example. It defines the `mq-deadline` scheduler for `/dev/sda` and the `none` scheduler for `/dev/sdb`. Keep in mind that the device name can be different depending on the device type. This feature is available on SLE 12 and up.

```

w /sys/block/sda/queue/scheduler - - - - mq-deadline
w /sys/block/sdb/queue/scheduler - - - - none

```

### 4.3.2. Asynchronous I/O

Many of the virtual disk back-ends use Linux Asynchronous I/O (aio) in their implementation. By default, the maximum number of aio contexts is set to 65536, which can be exceeded when running hundreds of VM Guests using virtual disks serviced by Linux Asynchronous I/O. When running large numbers of VM Guests on a VM Host Server, consider increasing `/proc/sys/fs/aio-max-nr`.

#### Procedure 3. Checking and changing aio-max-nr at runtime

1. To check your current aio-max-nr setting run:

```
>cat /proc/sys/fs/aio-max-nr
65536
```

2. You can change aio-max-nr at runtime with the following command:

```
# echo 131072 > /proc/sys/fs/aio-max-nr
```

To permanently set aio-max-nr, add an entry to a custom sysctl file. For example, include the following to `/etc/sysctl.d/aio-max-nr.conf`:

```
fs.aio-max-nr = 1048576
```

### 4.3.3. I/O Virtualization

SUSE products support multiple I/O virtualization technologies. The following table lists advantages and disadvantages of each technology. For more information about I/O in virtualization refer to the *the section called “I/O virtualization”*.

**Table 2. I/O Virtualization solutions**

Technology	Advantage	Disadvantage
Device Assignment (pass-through)	Device accessed directly by the guest	No sharing among multiple guests
	High performance	Live migration is complex
		PCI device limit is 8 per guest
		Limited number of slots on a server
Full virtualization (IDE, SATA, SCSI, e1000)	VM Guest compatibility	Bad performance
	Easy for live migration	Emulated operation

Technology	Advantage	Disadvantage
Para-virtualization (virtio-blk, virtio-net, virtio-scsi)	Good performance	Modified guest (PV drivers)
	Easy for live migration	
	Efficient host communication with VM Guest	

#### 4.4. Storage and file system

Storage space for VM Guests can either be a block device, for example, a partition on a physical disk, or an image file on the file system:

**Table 3. Block devices compared to disk images**

Technology	Advantages	Disadvantages
Block devices	<ul style="list-style-type: none"> <li>• Better performance</li> <li>• Use standard tools for administration/disk modification</li> <li>• Accessible from host (pro and con)</li> </ul>	<ul style="list-style-type: none"> <li>• Device management</li> </ul>
Image files	<ul style="list-style-type: none"> <li>• Easier system management</li> <li>• Easily move, clone, expand, back up domains</li> <li>• Comprehensive toolkit (guestfs) for image manipulation</li> <li>• Reduce overhead through sparse files</li> <li>• Fully allocate for best performance</li> </ul>	<ul style="list-style-type: none"> <li>• Lower performance than block devices</li> </ul>

For detailed information about image formats and maintaining images, refer to *the section called “VM Guest images”*.

If your image is stored on an NFS share, check certain server and client parameters to improve access to the VM Guest image.

#### 4.4.1. NFS read/write (client)

Options `rsize` and `wsize` specify the size of the chunks of data that the client and server pass back and forth to each other. You should ensure NFS read/write sizes are sufficiently large, especially for large I/O. Change the `rsize` and `wsize` parameter in your `/etc/fstab` by increasing the value to 16 KB. This will ensure that all operations can be frozen if there is any instance of hanging.

```
nfs_server:/exported/vm_images① /mnt/images② nfs③ rw④,hard⑤,sync⑥,
rsize=8192⑦,wsize=8192⑧ 0 0
```

- ① NFS server's host name and export path.
- ② Where to mount the NFS exported share.
- ③ This is an `nfs` mount point.
- ④ This mount point will be accessible in read/write.
- ⑤ Determines the recovery behavior of the NFS client after an NFS request times out. `hard` is the best option to avoid data corruption.
- ⑥ Any system call that writes data to files on that mount point causes that data to be flushed to the server before the system call returns control to user space.
- ⑦ Maximum number of bytes in each network READ request that the NFS client can receive when reading data from a file on an NFS server.
- ⑧ Maximum number of bytes per network WRITE request that the NFS client can send when writing data to a file on an NFS server.

#### 4.4.2. NFS threads (server)

Your NFS server should have enough NFS threads to handle multi-threaded workloads. Use the `nfsstat` tool to get RPC statistics on your server:

```
>sudo nfsstat -rc
Client rpc stats:
calls      retrans      authrefrsh
6401066      198          0          0
```

If the `retrans` is equal to 0, everything is fine. Otherwise, the client needs to retransmit, so increase the `USE_KERNEL_NFSD_NUMBER` variable in `/etc/sysconfig/nfs`, and adjust accordingly until `retrans` is equal to 0.

## 4.5. CPUs

Host CPU “components” will be “translated” to virtual CPUs in a VM Guest when being assigned. These components can either be:

- *CPU processor*: this describes the main CPU unit, which normally has multiple cores and may support Hyper-Threading.
- *CPU core*: a main CPU unit can provide more than one core, and the proximity of cores speeds up the computation process and reduces energy costs.
- *CPU Hyper-Threading*: this implementation is used to improve the parallelization of computations, but this is not as efficient as a dedicated core.

### 4.5.1. Assigning CPUs

CPU overcommit occurs when the cumulative number of virtual CPUs of all VM Guests becomes higher than the number of host CPUs. Best performance is achieved when there is no overcommit and each virtual CPU matches one hardware processor or core on the VM Host Server. VM Guests running on an overcommitted host will experience increased latency and a negative effect on per-VM Guest throughput is often observed. Therefore, try to avoid overcommitting CPUs.

Deciding whether to allow CPU overcommit or not requires good a priori knowledge of workload as a whole. For example, if you know that all the VM Guests' virtual CPUs will not be loaded over 50%, then you can assume that overcommitting the host by a factor of 2 (which means having 128 virtual CPUs in total, on a host with 64 CPUs) will work well. However, if you know that all the virtual CPUs of the VM Guests will try to run at 100% for most of the time, then even having one virtual CPU more than the host has CPUs is already a misconfiguration.

Overcommitting to a point where the cumulative number of virtual CPUs is higher than 8 times the number of physical cores of the VM Host Server may lead to a malfunctioning and unstable system and should hence be avoided.

Unless you know exactly how many virtual CPUs are required for a VM Guest, start with one. Target a CPU workload of approximately 70% inside your VM (see *the section called “Processes”* for information on monitoring tools). If you allocate more processors than needed in the VM Guest, this will negatively affect the performance of host and guest. Cycle efficiency will be degraded, as the unused vCPU will still cause timer interrupts. In case you primarily run single threaded applications on a VM Guest, a single virtual CPU is the best choice.

A single VM Guest with more virtual CPUs than the VM Host Server has CPUs is always a misconfiguration.

## 4.5.2. VM Guest CPU configuration

This section describes how to choose and configure a CPU type for a VM Guest. You will also learn how to pin virtual CPUs to physical CPUs on the host system. For more information about virtual CPU configuration and tuning parameters, refer to the libvirt documentation at <https://libvirt.org/formatdomain.html#elementsCPU>.

### 4.5.2.1. Virtual CPU models and features

The CPU model and topology can be specified individually for each VM Guest. Configuration options range from selecting specific CPU models to excluding certain CPU features. Predefined CPU models are listed in files in the directory `/usr/share/libvirt/cpu_map/`. A CPU model and topology that is similar to the host generally provides the best performance. The host system CPU model and topology can be displayed by running **virsh capabilities**.

Changing the default virtual CPU configuration will require a VM Guest shutdown when migrating it to a host with different hardware. More information on VM Guest migration is available in *Chapter 17, Migrating VM Guests*.

To specify a particular CPU model for a VM Guest, add a respective entry to the VM Guest configuration file. The following example configures a Broadwell CPU with the invariant TSC feature:

```
<cpu mode='custom' match='exact'>
  <model>Broadwell</model>
  <feature name='invsc' />
</cpu>
```

For a virtual CPU that most closely resembles the host physical CPU, `<cpu mode='host-passthrough'>` can be used. A host-passthrough CPU model may not exactly resemble the host physical CPU, since, by default, KVM will mask any non-migratable features. For example, `invsc` is not included in the virtual CPU feature set. Changing the default KVM behavior is not directly supported through libvirt, although it does allow arbitrary pass-through of KVM command-line arguments. Continuing with the `invsc` example, you can achieve pass-through of the host CPU (including `invsc`) with the following command-line pass-through in the VM Guest configuration file:

```
<domain type='kvm' xmlns:qemu='http://libvirt.org/schemas/domain/qemu/1.0'>
  <qemu:commandline>
    <qemu:arg value=''-cpu' />
    <qemu:arg value='host,migratable=off,+invsc' />
  </qemu:commandline>
  ...
</domain>
```

## The host-passthrough mode



Since host-passthrough exposes the physical CPU details to the virtual CPU, migration to dissimilar hardware is not possible. See *the section called “Virtual CPU migration considerations”* for more information.

### 4.5.2.2. Virtual CPU pinning

Virtual CPU pinning is used to constrain virtual CPU threads to a set of physical CPUs. The `vcpupin` element specifies the physical host CPUs that a virtual CPU can use. If this element is not set and the attribute `cpuset` of the `vcpu` element is not specified, the virtual CPU is free to use any of the physical CPUs.

CPU intensive workloads can benefit from virtual CPU pinning by increasing the physical CPU cache hit ratio. To pin a virtual CPU to a specific physical CPU, run the following commands:

```
>virsh vcpupin DOMAIN_ID --vcpu vCPU_NUMBER
VCPU: CPU Affinity
-----
0: 0-7
#virsh vcpupin SLE15 --vcpu 0 0 --config
```

The last command generates the following entry in the XML configuration:

```
<cputune>
  <vcpu pin='0' cpuset='0' />
</cputune>
```

#### Virtual CPU pinning on NUMA nodes



To confine a VM Guest's CPUs and its memory to a NUMA node, you can use virtual CPU pinning and memory allocation policies on a NUMA system. See *the section called “NUMA tuning”* for more information related to NUMA tuning.

#### Virtual CPU pinning and live migration



Even though `vcpu pin` can improve performance, it can complicate live migration. See *the section called “Virtual CPU migration considerations”* for more information on virtual CPU migration considerations.

### 4.5.2.3. Virtual CPU migration considerations

Selecting a virtual CPU model containing all the latest features may improve performance of a VM Guest workload, but often at the expense of migratability. Unless all hosts in the cluster contain the

latest CPU features, migration can fail when a destination host lacks the new features. If migratability of a virtual CPU is preferred over the latest CPU features, a normalized CPU model and feature set should be used. The **virsh cpu-baseline** command can help define a normalized virtual CPU that can be migrated across all hosts. The following command, when run on each host in the migration cluster, illustrates the collection of all hosts' capabilities in `all-hosts-caps.xml`.

```
>sudo virsh capabilities >> all-hosts-cpu-caps.xml
```

With the capabilities of each host collected in `all-hosts-caps.xml`, use **virsh cpu-baseline** to create a virtual CPU definition that will be compatible across all hosts.

```
>sudo virsh cpu-baseline all-hosts-caps.xml
```

The resulting virtual CPU definition can be used as the `cpu` element in the VM Guest configuration file.

At a logical level, virtual CPU pinning is a form of hardware pass-through. CPU pinning couples physical resources to virtual resources, which can also be problematic for migration. For example, the migration will fail if the requested physical resources are not available on the destination host, or if the source and destination hosts have different NUMA topologies. For more recommendations about Live Migration, see the section called “*Migration requirements*”.

## 4.6. NUMA tuning

NUMA is an acronym for Non Uniform Memory Access. A NUMA system has multiple physical CPUs, each with local memory attached. Each CPU can also access other CPUs' memory, known as “remote memory access”, but it is much slower than accessing local memory. NUMA systems can negatively affect VM Guest performance if not tuned properly. Although ultimately tuning is workload dependent, this section describes controls that should be considered when deploying VM Guests on NUMA hosts. Always consider your host topology when configuring and deploying VMs.

SUSE Linux Enterprise Server contains a NUMA auto-balancer that strives to reduce remote memory access by placing memory on the same NUMA node as the CPU processing it. Standard tools such as **cgset** and virtualization tools such as libvirt provide mechanisms to constrain VM Guest resources to physical resources.

**numactl** is used to check for host NUMA capabilities:

```
>sudo numactl --hardware
available: 4 nodes (0-3)
node 0 cpus: 0 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 72 73 74 75 76 77 78
79 80 81 82 83 84 85 86 87 88 89
node 0 size: 31975 MB
node 0 free: 31120 MB
node 1 cpus: 18 19 20 21 22 23 24 25 26 27 28 29 30 31 32 33 34 35 90 91 92 93
94 95 96 97 98 99 100 101 102 103 104 105 106 107
node 1 size: 32316 MB
node 1 free: 31673 MB
node 2 cpus: 36 37 38 39 40 41 42 43 44 45 46 47 48 49 50 51 52 53 108 109 110
111 112 113 114 115 116 117 118 119 120 121 122 123 124 125
node 2 size: 32316 MB
node 2 free: 31726 MB
node 3 cpus: 54 55 56 57 58 59 60 61 62 63 64 65 66 67 68 69 70 71 126 127 128
129 130 131 132 133 134 135 136 137 138 139 140 141 142 143
node 3 size: 32314 MB
node 3 free: 31387 MB
node distances:
node 0 1 2 3
0: 10 21 21 21
1: 21 10 21 21
2: 21 21 10 21
3: 21 21 21 10
```

The **numactl** output shows this is a NUMA system with 4 nodes or cells, each containing 36 CPUs and approximately 32G memory. **virsh capabilities** can also be used to examine the systems NUMA capabilities and CPU topology.

#### 4.6.1. NUMA balancing

On NUMA machines, there is a performance penalty if remote memory is accessed by a CPU. Automatic NUMA balancing scans a task's address space and unmaps pages. By doing so, it detects whether pages are properly placed or whether to migrate the data to a memory node local to where the task is running. In defined intervals (configured with `numa_balancing_scan_delay_ms`), the task scans the next scan size number of pages (configured with `numa_balancing_scan_size_mb`) in its address space. When the end of the address space is reached, the scanner restarts from the beginning.

Higher scan rates cause higher system overhead as page faults must be trapped and data needs to be migrated. However, the higher the scan rate, the more quickly a task's memory migrates to a local node when the workload pattern changes. This minimizes the performance impact caused by remote memory accesses. These **sysctl** directives control the thresholds for scan delays and the number of pages scanned:

```
>sudo sysctl -a | grep numa_balancing
kernel.numa_balancing = 1❶
kernel.numa_balancing_scan_delay_ms = 1000❷
kernel.numa_balancing_scan_period_max_ms = 60000❸
kernel.numa_balancing_scan_period_min_ms = 1000❹
kernel.numa_balancing_scan_size_mb = 256❺
```

❶ Enables/disables automatic page fault-based NUMA balancing

❷ Starting scan delay used for a task when it initially forks

- ❸ Maximum time in milliseconds to scan a task's virtual memory
- ❹ Minimum time in milliseconds to scan a task's virtual memory
- ❺ Size in megabytes' worth of pages to be scanned for a given scan

For more information, see *Chapter 11, Automatic Non-Uniform Memory Access (NUMA) balancing*.

The main goal of automatic NUMA balancing is either to reschedule tasks on the same node's memory (so the CPU follows the memory), or to copy the memory's pages to the same node (so the memory follows the CPU).

### Task placement



There are no rules to define the best place to run a task, because tasks could share memory with other tasks. For the best performance, we recommend to group tasks sharing memory on the same node. Check NUMA statistics with `# cat /proc/vmstat | grep numa_`

#### 4.6.2. Memory allocation control with the CPUset controller

The cgroups cpuset controller can be used confine memory used by a process to a NUMA node. There are three cpuset memory policy modes available:

- **interleave**: this is a memory placement policy which is also known as round-robin. This policy can provide substantial improvements for jobs that need to place thread local data on the corresponding node. When the interleave destination is not available, it will be moved to another node.
- **bind**: this will place memory only on one node, which means in case of insufficient memory, the allocation will fail.
- **preferred**: this policy will apply a preference to allocate memory to a node. If there is not enough space for memory on this node, it will fall back to another node.

You can change the memory policy mode with the **cgset** tool from the **libcgroup-tools** package:

```
>sudo cgset -r cpuset.mems=NODE sysdefault/libvirt/qemu/KVM_NAME/emulator
```

To migrate pages to a node, use the **migratepages** tool:

```
>migratepages PIDFROM-NODETO-NODE
```

To check everything is fine, use: **cat /proc/PID/status | grep Cpus**.

## Kernel NUMA/cpuset memory policy



For more information, see [Kernel NUMA memory policy](#) and [cpusets memory policy](#). Check also the [Libvirt NUMA Tuning documentation](#).

### 4.6.3. VM Guest: NUMA related configuration

libvirt allows to set up virtual NUMA and memory access policies. Configuring these settings is not supported by **virt-install** or **virt-manager** and needs to be done manually by editing the VM Guest configuration file with **virsh edit**.

#### 4.6.3.1. VM Guest virtual NUMA topology

Creating a VM Guest virtual NUMA (vNUMA) policy that resembles the host NUMA topology can often increase performance of traditional large, scale-up workloads. VM Guest vNUMA topology can be specified using the `numa` element in the XML configuration:

```
<cpu>
  ...
  <numa>
    <cell id="0" cpus='0-1' memory='512000' unit='KiB' />
    <cell id="1" cpus='2-3' memory='256000' unit='KiB' />
    <memAccess='shared' />
  </numa>
  ...
</cpu>
```

- ❶ Each `cell` element specifies a vNUMA cell or node
- ❷ All cells should have an `id` attribute, allowing to reference the cell in other configuration blocks. Otherwise, cells are assigned IDs in ascending order starting from 0.
- ❸ The CPU or range of CPUs that are part of the node
- ❹ The node memory
- ❺ Units in which node memory is specified
- ❻ Optional attribute which can control whether the memory is to be mapped as shared or private. This is valid only for hugepages-backed memory.

To find where the VM Guest has allocated its pages. use: **cat /proc/PID/numa\_maps** and **cat /sys/fs/cgroup/memory/sysdefault/libvirt/qemu/KVM\_NAME/memory.numa\_stat**.

## NUMA specification



The libvirt VM Guest NUMA specification is currently only available for QEMU/KVM.

### 4.6.3.2. Memory allocation control with libvirt

If the VM Guest has a vNUMA topology (see the section called “VM Guest virtual NUMA topology”), memory can be pinned to host NUMA nodes using the `numatune` element. This method is currently only available for QEMU/KVM guests. See *Non-vNUMA VM Guest* for how to configure non-vNUMA VM Guests.

```
<numatune>
  <memory mode="strict"① nodeset="1-4,^3"②/>
  <memnode③ cellid="0"④ mode="strict" nodeset="1"/>
  <memnode cellid="2" placement="strict"⑤ mode="preferred" nodeset="2"/>
</numatune>
```

- ① Policies available are: `interleave` (round-robin like), `strict` (default) or `preferred`.
- ② Specify the NUMA nodes.
- ③ Specify memory allocation policies for each guest NUMA node (if this element is not defined, then this will fall back and use the `memory` element).
- ④ Addresses the guest NUMA node for which the settings are applied.
- ⑤ The placement attribute can be used to indicate the memory placement mode for a domain process, the value can be `auto` or `strict`.

## Non-vNUMA VM Guest



On a non-vNUMA VM Guest, pinning memory to host NUMA nodes is done as in the following example:

```
<numatune>
  <memory mode="strict" nodeset="0-1"/>
</numatune>
```

In this example, memory is allocated from the host nodes 0 and 1. If these memory requirements cannot be fulfilled, starting the VM Guest will fail. `virt-install` also supports this configuration with the `--numatune` option.

## Memory and CPU across NUMA nodes



You should avoid allocating VM Guest memory across NUMA nodes, and prevent virtual CPUs from floating across NUMA nodes.

## 5. VM Guest images

Images are virtual disks used to store the operating system and data of VM Guests. They can be created, maintained and queried with the **qemu-img** command. Refer to the section called “*Creating, converting, and checking disk images*” for more information on the **qemu-img** tool and examples.

### 5.1. VM Guest image formats

Certain storage formats which QEMU recognizes have their origins in other virtualization technologies. By recognizing these formats, QEMU can use either data stores or entire guests that were originally targeted to run under these other virtualization technologies. Certain formats are supported only in read-only mode. To use them in read/write mode, convert them to a fully supported QEMU storage format (using **qemu-img**). Otherwise they can only be used as read-only data store in a QEMU guest.

Use **qemu-img info VMGUEST.IMG** to get information about an existing image, such as the format, the virtual size, the physical size and snapshots, if available.

#### Performance



It is recommended to convert the disk images to either raw or qcow2 to achieve good performance.

#### Encrypted images cannot be compressed



When you create an image, you cannot use compression (-c) in the output file together with the encryption option (-e).

#### 5.1.1. Raw format

- This format is simple and easily exportable to all other emulators/hypervisors.
- It provides the best performance (least I/O overhead).
- It occupies all allocated space on the file system.
- The raw format allows to copy a VM Guest image to a physical device (**dd if=VMGUEST.RAW of=/dev/sda**).

- It is byte-for-byte the same as what the VM Guest sees, so this wastes a lot of space.

### 5.1.2. qcow2 format

- Use this to have smaller images (useful if your file system does not support holes).
- It has optional AES encryption (now deprecated).
- Zlib-based compression option.
- Supports multiple VM snapshots (internal, external).
- Improved performance and stability.
- Supports changing the backing file.
- Supports consistency checks.
- Less performance than raw format.

#### l2-cache-size

qcow2 can provide the same performance for random read/write access as raw format, but it needs a well-sized cache size. By default, cache size is set to 1 MB. This will give good performance up to a disk size of 8 GB. If you need a bigger disk size, you need to adjust the cache size. For a disk size of 64 GB ( $64*1024 = 65536$ ), you need  $65536 / 8192 B = 8$  MB of cache (`-drive format=qcow2, l2-cache-size=8M`).

#### Cluster size

The qcow2 format offers the capability to change the cluster size. The value must be between 512 KB and 2 MB. Smaller cluster sizes can improve the image file size, whereas larger cluster sizes provide better performance.

#### Preallocation

An image with preallocated metadata is initially larger but can improve performance when the image needs to grow.

#### Lazy refcounts

Reference count updates are postponed with the goal of avoiding metadata I/O and improving performance. This is particularly beneficial with `cache=writethrough`. This option does not batch metadata updates, but if in case of host crash, the reference count tables must be rebuilt, this is done automatically at the next open with `qemu-img check -r all` and takes a certain amount of time.

### 5.1.3. qed format

qed is a follow-on qcow (QEMU Copy On Write) format. Because qcow2 provides all the benefits of qed and more, qed is now deprecated.

#### 5.1.4. VMDK format

VMware 3, 4 or 6 image format for exchanging images with that product.

### 5.2. Overlay disk images

The qcow2 and qed formats provide a way to create a base image (also called backing file) and overlay images on top of the base image. A backing file is useful for reverting to a known state and discarding the overlay. If you write to the image, the backing image will be untouched and all changes will be recorded in the overlay image file. The backing file will never be modified unless you use the `commit` monitor command (or `qemu-img commit`).

To create an overlay image:

```
#qemu-img create -o①backing_file=vmguest.raw②,backing_fmt=raw③\  
-f④ qcow2 vmguest.cow⑤
```

- ① Use `-o` ? for an overview of available options.
- ② The backing file name.
- ③ Specify the file format for the backing file.
- ④ Specify the image format for the VM Guest.
- ⑤ Image name of the VM Guest, it will only record the differences from the backing file.

## Backing image path



You should not change the path to the backing image, otherwise you will need to adjust it. The path is stored in the overlay image file. To update the path, make a symbolic link from the original path to the new path and then use the **qemu-img** **rebase** option.

```
#ln -sf /var/lib/images/OLD_PATH/vmguest.raw \
/var/lib/images/NEW_PATH/vmguest.raw
#qemu-img rebase❶ -u❷ -b❸ \
/var/lib/images/OLD_PATH/vmguest.raw /var/lib/images/NEW_PATH/
vmguest.cow
```

- ❶ The **rebase** subcommand tells **qemu-img** to change the backing file image.
- ❷ The **-u** option activates the unsafe mode (see note below). There are two different modes in which rebase can operate:
  - *Safe*: this is the default mode and performs a real rebase operation. The safe mode is a time-consuming operation.
  - *Unsafe*: the unsafe mode (**-u**) only changes the backing files name and the format of the file name, making no checks on the file's contents. Use this mode to rename or move a backing file.
- ❸ The backing image to be used is specified with **-b** and the image path is the last argument of the command.

A common use is to initiate a new guest with the backing file. Let's assume we have a **sle15\_base.img** VM Guest ready to be used (fresh installation without any modification). This will be our backing file. Now you need to test a new package, on an updated system and on a system with a different kernel. We can use **sle15\_base.img** to instantiate the new SUSE Linux Enterprise VM Guest by creating a **qcow2** overlay file pointing to this backing file (**sle15\_base.img**).

In our example, we will use **sle15\_updated.qcow2** for the updated system, and **sle15\_kernel.qcow2** for the system with a different kernel.

To create the two thin provisioned systems, use the **qemu-img** command line with the **-b** option:

```
#qemu-img create -b /var/lib/libvirt/sle15_base.img -f qcow2 \
/var/lib/libvirt/sle15_updated.qcow2
Formatting 'sle15_updated.qcow2', fmt=qcow2 size=17179869184
backing file='sle15_base.img' encryption=off cluster_size=65536
lazy_refcounts=off nocow=off
#qemu-img create -b /var/lib/libvirt/sle15_base.img -f qcow2 \
/var/lib/libvirt/sle15_kernel.qcow2
Formatting 'sle15_kernel.qcow2', fmt=qcow2 size=17179869184
backing file='vmguest-sle15_base.img' encryption=off cluster_size=65536
lazy_refcounts=off nocow=off
```

The images are now usable, and you can do your test without touching the initial `sle15_base.img` backing file. All changes will be stored in the new overlay images. You can also use these new images as a backing file and create a new overlay.

```
#qemu-img create -b sle15_kernel.qcow2 -f qcow2 sle15_kernel_TEST.qcow2
```


When using `qemu-img info` with the option `--backing-chain`, it will return all information about the entire backing chain recursively:

```
#qemu-img info --backing-chain
/var/lib/libvirt/images/sle15_kernel_TEST.qcow2
image: sle15_kernel_TEST.qcow2
file format: qcow2
virtual size: 16G (17179869184 bytes)
disk size: 196K
cluster_size: 65536
backing_file: sle15_kernel.qcow2
Format specific information:
compat: 1.1
lazy refcounts: false

image: sle15_kernel.qcow2
file format: qcow2
virtual size: 16G (17179869184 bytes)
disk size: 196K
cluster_size: 65536
backing_file: SLE15.qcow2
Format specific information:
compat: 1.1
lazy refcounts: false

image: sle15_base.img
file format: qcow2
virtual size: 16G (17179869184 bytes)
disk size: 16G
cluster_size: 65536
Format specific information:
compat: 1.1
lazy refcounts: true
```

**Figure 1. Understanding image overlay**



### 5.3. Opening a VM Guest image

To access the file system of an image, use the `guestfs-tools`. If you do not have this tool installed on your system, you can mount an image with other Linux tools. Avoid accessing an untrusted or unknown VM Guest's image system because this can lead to security issues (for more information, read [D. Berrangé's post](#)).

#### 5.3.1. Opening a raw image

##### Procedure 4. Mounting a raw image

1. To mount the image, find a free loop device. The following command displays the first unused loop device, `/dev/loop1` in this example.

```
#losetup -f
/dev/loop1
```

2. Associate an image (SLE15.raw in this example) with the loop device:

```
#losetup /dev/loop1 SLE15.raw
```

3. Check whether the image has successfully been associated with the loop device by getting detailed information about the loop device:

```
#losetup -l
NAME      SIZELIMIT OFFSET AUTOCLEAR RO BACK-FILE
/dev/loop1      0      0      0 0  /var/lib/libvirt/images/
SLE15.raw
```

4. Check the image's partitions with `kpartx`:

```
#kpartx -a① -v② /dev/loop1
add map loop1p1 (254:1): 0 29358080 linear /dev/loop1 2048
```

- ❶ Add partition device mappings.
- ❷ Verbose mode.

5. Now mount the image partitions (to `/mnt/sle15mount` in the following example):

```
#mkdir /mnt/sle15mount
#mount /dev/mapper/loop1p1 /mnt/sle15mount
```



### Raw image with LVM

If your raw image contains an LVM volume group, you should use LVM tools to mount the partition. Refer to the section called “*Opening images containing LVM*”.

## Procedure 5. Unmounting a raw image

1. Unmount all mounted partitions of the image, for example:

```
#umount /mnt/sle15mount
```

2. Delete partition device mappings with **kpartx**:

```
#kpartx -d /dev/loop1
```

3. Detach the devices with **losetup**

```
#losetup -d /dev/loop1
```

### 5.3.2. Opening a qcow2 image

## Procedure 6. Mounting a qcow2 image

1. First, you need to load the **nbd** (network block devices) module. The following example loads it with support for 16 block devices (`max_part=16`). Check with **dmesg** whether the operation was successful:

```
#modprobe nbd max_part=16
#dmesg | grep nbd
[89155.142425] nbd: registered device at major 43
```

2. Connect the VM Guest image, for example, `SLE15.qcow2`, to an NBD device (`/dev/nbd0` in the following example) with the **qemu-nbd** command. Use a free NBD device:

```
#qemu-nbd -c❶ /dev/nbd0❷ SLE15.qcow2❸
```

- ❶ Connect `SLE15.qcow2` to the local NBD device `/dev/nbd0`
- ❷ NBD device to use
- ❸ VM Guest image to use

## Checking for a free NBD device



To check whether an NBD device is free, run the following command:

```
#lsof /dev/nbd0
COMMAND PID USER FD TYPE DEVICE SIZE/OFF NODE NAME
qemu-nbd 15149 root 10u BLK 43,0 0t0 47347 /dev/nbd0
```

If the command produces an output like in the example above, the device is busy (not free). This can also be confirmed by the presence of the `/sys/devices/virtual/block/nbd0/pid` file.

### 3. Inform the operating system about partition table changes with **partprobe**:

```
#partprobe /dev/nbd0 -s
/dev/nbd0: msdos partitions 1 2
#dmesg | grep nbd0 | tail -1
[89699.082206] nbd0: p1 p2
```

### 4. In the example above, the SLE15.qcow2 contains two partitions: `/dev/nbd0p1` and `/dev/nbd0p2`. Before mounting these partitions, use **vgscan** to check whether they belong to an LVM volume:

```
#vgscan -v
  Wiping cache of LVM-capable devices
  Wiping internal VG cache
  Reading all physical volumes. This may take a while...
  Using volume group(s) on command line.
  No volume groups found.
```

### 5. If no LVM volume has been found, you can mount the partition with **mount**:

```
#mkdir /mnt/nbd0p2
# mount /dev/nbd0p1 /mnt/nbd0p2
```

Refer to the section called “*Opening images containing LVM*” for information on how to handle LVM volumes.

## Procedure 7. Unmounting a qcow2 image

### 1. Unmount all mounted partitions of the image, for example:

```
#umount /mnt/nbd0p2
```

### 2. Disconnect the image from the `/dev/nbd0` device.

```
#qemu-nbd -d /dev/nbd0
```

### 5.3.3. Opening images containing LVM

#### Warning



If your VM Host Server uses the VG name system, and the guest image also uses the VG name system, LVM will complain during its activation. A workaround is to temporarily rename the guest VG, while a correct approach is to use different VG names for the guests than for the VM Host Server.

#### Procedure 8. Mounting images containing LVM

1. To check images for LVM groups, use **vgscan -v**. If an image contains LVM groups, the output of the command looks like the following:

```
#vgscan -v
Wiping cache of LVM-capable devices
Wiping internal VG cache
Reading all physical volumes. This may take a while...
Finding all volume groups
Finding volume group "system"
Found volume group "system" using metadata type lvm2
```

2. The system LVM volume group has been found on the system. You can get more information about this volume with **vgdisplay VOLUMEGROUPNAME** (in our case **VOLUMEGROUPNAME** is system). You should activate this volume group to expose LVM partitions as devices so the system can mount them. Use **vgchange**:

```
#vgchange -ay -v
Finding all volume groups
Finding volume group "system"
Found volume group "system"
activation/volume_list configuration setting not defined: Checking only
host tags for system/home
Creating system-home
Loading system-home table (254:0)
Resuming system-home (254:0)
Found volume group "system"
activation/volume_list configuration setting not defined: Checking only
host tags for system/root
Creating system-root
Loading system-root table (254:1)
Resuming system-root (254:1)
Found volume group "system"
activation/volume_list configuration setting not defined: Checking only
host tags for system/swap
Creating system-swap
Loading system-swap table (254:2)
Resuming system-swap (254:2)
Activated 3 logical volumes in volume group system
 3 logical volume(s) in volume group "system" now active
```

3. All partitions in the volume group will be listed in the **/dev/mapper** directory. You can simply mount them now.

```
#ls /dev/mapper/system-*
/dev/mapper/system-home  /dev/mapper/system-root  /dev/mapper/system-swap

#mkdir /mnt/system-root
#mount  /dev/mapper/system-root /mnt/system-root

#ls /mnt/system-root/
bin  dev  home  lib64  mnt  proc  root  sbin  srv  tmp  var
boot  etc  lib  lost+found  opt  read-write  run  selinux  sys  usr
```

### Procedure 9. Unmounting images containing LVM

1. Unmount all partitions (with **umount**)

```
#umount /mnt/system-root
```

2. Deactivate the LVM volume group (with **vgchange -an VOLUMEGROUPNAME**)

```
#vgchange -an -v system
Using volume group(s) on command line
Finding volume group "system"
Found volume group "system"
Removing system-home (254:0)
Found volume group "system"
Removing system-root (254:1)
Found volume group "system"
Removing system-swap (254:2)
Deactivated 3 logical volumes in volume group system
0 logical volume(s) in volume group "system" now active
```

3. Now you have two choices:

- With a qcow2 image, proceed as described in step 2 (**qemu-nbd -d /dev/nbd0**).
- With a raw image, proceed as described in step 2 (**kpartx -d /dev/loop1; losetup -d /dev/loop1**).

#### Check for a successful unmount



You should double-check that unmounting succeeded by using a system command like **losetup**, **qemu-nbd**, **mount** or **vgscan**. If this is not the case, you may have trouble using the VM Guest because its system image is used in different places.

## 5.4. File system sharing

You can access a host directory in the VM Guest using the **filesystem** element. In the following example we will share the `/data/shared` directory and mount it in the VM Guest. The `accessmode` parameter only works with `type='mount'` for the QEMU/KVM driver (most other values for `type` are exclusively used for the LXC driver).

```
<filesystem type='mount'❶ accessmode='mapped'❷>
  <source dir='/data/shared'❸>
  <target dir='shared'❹/>
</filesystem>
```

- ❶ A host directory to mount VM Guest.
- ❷ Access mode (the security mode) set to `mapped` will give access with the permissions of the hypervisor. Use `passthrough` to access this share with the permissions of the user inside the VM Guest.
- ❸ Path to share with the VM Guest.
- ❹ Name or label of the path for the mount command.

To mount the shared directory on the VM Guest, use the following commands: under the VM Guest, now you need to mount the target `dir='shared'`:

```
#mkdir /opt/mnt_shared
#mount shared -t 9p /opt/mnt_shared -o trans=virtio
```

See [libvirt File System](#) and [QEMU 9psetup](#) for more information.

## 6. VM Guest configuration

### 6.1. Virtio driver

To increase VM Guest performance, we recommend using paravirtualized drivers within the VM Guests. The virtualization standard for such drivers for KVM are the `virtio` drivers, which are designed for running in a virtual environment. Xen uses similar paravirtualized device drivers (like [VMDP](#) in a Windows\* guest).

#### 6.1.1. `virtio blk`

`virtio_blk` is the `virtio` block device for disk. To use the `virtio blk` driver for a block device, specify the `bus='virtio'` attribute in the `disk` definition:

```
<disk type='....' device='disk'>
  ...
  <target dev='vda' bus='virtio' />
</disk>
```

#### Disk device names



`virtio` disk devices are named `/dev/vd[a-z][1-9]`. If you migrate a Linux guest from a non-`virtio` disk, you need to adjust the `root=` parameter in the GRUB configuration, and regenerate the `initrd` file. Otherwise the system cannot boot. On VM Guests with other operating systems, the boot loader may need to be adjusted or re-installed accordingly, too.

## Using **virtio** disks with **qemu-system-ARCH**



When running **qemu-system-ARCH**, use the **-drive** option to add a disk to the VM Guest. See the section called “*Basic installation with **qemu-system-ARCH***” for an example. The **-hd[abcd]** option will not work for virtio disks.

### 6.1.2. **virtio net**

**virtio\_net** is the virtio network device. The kernel modules should be loaded automatically in the guest at boot time. You need to start the service to make the network available.

```
<interface type='network'>
  ...
  <model type='virtio' />
</interface>
```

### 6.1.3. **virtio balloon**

The virtio balloon is used for host memory over-commits for guests. For Linux guests, the balloon driver runs in the guest kernel, whereas for Windows guests, the balloon driver is in the VMDP package. **virtio\_balloon** is a PV driver to give or take memory from a VM Guest.

- *Inflate balloon*: return memory from guest to host kernel (for KVM) or to hypervisor (for Xen)
- *Deflate balloon*: Guest will have more available memory

It is controlled by the **currentMemory** and **memory** options.

```
<memory unit='KiB'>16777216</memory>
<currentMemory unit='KiB'>1048576</currentMemory>
[...]
<devices>
  <memballoon model='virtio' />
</devices>
```

You can also use **virsh** to change it:

```
>virsh setmem DOMAIN_IDMEMORY in KB
```

### 6.1.4. Checking virtio presence

You can check the virtio block PCI with:

```
>find /sys/devices/ -name virtio*
/sys/devices/pci0000:00/0000:00:06.0/virtio0
/sys/devices/pci0000:00/0000:00:07.0/virtio1
/sys/devices/pci0000:00/0000:00:08.0/virtio2
```

To find the block device associated with vdX:

```
>find /sys/devices/ -name virtio* -print -exec ls {}/block 2>/dev/null \
/sys/devices/pci0000:00/0000:00:06.0/virtio0
/sys/devices/pci0000:00/0000:00:07.0/virtio1
/sys/devices/pci0000:00/0000:00:08.0/virtio2
vda
```

To get more information on the virtio block:

```
>udevadm info -p /sys/devices/pci0000:00/0000:00:08.0/virtio2
P: /devices/pci0000:00/0000:00:08.0/virtio2
E: DEVPATH=/devices/pci0000:00/0000:00:08.0/virtio2
E: DRIVER=virtio_blk
E: MODALIAS=virtio:d00000002v00001AF4
E: SUBSYSTEM=virtio
```

To check all virtio drivers being used:

```
>find /sys/devices/ -name virtio* -print -exec ls -l {}/driver 2>/dev/null \
/sys/devices/pci0000:00/0000:00:06.0/virtio0
lrwxrwxrwx 1 root root 0 Jun 17 15:48 /sys/devices/pci0000:00/0000:00:06.0/virtio0/driver -> ../../../../../../bus/virtio/drivers/virtio_console
/sys/devices/pci0000:00/0000:00:07.0/virtio1
lrwxrwxrwx 1 root root 0 Jun 17 15:47 /sys/devices/pci0000:00/0000:00:07.0/virtio1/driver -> ../../../../../../bus/virtio/drivers/virtio_balloon
/sys/devices/pci0000:00/0000:00:08.0/virtio2
lrwxrwxrwx 1 root root 0 Jun 17 14:35 /sys/devices/pci0000:00/0000:00:08.0/virtio2/driver -> ../../../../../../bus/virtio/drivers/virtio_blk
```

### 6.1.5. Find device driver options

Virtio devices and other drivers have multiple options. To list them all, use the `help` parameter of the `qemu-system-ARCH` command.

```
>qemu-system-x86_64 -device virtio-net,help
virtio-net-pci.ioeventfd=on/off
virtio-net-pci.vectors=uint32
virtio-net-pci.indirect_desc=on/off
virtio-net-pci.event_idx=on/off
virtio-net-pci.any_layout=on/off
....
```

## 6.2. Cirrus video driver

To get 16-bit color, high compatibility and better performance, we recommend using the cirrus video driver.



**libvirt**

libvirt ignores the `vram` value because video size has been hardcoded in QEMU.

```
<video>
  <model type='cirrus' vram='9216' heads='1' />
</video>
```

### 6.3. Better entropy

Virtio RNG (random number generator) is a paravirtualized device that is exposed as a hardware RNG device to the guest. On the host side, it can be wired up to one of several sources of entropy (including a real hardware RNG device and the host's `/dev/random`) if hardware support does not exist. The Linux kernel contains the guest driver for the device from version 2.6.26 and higher.

The system entropy is collected from several non-deterministic hardware events and is mainly used by cryptographic applications. The virtual random number generator device (paravirtualized device) allows the host to pass through entropy to VM Guest operating systems. This results in a better entropy in the VM Guest.

To use Virtio RNG, add an RNG device in **virt-manager** or directly in the VM Guest's XML configuration:

```
<devices>
  <rng model='virtio'>
    <backend model='random'>/dev/random</backend>
  </rng>
</devices>
```

The host now should use `/dev/random`:

```
>lsof /dev/random
qemu-syst 4926 qemu      6r      CHR      1,8      0t0  8199 /dev/random
```

On the VM Guest, the source of entropy can be checked with:

```
>cat /sys/devices/virtual/misc/hw_random/rng_available
```

The current device used for entropy can be checked with:

```
>cat /sys/devices/virtual/misc/hw_random/rng_current
virtio_rng.0
```

You should install the `rng-tools` package on the VM Guest, enable the service, and start it. Under SUSE Linux Enterprise Server 15, do the following:

```
#zypper in rng-tools
#systemctl enable rng-tools
#systemctl start rng-tools
```

### 6.4. Disable unused tools and devices

Per host, use one virtualization technology only. For example, do not use KVM and Xen on the same host. Otherwise, you may find yourself with a reduced amount of available resources, increased security risk and a longer software update queue. Even when the amount of resources allocated to each of the technologies is configured carefully, the host may suffer from reduced overall availability and degraded performance.

Minimize the amount of software and services available on hosts. Most default installations of operating systems are not optimized for VM usage. Install what you really need and remove all other components in the VM Guest.

Windows\* Guest:

- Disable the screen saver
- Remove all graphical effects
- Disable indexing of hard disks if not necessary
- Check the list of started services and disable the ones you do not need
- Check and remove all unneeded devices
- Disable system update if not needed, or configure it to avoid any delay while rebooting or shutting down the host
- Check the Firewall rules
- Schedule backups and anti-virus updates appropriately
- Install the [VMDP](#) paravirtualized driver for best performance
- Check the operating system recommendations, such as on the [Microsoft Windows\\* 7 better performance](#) Web page.

Linux Guest:

- Remove or do not start the X Window System if not necessary
- Check the list of started services and disable the ones you do not need
- Check the OS recommendations for kernel parameters that enable better performance
- Only install software that you really need
- Optimize the scheduling of predictable tasks (system updates, hard disk checks, etc.)

## 6.5. Updating the guest machine type

QEMU machine types define details of the architecture that are particularly relevant for migration and session management. As changes or improvements to QEMU are made, new machine types are added. Old machine types are still supported for compatibility reasons, but to use improvements, we recommend to always migrate to the latest machine type when upgrading.

Changing the guest's machine type for a Linux guest is transparent. For Windows\* guests, we recommend creating a snapshot or backup of the guest—in case Windows\* has issues with the changes it detects, and subsequently the user reverts to the original machine type the guest was created with.

## Changing the machine type



Refer to the section called “*Changing the machine type*” for documentation.

## 7. VM Guest-specific configurations and settings

### Note



This section applies to QEMU / KVM hypervisor only.

### 7.1. ACPI testing

The ability to change a VM Guest's state heavily depends on the operating system. It is important to test this feature before any use of your VM Guests in production. For example, most Linux operating systems disable this capability by default, so this requires you to enable this operation (normally through Polkit).

ACPI must be enabled in the guest for a graceful shutdown to work. To check if ACPI is enabled, run:

```
>virsh dumpxml VMNAME | grep acpi
```

If nothing is printed, ACPI is not enabled for your machine. Use **virsh edit** to add the following XML under <domain>:

```
<features>
  <acpi/>
</features>
```

If ACPI was enabled during a Windows Server\* guest installation, it is not sufficient to turn it on in the VM Guest configuration only. For more information, see <https://support.microsoft.com/en-us/kb/309283>.

Regardless of the VM Guest's configuration, a graceful shutdown is always possible from within the guest operating system.

### 7.2. Keyboard layout

Though it is possible to specify the keyboard layout from a **qemu-system-ARCH** command, we recommend configuring it in the libvirt XML file. To change the keyboard layout while connecting to a remote VM Guest using vnc, edit the VM Guest XML configuration file. For example, to add an en-us keymap, add in the <devices> section:

```
<graphics type='vnc' port=''-1' autoport='yes' keymap='en-us' />
```

Check the vncdisplay configuration and connect to your VM Guest:

```
>virsh vncdisplay sles15 127.0.0.1:0
```

### 7.3. Spice default listen URL

If no network interface other than `lo` is assigned an IPv4 address on the host, the default address on which the spice server listens will not work. An error like the following one will occur:

```
>virsh start sles15
error: Failed to start domain sles15
error: internal error: process exited while connecting to monitor: ((null):26929): Spice-Warning **: reds.c:2330:reds_init_socket: getaddrinfo(127.0.0.1,5900): Address family for hostname not supported
2015-08-12T11:21:14.221634Z qemu-system-x86_64: failed to initialize spice server
```

To fix this, you can change the default `spice_listen` value in `/etc/libvirt/qemu.conf` using the local IPv6 address `::1`. The spice server listening address can also be changed on a per VM Guest basis, use `virsh edit` to add the `listen` XML attribute to the `graphics type='spice'` element:

```
<graphics type='spice' listen='::1' autoport='yes' />>
```

### 7.4. XML to QEMU command line

Sometimes it could be useful to get the QEMU command line to launch the VM Guest from the XML file.

```
>virsh domxml-to-native❶ qemu-argv❷ SLE15.xml❸
```

- ❶ Convert the XML file in domain XML format to the native guest configuration
- ❷ For the QEMU/KVM hypervisor, the format argument needs to be `qemu-argv`
- ❸ Domain XML file to use

```
>sudo virsh domxml-to-native qemu-argv /etc/libvirt/qemu/SLE15.xml
LC_ALL=C PATH=/usr/local/sbin:/usr/local/bin:/usr/sbin:/usr/bin:/sbin:/bin \
  QEMU_AUDIO_DRV=none /usr/bin/qemu-system-x86_64 -name SLE15 -machine \
  pc-i440fx-2.3,accel=kvm,usb=off -cpu SandyBridge -m 4048 -realtime \
  mlock=off -smp 4,sockets=4,cores=1,threads=1 -uuid
8616d00f-5f05-4244-97cc-86aeaed8aea7 \
  -no-user-config -nodefaults -chardev socket,id=charmonitor,path=/var/lib/lib-
virt/qemu/SLE15.monitor,server,nowait \
  -mon chardev=charmonitor,id=monitor,mode=control -rtc base=utc,driftfix=slew
\
  -global kvm-pit.lost_tick_policy=discard -no-hpet \
  -no-shutdown -global PIIIX4_PM.disable_s3=1 -global PIIIX4_PM.disable_s4=1 \
  -boot strict=on -device ich9-usb-ehci1,id=usb,bus=pci.0,addr=0x4.0x7 \
  -device ich9-usb-uhci1,masterbus=usb.0,firstport=0,bus=pci.
0,multifunction=on,addr=0x4 \
  -device ich9-usb-uhci2,masterbus=usb.0,firstport=2,bus=pci.0,addr=0x4.0x1 \
  -device ich9-usb-uhci3,masterbus=usb.0,firstport=4,bus=pci.0,addr=0x4.0x2 \
  -drive file=/var/lib/libvirt/images/SLE15.qcow2,if=none,id=drive-virtio-
disk0,format=qcow2,cache=none \
  -device virtio-blk-pci,scsi=off,bus=pci.0,addr=0x6,drive=drive-virtio-
disk0,id=virtio-disk0,bootindex=2 \
  -drive if=none,id=drive-ide0-0-1,readonly=on,format=raw \
  -device ide-cd,bus=ide.0,unit=1,drive=drive-ide0-0-1,id=ide0-0-1 -netdev
tap,id=hostnet0 \
  -device virtio-net-pci,netdev=hostnet0,id=net0,mac=52:54:00:28:04:a9,bus=pci.
0,addr=0x3,bootindex=1 \
  -chardev pty,id=charserial0 -device isa-serial,chardev=charserial0,id=serial0
\
  -vnc 127.0.0.1:0 -device cirrus-vga,id=video0,bus=pci.0,addr=0x2 \
  -device virtio-balloon-pci,id=balloon0,bus=pci.0,addr=0x5 -msg timestamp=on
```

## 7.5. Change kernel parameters at boot time

### 7.5.1. SUSE Linux Enterprise 11

To change the value for SLE 11 products at boot time, you need to modify your /boot/grub/menu.lst file by adding the **OPTION=**parameter. Then reboot your system.

### 7.5.2. SUSE Linux Enterprise 12 and 15

To change the value for SLE 12 and 15 products at boot time, you need to modify your /etc/default/grub file. Find the variable starting with **GRUB\_CMDLINE\_LINUX\_DEFAULT** and add at the end **OPTION=**parameter (or change it with the correct value if it is already available).

Now you need to regenerate your grub2 configuration:

```
# grub2-mkconfig -o /boot/grub2/grub.cfg
```

Then reboot your system.

## 7.6. Add a device to an XML configuration

To create a new VM Guest based on an XML file, you can specify the QEMU command line using the special tag **qemu:commandline**. For example, to add a virtio-balloon-pci, add this block at the end of the XML configuration file (before the **</domain>** tag):

```
<qemu:commandline>
  <qemu:arg value='-device' />
  <qemu:arg value='virtio-balloon-pci,id=balloon0' />
</qemu:commandline>
```

## 7.7. Adding and removing CPUs

Certain virtualization environments allow adding or removing CPUs while the virtual machine is running.

For the safe removal of CPUs, deactivate them first by executing

```
#echo 0 > /sys/devices/system/cpu/cpuX/online
```

Replace X with the CPU number. To bring a CPU back online, execute

```
#echo 1 > /sys/devices/system/cpu/cpuX/online
```

## 7.8. SCSI persistent reservation on a multipathed device

SCSI persistent reservations allow restriction of access to block devices in a shared storage setup. This avoids improper multiple parallel accesses to the same block device from software components on local or remote hosts, which could lead to device damage and data corruption.

Find more information on managing storage multipath I/O in *Chapter 18, Managing multipath I/O for devices*. Find more information about SCSI persistent reservations in *the section called “SCSI persistent reservations and **mpathpersist**”*.

For the virtualization scenario, QEMU's SCSI passthrough devices `scsi-block` and `scsi-generic` support passing guest persistent reservation requests to a privileged external helper program `qemu-pr-helper`. This needs to start before QEMU and creates a listener socket that accepts incoming connections for communication with QEMU.

### Live migration scenario with multipathed devices



We recommend using the multipath alias instead of `wwid`. It is useful in the VM Guest live migration scenario, because it makes sure that the storage paths are identical between the source and destination hosts.

Find more information about multipath in *the section called “Setting aliases for multipath maps”*.

## Procedure 10. Adding a SCSI persistent reservation in a VM Guest against the related multipathed device in the VM Host Server:

1. In the VM Host Server, create a multipath environment. For more information, refer to *the section called “Configuring the system for multipathing” and the section called “Multipath configuration”*.
2. In the VM Host Server, configure the `<reservations/>` sub-element of the `<source/>` element of the `<disk/>` element for the passed-through lun in your libvirt domain configuration. Refer to [libvirt Domain XML format](#).
3. In the VM Guest, install the `sg3_utils` package and reserve the SCSI disks on demand by using the `sg_persist` command.

### Example 1. Practical example

1. In the VM Host Server, verify that the `multipathd.service` is running, and that a multipathed disk exists and is named, for example, `storage1`.

```
>sudo systemctl status multipathd.service
● multipathd.service - Device-Mapper Multipath Device Controller
  Loaded: loaded (/usr/lib/systemd/system/multipathd.service; enabled; preset: disabled)
  Active: active (running) since Sat 2023-08-26 21:34:13 CST; 1 week
1 day ago
  TriggeredBy: ○ multipathd.socket
    Main PID: 79411 (multipathd)
    Status: "up"
    Tasks: 7
      CPU: 1min 43.514s
    CGroup: /system.slice/multipathd.service
            └─79411 /sbin/multipathd -d -s
```

```
>sudo multipath -ll
storage1 (36589cf000000537c47ad3eb2b20216e) dm-6 TrueNAS,iSCSI Disk
size=50G features='0' hwhandler='1' alua' wp=rw
  |--- policy='service-time 0' prio=50 status=active
  |   `--- 16:0:0:0 sdg 8:96 active ready running
  |--- policy='service-time 0' prio=50 status=enabled
  |   `--- 17:0:0:0 sdh 8:112 active ready running
```

2. In the VM Host Server, add a `<disk/>` element in the VM Guest configuration file by running `virsh edit`.

```
<disk type='block' device='lun' ①>
  <driver name='qemu' type='raw' />
  <source dev='/dev/mapper/storage1'>
    <reservations ② managed='yes' ③/>
  </source>
  <target dev='sda' bus='scsi' />
  <address type='drive' controller='0' bus='0' target='0' unit='0' ④/>
</disk>
```

- ① To support persistent reservations, the disks must be marked as `lun` with type `block` so that QEMU does SCSI passthrough.
- ② If present, it enables persistent reservations for SCSI-based disks. The element has one mandatory attribute `managed` with accepted values `yes` and `no`.

③ If managed is yes, libvirt prepares and manages any resources needed.

When the value of the attribute managed is no, then the hypervisor acts as a client and the path to the server socket must be provided in the child element source, which currently accepts only the following attributes:

**type**

The only valid option is unix.

**path**

The path to the server socket.

**mode**

The role of the hypervisor. Valid is client.

④ Verify that the virtio-scsi HBA that the disk attaches already exists and has available units (the maximum count of units per virtio-scsi HBA is 7). Otherwise, you need to manually add a virtio-scsi HBA to avoid automatically adding the LSI HBA by libvirt. For example:

```
<controller type='scsi' index='0' model='virtio-scsi'>
  <address type='pci' domain='0x0000' bus='0x03' slot='0x00' function='0x0' />
</controller>
```

```
>sudo virsh domblklist sles15sp5
  Target  Source
  -----
  vda      /mnt/images/sles15sp5/disk0.qcow2
  sda      /dev/mapper/storage1
```

3. In the VM Host Server, start the VM Guest. libvirt launches a qemu-pr-helper instance as the server role for the VM Guest sles15sp5, then launches the VM Guest sles15sp5 as the client role.

```
>sudo virsh start sles15sp5
  Domain 'sles15sp5' started
```

```
>sudo virsh list
  Id  Name      State
  -----
  4   sles15sp5  running
```

```
>sudo ps -eo pid,args | grep -v grep | grep qemu-pr-helper
  37063 /usr/bin/qemu-pr-helper -k /var/lib/libvirt/qemu/domain-4-
sles15sp5/pr-helper0.sock
```

```
>sudo virsh dumpxml sles15sp5 | grep -A11 "<disk type='block' de-  
vice='lun'>  
  <disk type='block' device='lun'>  
    <driver name='qemu' type='raw'/>  
    <source dev='/dev/mapper/storage1' index='1'>  
      <reservations managed='yes'>  
        <source type='unix' path='/var/lib/libvirt/qemu/domain-4-sles15sp5/  
pr-helper0.sock' mode='client'/>  
      </reservations>  
    </source>  
    <backingStore/>  
    <target dev='sda' bus='scsi' />  
    <alias name='scsi0-0-0-0' />  
    <address type='drive' controller='0' bus='0' target='0' unit='0' />  
  </disk>
```

4. In the VM Guest, reserve the scsi disk, for example, `sda`, with the key `123abc`.

```
>lsblk  
NAME  MAJ:MIN RM  SIZE  RO  TYPE  MOUNTPOINTS  
sda    8:0    0   50G  0  disk  
vda   253:0   0   20G  0  disk  
|---vda1 253:1   0   8M  0  part  
|---vda2 253:2   0   2G  0  part  [SWAP]  
|---vda3 253:3   0   18G 0  part  /
```

```
>sudo sg_persist --verbose --out --register --param-sark=123abc /dev/sda  
  inquiry cdb: [12 00 00 00 24 00]  
  TrueNAS iSCSI Disk 0123  
  Peripheral device type: disk  
  Persistent reservation out cdb: [5f 00 00 00 00 00 00 00 18 00]  
PR out: command (Register) successful
```

```
>sudo sg_persist --verbose --in -k /dev/sda  
  inquiry cdb: [12 00 00 00 24 00]  
  TrueNAS iSCSI Disk 0123  
  Peripheral device type: disk  
  Persistent reservation in cdb: [5e 00 00 00 00 00 00 20 00 00]  
PR generation=0x5, 2 registered reservation keys follow:  
  0x123abc  
  0x123abc
```

5. In the VM Guest, release the `sda` disk with the key `123abc`.

```
>sudo sg_persist --verbose --out --clear --param-rk=123abc /dev/sda  
  inquiry cdb: [12 00 00 00 24 00]  
  TrueNAS iSCSI Disk 0123  
  Peripheral device type: disk  
  Persistent reservation out cdb: [5f 03 00 00 00 00 00 00 18 00]  
PR out: command (Clear) successful
```

```
>sudo sg_persist --verbose --in -k /dev/sda  
  inquiry cdb: [12 00 00 00 24 00]  
  TrueNAS iSCSI Disk 0123  
  Peripheral device type: disk  
  Persistent reservation in cdb: [5e 00 00 00 00 00 00 20 00 00]  
PR generation=0x6, there are NO registered reservation keys
```

## 8. More information

- [Increasing memory density using KSM](#)
- [linux-kvm.org KSM](#)
- [KSM's kernel documentation](#)

- [ksm - dynamic page sharing driver for Linux v4](#)
- [Memory Ballooning](#)
- [libvirt virtio](#)
- [BFQ \(Budget Fair Queueing\)](#)
- [Documentation for sysctl](#)
- [LWN Random Number](#)
- [Kernel Parameters](#)
- [Huge pages Administration \(Mel Gorman\)](#)
- [kernel hugetlbpge](#)

## 9. Legal notice

Copyright© 2006– 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.

## 10. 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:

1. 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.
2. 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.

3. State on the Title page the name of the publisher of the Modified Version, as the publisher.
4. Preserve all the copyright notices of the Document.
5. Add an appropriate copyright notice for your modifications adjacent to the other copyright notices.
6. 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.
7. Preserve in that license notice the full lists of Invariant Sections and required Cover Texts given in the Document's license notice.
8. Include an unaltered copy of this License.
9. 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.
10. 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.
11. 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.
12. 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.
13. Delete any section Entitled "Endorsements". Such a section may not be included in the Modified Version.
14. Do not retitle any existing section to be Entitled "Endorsements" or to conflict in title with any Invariant Section.
15. 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.



# Using KubeVirt on SUSE Linux Enterprise

Jim Fehlig

Vasily Ulyanov

Koordinierungsstelle  
für IT-Standards

KubeVirt is a virtual machine management add-on for Kubernetes. KubeVirt extends Kubernetes by adding additional virtualization resource types through Kubernetes' Custom Resource Definitions (CRD) API. Along with the Custom Resources, KubeVirt includes controllers and agents that provide virtual machine management capabilities on the cluster. By using this mechanism, the Kubernetes API can be used to manage virtual machine resources similar to other Kubernetes resources.

## 1. KubeVirt components

KubeVirt consists of two RPM-based packages and six container images that provide the Kubernetes virtual machine management extension. The RPM packages include `kubevirt-virtctl` and `kubevirt-manifests`. The container images include `virt-api`, `virt-controller`, `virt-handler`, `virt-launcher`, and `virt-operator`, `libguestfs-tools`.

`kubevirt-virtctl` can be installed on any machine with administrator access to the cluster. It contains the `virtctl` tool, which provides syntactic sugar on top of the `kubectl` tool for virtual machine resources. Although the `kubectl` tool can be used to manage virtual machines, it is a bit awkward since, unlike standard Kubernetes resources, virtual machines maintain state. Migration is also unique to virtual machines. If a standard Kubernetes resource needs to be evacuated from a cluster node, it is destroyed and started again on an alternate node. Since virtual machines are stateful, they cannot be destroyed and must be live-migrated away if a node is under evacuation. The `virtctl` tool abstracts the complexity of managing virtual machines with `kubectl`. It can be used to stop, start, pause, unpause and migrate virtual machines. `virtctl` also provides access to the virtual machine's serial console and graphics server.

`kubevirt-manifests` contains the manifests, or recipes, for installing KubeVirt. The most interesting files are `/usr/share/kube-virt/manifests/release/kubevirt-cr.yaml` and `/usr/share/kube-virt/manifests/release/kubevirt-operator.yaml`. `kubevirt-cr.yaml` contains the KubeVirt Custom Resource definition that represents the KubeVirt service. `kubevirt-operator.yaml` is the recipe for deploying the KubeVirt operator, which deploys the KubeVirt service to the cluster and manages its' lifecycle.

`virt-api` is a cluster component that provides the Kubernetes API extension for virtual machine resources. Like `virt-api`, `virt-controller` is a cluster component that watches for new ob-

jects created via `virt-api`, or updates to existing objects, and takes action to ensure the object state matches the requested state. `virt-handler` is a DaemonSet and a node component that has the job of keeping the cluster-level virtual machine object in sync with the `libvirt` domain running in `virt-launcher`. `virt-handler` can also perform node-centric operations like configuring networking and/or storage on the node per the virtual machine specification. `virt-launcher` is also a node component and has the job of running `libvirt` plus `qemu` to provide the virtual machine environment. `virt-launcher` is a lowly pod resource. `libguestfs-tools` is a component providing a set of utilities for accessing and modifying VM disk images.

`virt-operator` implements the Kubernetes operator pattern. Operators encode the human knowledge required to deploy, run and maintain an application. Operators are a Kubernetes Deployment resource type and are often used to manage the custom resources and custom controllers that together provide a more complex Kubernetes application such as KubeVirt.

## 2. Installing KubeVirt on Kubernetes

KubeVirt can be installed on a Kubernetes cluster by installing the `kubevirt-manifests` package on an admin node, applying the `virt-operator` manifest, and creating the KubeVirt custom resource. For example, on a cluster admin node execute the following:

```
>sudo zypper install kubevirt-manifests
>kubectl apply -f /usr/share/kube-virt/manifests/release/kubevirt-operator.yaml
>kubectl apply -f /usr/share/kube-virt/manifests/release/kubevirt-cr.yaml
```

After creating the KubeVirt custom resource, `virt-operator` deploys the remaining KubeVirt components. Progress can be monitored by viewing the status of the resources in the `kubevirt` namespace:

```
>kubectl get all -n kubevirt
```

The cluster is ready to deploy virtual machines once `virt-api`, `virt-controller`, and `virt-handler` are READY with STATUS “Running”.

Alternatively it is possible to wait until KubeVirt custom resource becomes available:

```
>kubectl -n kubevirt wait kv kubevirt --for condition=Available
```

Some KubeVirt functionality is disabled by default and must be enabled via feature gates. For example, live migration and the use of `HostDisk` for virtual machine disk images are disabled. Enabling KubeVirt feature gates can be done by altering an existing KubeVirt custom resource and specifying the list of features to enable. For example, you can enable live migration and the use of `HostDisks`:

```
>kubectl edit kubevirt kubevirt -n kubevirt
...
spec:
  configuration:
    developerConfiguration:
      featureGates:
        - HostDisk
        - LiveMigration
```

#### Note



The names of feature gates are case-sensitive.

### 3. Updating the KubeVirt deployment

Updating KubeVirt is similar to the initial installation. The updated operator manifest from the `kubevirt-manifests` package is applied to the cluster.

```
>sudo zypper update kubevirt-manifests
>kubectl apply -f /usr/share/kube-virt/manifests/release/kubevirt-operator.yaml
```

### 4. Deleting KubeVirt from a cluster

KubeVirt can be deleted from a cluster by deleting the custom resource and operator:

```
>kubectl delete -n kubevirt kubevirt kubevirt # or alternatively: kubectl delete
-f /usr/share/kube-virt/manifests/release/kubevirt-cr.yaml
>kubectl delete -f /usr/share/kube-virt/manifests/release/kubevirt-operator.yaml
```

#### Note



It is important to delete the custom resource first otherwise it gets stuck in the Terminating state. To fix that the resource finalizer needs to be manually deleted:

```
>kubectl -n kubevirt patch kv kubevirt --type=json -p '[{"op": "re-
move", "path": "/metadata/finalizers"}]'
```

After deleting the resources from Kubernetes cluster the installed KubeVirt RPMs can be removed from the system:

```
>sudo zypper rm kubevirt-manifests kubevirt-virtctl
```

### 5. Containerized Data Importer

Containerized Data Importer (CDI) is an add-on for Kubernetes focused on persistent storage management. It is primarily used for building and importing Virtual Machine Disks for KubeVirt.

#### 5.1. Installing CDI

CDI can be installed on a Kubernetes cluster in a way similar to KubeVirt by installing the RPMs and applying the operator and custom resource manifests using `kubectl`:

```
>sudo zypper in containerized-data-importer-manifests
>kubectl apply -f /usr/share/cdi/manifests/release/cdi-operator.yaml
>kubectl apply -f /usr/share/cdi/manifests/release/cdi-cr.yaml
```

## 5.2. Updating and deleting CDI:

To update CDI:

```
>sudo zypper update containerized-data-importer-manifests
>kubectl apply -f /usr/share/cdi/manifests/release/cdi-operator.yaml
```

To delete CDI:

```
>kubectl delete -f /usr/share/cdi/manifests/release/cdi-cr.yaml
>kubectl delete -f /usr/share/cdi/manifests/release/cdi-operator.yaml
>sudo zypper rm containerized-data-importer-manifests
```

## 6. Running virtual machines

Two of the most interesting custom resources provided by KubeVirt are *VirtualMachine* (VM) and *VirtualMachineInstance* (VMI). As the names imply, a VMI is a running instance of a VM. The lifecycle of a VMI can be managed independently from a VM, but long-lived, stateful virtual machines are managed as a VM. The VM is deployed to the cluster in a shutoff state, then activated by changing the desired state to running. Changing a VM resource state can be done with the standard Kubernetes client tool `kubectl` or with the client `virtctl` provided by KubeVirt.

The VM and VMI custom resources make up part of the KubeVirt API. To create a virtual machine, a VM or VMI manifest must be created that adheres to the API. The API supports setting a wide variety of the common virtual machine attributes, for example, model of vCPU, number of vCPUs, amount of memory, disks, network ports, etc. Below is a simple example of a VMI manifest for a virtual machine with one Nehalem CPU, 2G of memory, one disk, and one network interface:

```

apiVersion: kubevirt.io/v1
kind: VirtualMachineInstance
metadata:
  labels:
    special: vmi-host-disk
  name: sles15sp2
spec:
  domain:
    cpu:
      model: Nehalem-IBRS
  devices:
    disks:
      - disk:
          bus: virtio
          name: host-disk
  interfaces:
    - name: green
      masquerade: {}
  ports:
    - port: 80
  machine:
    type: ""
  resources:
    requests:
      memory: 2048M
  terminationGracePeriodSeconds: 0
  networks:
    - name: green
      pod: {}
  volumes:
    - hostDisk:
        path: /hostDisks/sles15sp2/disk.raw
        type: Disk
        shared: true
        name: host-disk

```

Applying this VMI manifest to the cluster creates a virt-launcher container running `libvirt` and `qemu`, providing the familiar KVM virtual machine environment.

```

>kubectl apply -f sles15sp2vmi.yaml
>kubectl get vmis

```

Similar to other Kubernetes resources, VMs and VMIs can be managed with the `kubectl` client tool. Any `kubectl` operation that works with resource types works with the KubeVirt custom resources, for example, `describe`, `delete`, `get`, `log`, `patch`, etc. VM resources are a bit more awkward to manage with `kubectl`. Since a VM resource can be in a shutoff state, turning it on requires patching the manifest to change the desired state to running. Find an example below:

```

>kubectl patch vm sles15sp2 --type merge -p '{"spec":{"running":true}}'

```

The `virtctl` tool included in the `kubevirt-virtctl` package provides syntactic sugar on top of `kubectl` for VM and VMI resources, allowing them to be stopped, started, paused, unpause and migrated. `virtctl` also provides access to the virtual machine's serial console and graphics server. Find an example below:

```
>virtctl start VM
>virtctl console VMI
>virtctl stop VM
>virtctl pause VM|VMI
>virtctl unpause VM|VMI
>virtctl vnc VMI
>virtctl migrate VM
```

## 7. Live migration

KubeVirt supports live migration of VMs. Though this functionality must first be activated by adding `LiveMigration` to the list of feature gates in the KubeVirt custom resource.

```
>kubectl edit kubevirt kubevirt -n kubevirt
```

```
spec:
  configuration:
    developerConfiguration:
      featureGates:
        - LiveMigration
```

### 7.1. Prerequisites

- All the Persistent Volume Claims (PVCs) used by a VM must have `ReadWriteMany` (RWX) access mode.
- VM pod network binding must be of type `masquerade`:

```
spec:
  domain:
    devices:
      interfaces:
        - name: green
          masquerade: {}
```

Whether live migration is possible or not can be checked via the `VMI.status.conditions` field of a running VM spec:

```
>kubectl describe vmi sles15sp2
```

```
Status:
  Conditions:
    Status: True
    Type: LiveMigratable
  Migration Method: BlockMigration
```

### 7.2. Initiating live migration

Live migration of a VMI can be initiated by applying the following yaml file:

```
apiVersion: kubevirt.io/v1
kind: VirtualMachineInstanceMigration
metadata:
  name: migration-job
spec:
  vmiName: sles15sp2
```

```
>kubectl apply -f migration-job.yaml
```

Alternatively it is possible to migrate a VM using `virtctl` tool:

```
>virtctl migrate VM
```

### 7.3. Cancelling live migration

Live migration can be canceled by deleting the existing migration object:

```
>kubectl delete VirtualMachineInstanceMigration migration-job
```

## 8. Volume hotplugging

KubeVirt allows hotplugging additional storage into a running VM. Both block and file system volume types are supported. The hotplug volumes feature can be activated via the `HotplugVolumes` feature gate:

```
>kubectl edit kubevirt kubevirt -n kubevirt
```

```
spec:
  configuration:
    developerConfiguration:
      featureGates:
        - HotplugVolumes
```

Assuming that `hp-volume` is an existing `DataVolume` or `PVC`, `virtctl` can be used to operate with the volume on a running VM:

```
>virtctl addvolume sles15sp2 --volume-name=hp-volume
>virtctl removevolume sles15sp2 --volume-name=hp-volume
```

## 9. Running Windows VMs with VMDP ISO

The VMDP ISO is provided in the form of a container image which can be consumed by KubeVirt. To run a Windows VM with VMDP ISO attached, the corresponding `containerDisk` needs to be added to the VM definition:

```
spec:
  domain:
    devices:
      disks:
        - name: vmdp
          cdrom:
            bus: sata
  volumes:
    - containerDisk:
        image: registry.suse.com/suse/vmdp/vmdp:latest
        name: vmdp
```

**Note**

The sequence in which the disks are defined affects the boot order. It is possible to specify the `bootOrder` explicitly or otherwise sort the disk items as needed.

## 10. Supported features

- Guest Agent Information
- Live migration
- Hotplug volumes
- VMI Dedicated CPU resource

### 10.1. VMI virtual hardware

- machine type
- BIOS/UEFI/SMBIOS
- cpu
- clock
- RNG
- CPU/Memory limits and requirements
- tablet input
- hugepage

### 10.2. VMI disks and volumes

Disk types:

- lun
- disk
- cdrom

Volume sources:

- cloudInitNoCloud
- cloudInitConfigDrive
- persistentVolumeClaim
- dataVolume
- ephemeral
- containerDisk
- emptyDisk

- hostDisk
- configMap
- secret
- serviceAccount
- downwardMetrics

High performance features:

- IO threads
- Virtio Block Multi-Queue
- Disk cache

### 10.3. VMI interfaces and networks

Network (back-end) types:

- pod
- multus

Interface (front-end) types:

- bridge
- masquerade

## 11. Debugging

If issues are encountered the following debug resources are available to help identify the problem.

The status of all KubeVirt resources can be examined with the **kubectl get** command:

```
>kubectl get all -n kubevirt
```

Resources with failed status can be further queried by examining their definition and expanded status information.

```
>kubectl describe deployment virt-operator
>kubectl get deployment virt-operator -o yaml -n kubevirt
>kubectl describe pod virt-handler-xbjkg -n kubevirt
>kubectl get pod virt-handler-xbjkg -o yaml -n kubevirt
```

Logs from the problematic KubeVirt pod can contain a wealth of information since stderr and service logging from within the pod is generally available via the Kubernetes log service:

```
>kubectl logs virt-operator-558c57bc4-mg68w -n kubevirt
>kubectl logs virt-handler-xbjkg -n kubevirt
```

If the underlying pod is running but there are problems with the service running in it, a shell can be accessed to inspect the pod environment and poke at its service:

```
>kubectl -n kubevirt exec -it virt-handler-xbjkg -- /bin/bash
```



# AMD Secure Encrypted Virtualization (AMD-SEV) Guide

Koordinierungsstelle  
für IT-Standards

AMD's Secure Encrypted Virtualization (SEV) allows the memory of virtual machines to be encrypted. SEV with Encrypted State (SEV-ES) goes one step further by encrypting the virtual machine's CPU register content. These technologies increase system security and are ideal for multi-tenant environments such as cloud computing. They enable protection from a variety of cross-VM and hypervisor-based attacks. As an example, a hostile VM that has escaped its hypervisor-enforced confines and is able to read arbitrary memory is unable to steal sensitive data from an SEV or SEV-ES VM.

This document aims to provide a basic understanding of how SEV and SEV-ES work, and how to enable and configure these features. It also mentions certain limitations and restrictions that the use of SEV and SEV-ES causes as compared to non-encrypted virtualization.

## 1. Introducing SEV

Encryption of computer data stored on disk is a widely deployed feature. However, data in RAM is stored in the clear. This can leave that data vulnerable to software or hardware probing by intruders on the host system, particularly in cloud computing environments where the physical resources are shared by many tenants. Consider a virtual machine of a hostile tenant escaping its sandbox because of a hypervisor bug and searching memory for sensitive data.

AMD's SEV (Secure Encrypted Virtualization) is a technology to protect Linux KVM virtual machines by transparently encrypting the memory of each VM with a unique key. SEV can also calculate a signature of the memory contents, which can be sent to the VM's owner as an attestation that the memory was encrypted correctly by the firmware. SEV is especially relevant to cloud computing environments, where VMs are hosted on remote servers which are not under the control of the VMs' owners. SEV can reduce the amount of trust VMs need to place in the hypervisor and administrator of their host system.

When a virtual machine is processing sensitive data, it can be present in CPU registers as well as memory. If the processing is halted, for example, to service an interrupt or share time with other processes, the virtual machine's CPU register contents are saved to hypervisor memory. This memory is readable by the hypervisor even if SEV is enabled. SEV-ES protects against this scenario by encrypting all CPU register contents when the processing of a virtual machine is halted.

SEV-ES builds upon SEV to provide an even smaller attack surface for virtual machines running in a multi-tenant environment.

## 2. VM host requirements

The VM host hardware must support AMD's SEV technology. To detect if the host hardware supports SEV, check that the `sev` attribute is in the capabilities of `libvirt` and that its value is set appropriately:

```
<domainCapabilities>
  ...
  <features>
  ...
  <sev supported='yes' />
  ...
  </sev>
  </features>
</domainCapabilities>
```

Additionally, ensure that the `kvm_amd` kernel module has the `sev` parameter enabled:

```
/sys/module/kvm_amd/parameters/sev = 1
```

## 3. VM requirements

The VM must be the modern Q35 machine type and must use UEFI firmware. `libvirt` can automatically select an appropriate SEV or SEV-ES enabled UEFI firmware, or one can be specified manually. Currently, the only firmware supported are `/usr/share/qemu/ovmf-x86_64-sev.bin` and `/usr/share/qemu/ovmf-x86_64-sev-code.bin`. See the section called “*Advanced UEFI configuration*” for more details on using UEFI firmware and the auto-selection feature.



### No IDE support in Q35

The Q35 machine type does not have an IDE controller and does not support IDE disks.

All `virtio-net` devices need to be configured with the iPXE option ROM disabled. iPXE is currently not compatible with SEV and SEV-ES. All memory regions used by the VM must be locked for Direct Memory Access (DMA) and to prevent swapping. This includes memory for the VM and any memory regions allocated by QEMU to support running the VM, such as UEFI pflash for firmware and variable store, video RAM, etc.

## 4. VM configuration

### Example 1. Sample configuration file

As an example, an SEV-encrypted VM configured with 4 GB of memory would contain the following XML configuration:

```

<domain type='kvm'>
    <memory unit='KiB'>4194304</memory>
    <currentMemory unit='KiB'>4194304</currentMemory>
    <memoryBacking>
        <locked/> ①
    </memoryBacking>
    <os>
        <type arch='x86_64' machine='pc-q35-2.11'>hvm</type>
        <loader readonly='yes' stateless='yes' type='pflash'>/usr/share/qemu/ovmf-x86_64-sev.bin</loader>
        <boot dev='hd' />
    </os>
    <launchSecurity ② type='sev'>
        <cbitpos>47</cbitpos> ③
        <reducedPhysBits>1</reducedPhysBits> ④
        <policy>0x0033</policy> ⑤
        <dhCert>AAAABBBB=CCCCCDDDDD</dhCert> ⑥
        <session>AAAABBBB=EEEEEEFFFF</session> ⑦
    </launchSecurity>
    <devices>
        <interface type='bridge'>
            <source bridge='br0' />
            <model type='virtio' />
            <rom enabled='no' /> ⑧
        </interface>
        ...
    </devices>
    ...
</domain>

```

- ① The `memoryBacking` element, along with its child element `locking`, is used to ease memory limit restrictions libvirt places on the VM's cgroup. Otherwise, VM creation would fail when QEMU attempts to lock the VM's memory regions along with other regions used to support the VM operation. See [https://libvirt.org/kbase/launch\\_security\\_sev.html#memory](https://libvirt.org/kbase/launch_security_sev.html#memory) for more information on VM memory configuration requirements for SEV VMs.
- ② The `launchSecurity type='sev'` element and its contents enable encryption of the VM's memory contents.
- ③ When memory encryption is enabled, one of the physical address bits (also known as the "C-bit") is used to mark if a memory page is protected. The required `cbitpos` element provides the location of the C-bit in a guest page table entry. For example, the value 47 indicates that bit position 47 in a page table entry determines whether that page is encrypted or not. The C-bit number is read from the host's CPUID and is thus hardware-dependent. The value of `cbitpos` is hypervisor-dependent, and can be obtained through the `sev` element in the capabilities of the domain.
- ④ When memory encryption is enabled, we lose certain bits of the physical address space. The required `reducedPhysBits` element provides this physical address bit reduction. Similarly to `cbitpos`, the value of `reducedPhysBits` is processor-family-dependent and can be obtained through the `sev` element in the domain capabilities.
- ⑤

The required policy element provides the guest policy which must be maintained by the SEV firmware. This policy is enforced by the firmware, and restricts what configuration and operational commands can be performed on the VM by the hypervisor. The guest policy provided when starting the VM is bound to that VM and cannot be changed throughout its lifetime.

- ❶ The optional dhCert element provides the guest owner's base64-encoded Diffie-Hellman (DH) key. The key is used to negotiate a master secret key between the SEV firmware and guest owner. This master secret key is then used to establish a trusted channel between the SEV firmware and guest owner.
- ❷ The optional session element provides the guest owner's base64-encoded session blob, as defined in the SEV API specification. See the LAUNCH\_START section of the SEV specification for the session-blob format.
- ❸ Besides the launchSecurity settings, SEV-encrypted VMs must have the iPXE option ROM disabled on all virtio-net devices. Currently, iPXE is not compatible with SEV-encrypted VMs.

The guest policy is four unsigned bytes with the following definition:

**Table 1. Guest policy definitions**

Bits	Definition
0	If set, debugging of the guest is disallowed
1	If set, sharing keys with other guests is disallowed
2	If set, SEV-ES is required
3	If set, sending the guest to another platform is disallowed
4	If set, the guest must not be transmitted to another platform that is not in the domain
5	If set, the guest must not be transmitted to another platform that is not SEV-capable
6-15	Reserved
16-32	The guest must not be transmitted to another platform with a lower firmware version

## 5. VM installation

**virt-install** supports the installation of SEV and SEV-ES virtual machines. In addition to your standard installation parameters, provide **virt-install** with options to satisfy the VM requirements and the `--launchSecurity` option.

The following example starts a network installation of a SLES 15 SP4 virtual machine protected with SEV-ES.

```
virt-install --name sles15sp4-sev-es --location http://192.168.0.1/install/1  
sles15sp4/x86_64 --disk size=20 --network=bridge=br0,model=virtio,rom.bar=off ①  
--vcpus 4 --memory 4096 --noautoconsole --events on_reboot=destroy --machine q35  
--memtune hard_limit=4563402 --launchSecurity sev,policy=0x07 ② --boot  
firmware=efi ③ --vnc --serial pty
```

- ① The iPXE option ROM is not compatible with SEV-encrypted VMs and must be disabled on all virtio-net devices. While libvirt supports disabling option ROMs using either the `enabled` or `bar` attributes of the `rom` element, `virt-install` only supports disabling option ROMs using the `bar` attribute.
- ② The `launchSecurity` option specifies the type and policy to be enforced by the SEV firmware. The policy setting is described in *Table 1, “Guest policy definitions”*.
- ③ The `boot` option allows specifying many boot-related settings, including the firmware used by the virtual machine. Specifying a firmware type `efi` allows libvirt's firmware auto-selection feature to select an appropriate SEV capable firmware for the virtual machine.

## 6. VM attestation

VM attestation is the process of verifying that trusted software components are correctly instantiated on a trusted compute platform. The process typically involves starting a VM in a paused state, retrieving a launch measurement of the instantiated software components, verifying the measurement, then providing a disk password or other key to unlock the VM and resume the paused boot process. The launch measurement includes cryptographic artifacts provided by the VM owner, with the cryptographic root of trust being the AMD SEV platform. The VM owner can be confident their software components have not been compromised and are running on a trusted platform once the launch measurement has been verified.

The overall attestation process is moderately complex with plenty of opportunity for error. Care must be taken to ensure the process itself is secure. For example, a secure attestation process cannot be executed directly on the hypervisor running the VM, since the goal is to demonstrate the hypervisor has not acted maliciously and contaminated the VM.

Although all the information and APIs required for attestation exist in SLES 15 SP4, SLES 15 SP5 introduces a simple utility called **`virt-qemu-sev-validate`** that can be used to satisfy several attestation use cases. For example, automated tests in quality assurance and small libvirt+KVM deployments that are not managed by large, commercial management stacks.

**`virt-qemu-sev-validate`** is included in the `libvirt-client-qemu` package and supports both offline and online attestation modes. **`virt-qemu-sev-validate`** requires all input for attestation as command-line parameters. Assuming it is invoked on a trusted machine, the results of

**virt-qemu-sev-validate** can be trusted since no information is retrieved from untrusted sources. Online mode is less secure, particularly when executed directly on the hypervisor running the VM.

Regardless of mode, the attestation process of a libvirt+KVM VM starts with creating a VM or Guest Owner (GO) certificate and session blob that is unique for each start of the VM. The certificate and blob can be created with the **sevctl** utility, available in the **sevctl** package. The following example illustrates the use of the **sevctl session** command to create all the prelaunch SEV-related artifacts. The *NAME* parameter is optional and allows a prefix to be specified for the artifact file names. Using the VM name as a prefix is convenient for matching artifacts with VMs later. The path to the platform Diffie-Hellman (DH) certificate and the desired SEV policy are required parameters.

```
#sevctl session --name test-sev /data/sev/pdh.cert 7
```

The **sevctl session** command produces four files: tik.bin, tek.bin, godh.b64 and session.b64. If the optional *NAME* parameter was used, the files are prefixed with the specified value. The transport integrity key (tik.bin) and transport encryption key (tek.bin) are used in the verification stage of the attestation process. The guest owner Diffie-Hellman key (godh.b64) and session blob (session.b64) are copied to the VM XML configuration before starting the VM. See the **dhCert** and **session** subelements of the **launchSecurity** element in the VM configuration section for more details.

After the VM session artifacts have been created and VM XML configuration updated, the VM can be started in a paused state, for example:

```
#virsh -c qemu+ssh://USER_NAME@HOST_NAME/system create --paused /path/to/vm.xml
```

Creating the VM in a paused state allows retrieving the launch measurement from the hypervisor and comparing it to a measurement calculated on a trusted host using trusted inputs. If the measurements compare, the VM owner can be confident the VM has been properly instantiated on the hypervisor and execution can safely be started. The following command demonstrates using the **virsh domLaunchSecInfo** command to retrieve the paused VM launch measurement and other SEV-related information from the hosting hypervisor.

```
#virsh -c qemu+ssh://username@hostname/system domLaunchSecInfo sevtest
sev-measurement: VZjxM-
Slu+UuYkWHN2mAxDVVYXRmL3wqTu84kwk+5QS+40Mii7hs6cMAMXNpmmmyS/
sev-api-major : 1
sev-api-minor : 51
sev-build-id   : 3
sev-policy     : 7
```

The retrieved launch measurement can then be given to **virt-qemu-sev-validate** to verify the VM has been securely instantiated. The following example demonstrates a full offline attestation of the measurement.

```
#virt-qemu-sev-validate --api-major 1 --api-minor 51 --build-id 3 --policy 7 \
--firmware /usr/share/qemu/ovmf-x86_64-sev.bin --tik sevtest_tik.bin --tek
sevtest_tek.bin --num-cpus 4 \
--cpu-family 25 --cpu-model 1 --cpu-stepping 1 \
--measurement QJ0oDpFmWj+bGZzFoMPbAxTuC6QD44W5w88x/hQM8toVsB75ci7V1YDfYoI9GtK
```

It is also possible to use **virt-qemu-sev-validate** in an online mode, where information needed to perform the VM attestation is retrieved from the hosting hypervisor. The following example demonstrates an online attestation of the VM, where only the hosting hypervisor URI, VM name, and associated TIK and TEK are specified. **virt-qemu-sev-validate** retrieves the remaining information, including the measurement itself, from the hosting hypervisor:

```
#virt-qemu-sev-validate --tik sevtest_tik.bin --tek sevtest_tek.bin \
--connect qemu+ssh://USER_NAME@HOST_NAME/system --domain sevtest
```

Once the VM launch measurement has been verified, the VM owner can optionally inject a secret in the VM and resume VM execution. An example of injecting a secret would be providing a key to unlock an encrypted root disk.

```
#virsh -c qemu+ssh://USER_NAME@HOST_NAME/system domsetlauchsecstate sevtest \
--secrethdr hdr-str --secret secret-str
#virsh -c qemu+ssh://USER_NAME@HOST_NAME/system resume sevtest
```

## 7. SEV with KubeVirt

KubeVirt supports running SEV guests starting from the version 0.49.0. The functionality can be activated by enabling the `WorkloadEncryptionSEV` feature gate:

```
>kubectl edit kubevirt kubevirt -n kubevirt
[...]
spec:
  configuration:
    developerConfiguration:
      featureGates:
        - WorkloadEncryptionSEV
[...]
```

To run an SEV-encrypted guest, the virtual machine specification must include the entry `sev: {}` under the `launchSecurity` domain element. Additionally, you need to configure the `firmware`/`bootloader` parameters to use the `efi` option with the `secureBoot` flag set to `disabled`. The corresponding YAML snippet looks similar to the following:

```
[...]
spec:
  domain:
    firmware:
      bootloader:
        efi:
          secureBoot: false
    launchSecurity:
      sev: {}
[...]
```

## 8. Current limitations

SUSE does not recommend using the SEV and SEV-ES features with SUSE Linux products on the first generation AMD EPYC™ 7000 series of processors, code name Naples. It is recommended to use at least the second generation 7002 series processors, code name Rome. Additionally, the following limitations are placed on SEV and SEV-ES VMs.

- The guest operating system running inside an SEV-encrypted VM must contain SEV support. SUSE Linux Enterprise Server 12 SP4 and newer, and all SUSE Linux Enterprise Server 15 releases support SEV.
- Any operations that involve saving and restoring the memory and state of an instance are currently not supported. This means that SEV-encrypted VMs cannot be resumed from snapshots, saved/restored, or live migrated. Encrypted VMs can be shutdown and restarted on another host as normal.
- SEV-encrypted VMs cannot contain directly accessible host devices (that is, PCI passthrough).
- SEV-encrypted VMs are not compatible with Secure Boot. UEFI firmware containing Secure Boot support does not work with SEV or SEV-ES VMs.
- SEV-ES VMs cannot be rebooted from within using **reboot**, **shutdown -r now**, etc. A reboot must be done by shutting down the VM and starting it again. This limitation does not apply to SEV VMs, only SEV-ES.

These limitations will be removed in the future as the hardware, firmware and specific layers of software receive new features.

## 9. More information

- <https://developer.amd.com/sev> — AMD-SEV landing page
- [AMD SEV-KM API Specification \(PDF\)](#)
- [AMD SEV GitHub repository containing examples and tools](#)
- [libvirt SEV configuration settings](#)
- [libvirt knowledge base article on AMD SEV](#)

## 10. Legal notice

Copyright© 2006– 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.

## 11. 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:

1. 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.
2. 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.
3. State on the Title page the name of the publisher of the Modified Version, as the publisher.
4. Preserve all the copyright notices of the Document.
5. Add an appropriate copyright notice for your modifications adjacent to the other copyright notices.
6. 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.
7. Preserve in that license notice the full lists of Invariant Sections and required Cover Texts given in the Document's license notice.
8. Include an unaltered copy of this License.
9. 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.
10. 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.

11. 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.
12. 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.
13. Delete any section Entitled "Endorsements". Such a section may not be included in the Modified Version.
14. Do not retitle any existing section to be Entitled "Endorsements" or to conflict in title with any Invariant Section.
15. 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.



## 1. Introduction

NVIDIA virtual GPU (vGPU) is a graphics virtualization solution that provides multiple virtual machines (VMs) simultaneous access to one physical Graphics Processing Unit (GPU) on the VM Host Server. This article refers to the Volta and Ampere GPU architecture.

## 2. Configuring vGPU manager in VM Host Server

### 2.1. Prepare VM Host Server environment

1. Verify that you have a compatible server and GPU cards. Check specifications for details:

- <https://docs.nvidia.com/grid/gpus-supported-by-vgpu.html>
- <https://docs.nvidia.com/grid/index.html>

2. Verify that VM Host Server is SUSE Linux Enterprise Server 15 SP3 or newer:

```
>cat /etc/issue
Welcome to SUSE Linux Enterprise Server 15 SP3 (x86_64) - Kernel \r (\l).
```

3. Get the vGPU drivers from NVIDIA. To get the software, please follow the steps at <https://docs.nvidia.com/grid/latest/grid-software-quick-start-guide/index.html#redeeming-pak-and-downloading-grid-software>. For example, to install vGPU 13.0, you need the following files:

```
NVIDIA-Linux-x86_64-470.63-vgpu-kvm.run # vGPU manager for the VM host
NVIDIA-Linux-x86_64-470.63.01-grid.run # vGPU driver for the VM guest
```

4. If you are using Ampere architecture GPU cards, verify that VM Host Server supports VT-D/IOMMU and SR-IOV technologies, and that they are enabled in BIOS.

5. Enable IOMMU. Verify that it is included in the boot command line:

```
cat /proc/cmdline
BOOT_IMAGE=/boot/vmlinuz-default [...] intel_iommu=on [...]
```

If not, add the following line to /etc/default/grub.

- For Intel CPUs:

```
GRUB_CMDLINE_LINUX="intel_iommu=on"
```

For AMD CPUs:

```
GRUB_CMDLINE_LINUX="amd_iommu=on"
```

Then generate new GRUB 2 configuration file and reboot:

```
>sudo grub2-mkconfig -o /boot/grub2/grub.cfg
>sudo systemctl reboot
```

### Tip



You can verify that IOMMU is loaded by running the following command:

```
sudo dmesg | grep -e IOMMU
```

6. Enable SR-IOV. Refer to <https://docs.nvidia.com/grid/13.0/grid-vgpu-user-guide/index.html#vgpu-types-tesla-v100-pcie> for useful information.
7. Disable the nouveau kernel module by adding the following line it to the top of the /etc/modprobe.d/50-blacklist.conf file:

```
blacklist nouveau
```

## 2.2. Remove conflicting and pre-installed NVIDIA drivers

On a fresh installation of SUSE Linux Enterprise Server 15 SP7, the open-source NVIDIA desktop driver such as nvidia-open-driver-G06-signed-kmp-default is often automatically installed if an NVIDIA GPU is detected. This pre-installed driver can conflict with the installation of the proprietary NVIDIA vGPU driver. The conflict can lead to errors during the build process of the vGPU kernel modules and/or issues with the correct driver being loaded, preventing the vGPU functionality from working correctly. This section describes how to identify and remove conflicting drivers before proceeding with the NVIDIA vGPU driver installation.

1. Check for pre-installed NVIDIA drivers that can potentially conflict with the vGPU driver that you want to install.
  1. Check for the presence of automatically installed open source NVIDIA driver.

```
>zypper info nvidia-open-driver-G06-signed-kmp-default
```

### Note



If the package is *not installed*, you can skip the rest of this section and proceed to *the section called “Install the NVIDIA KVM driver”*. However, if it is *installed*, perform the rest of the steps to uninstall conflicting drivers and clean up any residues.

2. List all installed NVIDIA-related packages to identify any other potential conflicts:

```
>zypper search --installed-only nvidia*
```

Look for any packages that might conflict with the vGPU driver, such as other desktop driver versions.

## 2. Remove the conflicting NVIDIA drivers.

### 1. Remove the open source NVIDIA driver that is automatically installed.

```
>sudozypper remove nvidia-open-driver-G06-signed-kmp-default
```

### 2. If you identified other conflicting packages in the previous step, remove them similarly.

```
>sudozypper remove CONFLICTING-NVIDIA-PACKAGECONFLICTING-NVIDIA-PACK-  
AGE
```

### 3. Uninstall any manually installed proprietary NVIDIA drivers that you have installed using .run files downloaded from NVIDIA's Web site.

```
>sudo/PATH/T0/NVIDIA-Linux-x86_64-XXX.XX.XX.run --uninstall
```

### 4. *OPTIONAL but RECOMMENDED:* Sometimes, driver installations can leave behind configuration files. To ensure a clean state, remove any potentially conflicting configuration files.

#### Warning



Remove driver configuration files only when you are absolutely sure that it relates to a potentially conflicting NVIDIA driver that you have already uninstalled.

### 1. Search for NVIDIA-related configuration directories.

```
>sudofind /etc/ -name "**nvidia**"
```

### 2. Remove the directories related to conflicting uninstalled NVIDIA drivers.

```
>sudorm -rf /etc/CONFLICTING/NVIDIA/DIRECTORY
```

### 5. Manually load the nvidia-modprobe utility before attempting the vGPU driver installation:

```
>sudomodprobe nvidia-modprobe
```

## 2.3. Install the NVIDIA KVM driver

### 1. Exit from the graphical mode:

```
>sudo init 3
```

### 2. Install kernel-default-devel and gcc packages and their dependencies:

```
>sudo zypper in kernel-default-devel gcc
```

### 3. Download the vGPU software from the NVIDIA portal. Make the NVIDIA vGPU driver executable and run it:

```
>chmod +x NVIDIA-Linux-x86_64-450.55-vgpu-kvm.run  
>sudo ./NVIDIA-Linux-x86_64-450.55-vgpu-kvm.run
```

You can find detailed information about the installation process in the log file `/var/log/nvidia-installer.log`

**Tip**



To enable dynamic kernel-module support, and thus have the module rebuilt automatically when a new kernel is installed, add the `--dkms` option:

```
>sudo ./NVIDIA-Linux-x86_64-450.55-vgpu-kvm.run --dkms
```

4. When the driver installation is finished, reboot the system:

```
>sudo systemctl reboot
```

## 2.4. Verify the driver installation

1. Verify loaded kernel modules:

```
>lsmod | grep nvidia
nvidia_vgpu_vfio      49152  9
nvidia                  14393344  229 nvidia_vgpu_vfio
mdev                      20480  2 vfio_mdev,nvidia_vgpu_vfio
vfio                     32768  6 vfio_mdev,nvidia_vgpu_vfio,vfio_iommu_type1
```

The modules containing the `vfio` string are required dependencies.

2. Print the GPU device status with the `nvidia-smi` command. The output should be similar to the following one:

```

>nvidia-smi
+-----+
| NVIDIA-SMI 470.63      Driver Version: 470.63      CUDA Version: N/
A   |
|-----+
+-----+
| GPU  Name      Persistence-M| Bus-Id      Disp.A | Volatile Uncorr.
ECC  |
| Fan  Temp  Perf  Pwr:Usage/Cap|          Memory-Usage | GPU-Util  Compute
M.   |
|          |          |          |          |          |          |          |
M.   |
|-----+
==| 0  NVIDIA A40          Off  | 00000000:31:00.0 Off  |
0  |
| 0%   46C    P0    39W / 300W |      0MiB / 45634MiB |      0%     De-
fault  |
|          |          |          |          |          |          |
N/A  |
+-----+
+-----+
| Process-
es:
| GPU  GI  CI          PID  Type  Process name          GPU Mem-
ory  |
|       ID  ID          |          |          |          |          |
age  |
|-----+
==|  No running processes
found  |
+-----+
---+

```

3. Check the sysfs file system. For Volta and earlier GPU cards, new directory `mdev_supported_types` is added, for example:

```
cd /sys/bus/pci/devices/00000000\:31\:00.0/mdev_supported_types
```

For Ampere GPU cards, the directory gets created automatically for each virtual function after SR-IOV is enabled.

### 3. Creating a vGPU device

#### 3.1. Create a legacy vGPU device without support for SR-IOV

All the NVIDIA Volta and earlier architecture GPUs work in this mode.

1. Obtain the Bus/Device/Function (BDF) numbers of the host GPU device:

```
>lspci | grep NVIDIA
84:00.0 3D controller: NVIDIA Corporation GV100GL [Tesla V100 PCIe 16GB]
(rev a1)
```

2. Check for the mdev supported devices and detailed information:

```
>ls /sys/bus/pci/devices/0000:84:00.0/mdev_supported_types/
nvidia-105  nvidia-106  nvidia-107  nvidia-108  nvidia-109  nvidia-110
[...]
```

The map of vGPU mdev devices and their type is as follows:

- nvidia-105 to nvidia-109: 1Q 2Q 4Q 8Q 16Q
- nvidia-110 to nvidia-114: 1A 2A 4A 8A 16A
- nvidia-115, nvidia-163, nvidia-217, nvidia-247: 1B 2B 2B4 1B4
- nvidia-299 to nvidia-301: 4C 8C 16C

Refer to <https://docs.nvidia.com/grid/latest/grid-vgpu-user-guide/index.html#vgpu-types-test-la-v100-pcie> for more details.

### 3. Inspect a vGPU device:

```
>cd /sys/bus/pci/devices/0000:03:00.0/mdev_supported_types/
>ls nvidia-105
>cat nvidia-105/description
num_heads=2, frl_config=60, framebuffer=1024M, max_resolution=4096x2160,
max_instance=16
>cat nvidia-105/name
GRID V100-1Q
```

### 4. Generate a unique ID and create an mdev device based on it:

```
>uuidgen
4f3b6e47-0baa-4900-b0b1-284c1ecc192f
>sudo echo "4f3b6e47-0baa-4900-b0b1-284c1ecc192f" > nvidia-105/create
```

### 5. Verify the new mdev device. You can inspect the content of the /sys/bus/mdev/devices directory:

```
>cd /sys/bus/mdev/devices
>ls -l
lrwxrwxrwx 1 root root 0 Aug 30 23:03 86380ffb-8f13-4685-9c48-0e0f4e65fb87
\
-> ../../devices/
pci0000:80/0000:80:02.0/0000:84:00.0/86380ffb-8f13-4685-9c48-0e0f4e65fb87
lrwxrwxrwx 1 root root 0 Aug 30 23:03 86380ffb-8f13-4685-9c48-0e0f4e65fb88
\
-> ../../devices/
pci0000:80/0000:80:02.0/0000:84:00.0/86380ffb-8f13-4685-9c48-0e0f4e65fb88
lrwxrwxrwx 1 root root 0 Aug 30 23:03 86380ffb-8f13-4685-9c48-0e0f4e65fb89
\
-> ../../devices/
pci0000:80/0000:80:02.0/0000:84:00.0/86380ffb-8f13-4685-9c48-0e0f4e65fb89
lrwxrwxrwx 1 root root 0 Aug 30 23:03 86380ffb-8f13-4685-9c48-0e0f4e65fb90
\
-> ../../devices/
pci0000:80/0000:80:02.0/0000:84:00.0/86380ffb-8f13-4685-9c48-0e0f4e65fb90
```

Or you can use the **mdevctl** command:

```
>sudo mdevctl list
86380ffb-8f13-4685-9c48-0e0f4e65fb90 0000:84:00.0 nvidia-299
86380ffb-8f13-4685-9c48-0e0f4e65fb89 0000:84:00.0 nvidia-299
86380ffb-8f13-4685-9c48-0e0f4e65fb87 0000:84:00.0 nvidia-299
86380ffb-8f13-4685-9c48-0e0f4e65fb88 0000:84:00.0 nvidia-299
```

## 6. Query the new vGPU device capability:

```
>sudo nvidia-smi vgpu -q
GPU 00000000:84:00.0
Active vGPUs : 1
vGPU ID : 3251634323
VM UUID : ee7b7a4b-388a-4357-a425-5318b2c65b3f
VM Name : sle15sp3
vGPU Name : GRID V100-4C
vGPU Type : 299
vGPU UUID : d471c7f2-0a53-11ec-af3-38b06df18e37
MDEV UUID : 86380ffb-8f13-4685-9c48-0e0f4e65fb87
Guest Driver Version : 460.91.03
License Status : Licensed
GPU Instance ID : N/A
Accounting Mode : Disabled
ECC Mode : N/A
Accounting Buffer Size : 4000
Frame Rate Limit : N/A
FB Memory Usage
    Total : 4096 MiB
    Used : 161 MiB
    Free : 3935 MiB
Utilization
    Gpu : 0 %
    Memory : 0 %
    Encoder : 0 %
    Decoder : 0 %
Encoder Stats
    Active Sessions : 0
    Average FPS : 0
    Average Latency : 0
FBC Stats
    Active Sessions : 0
    Average FPS : 0
    Average Latency : 0
```

## 3.2. Create a vGPU device with support for SR-IOV

All NVIDIA Ampere and newer architecture GPUs work in this mode.

1. Obtain the Bus/Device/Function (BDF) numbers of the host GPU device:

```
>lspci | grep NVIDIA
b1:00.0 3D controller: NVIDIA Corporation GA100 [A100 PCIe 40GB] (rev a1)
```

2. Enable virtual functions:

```
>sudo /usr/lib/nvidia/sriov-manage -e 00:b1:0000.0
```

### Note



This configuration is not persistent and must be re-enabled after the host reboot.

3. Obtain the Bus/Domain/Function (BDF) of virtual functions on the GPU:

```
>ls -l /sys/bus/pci/devices/0000:b1:00.0/ | grep virtfn
lrwxrwxrwx 1 root root          0 Sep 21 11:58 virtfn0 -> ../../0000:b1:00.4
lrwxrwxrwx 1 root root          0 Sep 21 11:58 virtfn1 -> ../../0000:b1:00.5
lrwxrwxrwx 1 root root          0 Sep 21 11:58 virtfn10 -> ../../0000:b1:01.6
lrwxrwxrwx 1 root root          0 Sep 21 11:58 virtfn11 -> ../../0000:b1:01.7
lrwxrwxrwx 1 root root          0 Sep 21 11:58 virtfn12 -> ../../0000:b1:02.0
lrwxrwxrwx 1 root root          0 Sep 21 11:58 virtfn13 -> ../../0000:b1:02.1
lrwxrwxrwx 1 root root          0 Sep 21 11:58 virtfn14 -> ../../0000:b1:02.2
lrwxrwxrwx 1 root root          0 Sep 21 11:58 virtfn15 -> ../../0000:b1:02.3
lrwxrwxrwx 1 root root          0 Sep 21 11:58 virtfn2 -> ../../0000:b1:00.6
lrwxrwxrwx 1 root root          0 Sep 21 11:58 virtfn3 -> ../../0000:b1:00.7
lrwxrwxrwx 1 root root          0 Sep 21 11:58 virtfn4 -> ../../0000:b1:01.0
lrwxrwxrwx 1 root root          0 Sep 21 11:58 virtfn5 -> ../../0000:b1:01.1
lrwxrwxrwx 1 root root          0 Sep 21 11:58 virtfn6 -> ../../0000:b1:01.2
lrwxrwxrwx 1 root root          0 Sep 21 11:58 virtfn7 -> ../../0000:b1:01.3
lrwxrwxrwx 1 root root          0 Sep 21 11:58 virtfn8 -> ../../0000:b1:01.4
lrwxrwxrwx 1 root root          0 Sep 21 11:58 virtfn9 -> ../../0000:b1:01.5
```

4. Create a vGPU device. Select the virtual function (VF) that you want to use to create the vGPU device and assign it a unique ID.

### Important



Each VF can only create one vGPU instance. To create more vGPU instances, you need to use a different VF.

```
>cd /sys/bus/pci/devices/0000:b1:00.0/virtfn1/mdev_supported_types
>for i in *; do echo "$i" $(cat $i/name) available: $(cat $i/avail*); done
nvidia-468 GRID A100-4C available: 0
nvidia-469 GRID A100-5C available: 0
nvidia-470 GRID A100-8C available: 0
nvidia-471 GRID A100-10C available: 1
nvidia-472 GRID A100-20C available: 0
nvidia-473 GRID A100-40C available: 0
nvidia-474 GRID A100-1-5C available: 0
nvidia-475 GRID A100-2-10C available: 0
nvidia-476 GRID A100-3-20C available: 0
nvidia-477 GRID A100-4-20C available: 0
nvidia-478 GRID A100-7-40C available: 0
nvidia-479 GRID A100-1-5CME available: 0
>uuidgen
f715f63c-0d00-4007-9c5a-b07b0c6c05de
>sudo echo "f715f63c-0d00-4007-9c5a-b07b0c6c05de" > nvidia-471/create
>sudo dmesg | tail
[...]
[ 3218.491843] vfio_mdev f715f63c-0d00-4007-9c5a-b07b0c6c05de: Adding to
iommu group 322
[ 3218.499700] vfio_mdev f715f63c-0d00-4007-9c5a-b07b0c6c05de: MDEV:
group_id = 322
[ 3599.608540] vfio_mdev f715f63c-0d00-4007-9c5a-b07b0c6c05de: Removing
from iommu group 322
[ 3599.616753] vfio_mdev f715f63c-0d00-4007-9c5a-b07b0c6c05de: MDEV: de-
taching iommu
[ 3626.345530] vfio_mdev f715f63c-0d00-4007-9c5a-b07b0c6c05de: Adding to
iommu group 322
[ 3626.353383] vfio_mdev f715f63c-0d00-4007-9c5a-b07b0c6c05de: MDEV:
group_id = 322
```

5. Verify the new vGPU device:

```
>cd /sys/bus/mdev/devices/
>ls
f715f63c-0d00-4007-9c5a-b07b0c6c05de
```

6. Query the new vGPU device capability:

```
>sudo nvidia-smi vgpu -q
GPU 00000000:B1:00.0
Active vGPUs : 1
vGPU ID : 3251634265
VM UUID : b0d9f0c6-a6c2-463e-967b-06cb206415b6
VM Name : sles15sp2-gehc-vm1
vGPU Name : GRID A100-10C
vGPU Type : 471
vGPU UUID : 444f610c-1b08-11ec-9554-ebd10788ee14
MDEV UUID : f715f63c-0d00-4007-9c5a-b07b0c6c05de
Guest Driver Version : N/A
License Status : N/A
GPU Instance ID : N/A
Accounting Mode : N/A
ECC Mode : Disabled
Accounting Buffer Size : 4000
Frame Rate Limit : N/A
FB Memory Usage
  Total : 10240 MiB
  Used : 0 MiB
  Free : 10240 MiB
Utilization
  Gpu : 0 %
  Memory : 0 %
  Encoder : 0 %
  Decoder : 0 %
Encoder Stats
  Active Sessions : 0
  Average FPS : 0
  Average Latency : 0
FBC Stats
  Active Sessions : 0
  Average FPS : 0
  Average Latency : 0
```

### 3.3. Creating a MIG-backed vGPU

#### Important



SR-IOV must be enabled to create vGPUs and assign them to guest VMs.

1. Enable MIG mode for a GPU:

```
>sudo nvidia-smi -i 0 -mig 1
Enabled MIG Mode for GPU 00000000:B1:00.0
All done.
```

2. Query the GPU instance profile:

GPU instance profiles:		ID	Instances	Memory	P2P	SM	DEC
ENC	OFA		Free/Total	GiB		CE	JPEG
		0	MIG 1g.5gb	19 7/7	4.75	No 14	0
		0				1	0
		0	MIG 1g.5gb+me	20 1/1	4.75	No 14	1
		1				1	1
		0	MIG 2g.10gb	14 3/3	9.75	No 28	1
		0				2	0
		0	MIG 3g.20gb	9 2/2	19.62	No 42	2
		0				3	0
		0	MIG 4g.20gb	5 1/1	19.62	No 56	2
		0				4	0
		0	MIG 7g.40gb	0 1/1	39.50	No 98	5
		1				7	1

3. Create a GPU instance specifying '5' as a GPU profile instance ID and optionally create a Compute Instance on it, either on the host server or within the guest:

```
>sudo nvidia-smi mig -cgi 5
Successfully created GPU instance ID 1 on GPU 0 using profile MIG 4g.20gb (ID 5)
>sudo nvidia-smi mig -cci -gi 1
Successfully created compute instance ID 0 on GPU 0 GPU instance ID 1 using profile MIG 4g.20gb (ID 3)
```

4. Verify the GPU instance:

```

>sudo nvidia-smi
Tue Sep 21 11:19:36 2021
+-----+
| NVIDIA-SMI 470.63      Driver Version: 470.63      CUDA Version: N/
A   |
+-----+
| GPU  Name      Persistence-M| Bus-Id      Disp.A | Volatile Uncorr.
ECC  |
| Fan  Temp     Perf  Pwr:Usage/Cap|      Memory-Usage | GPU-Util  Compute
M.   |
| M.   |                               |              |              |
+-----+
==| 0  NVIDIA A100-PCI...  On  | 00000000:B1:00.0 Off |
On | N/A   38C     P0     38W / 250W |      0MiB / 40536MiB |      N/A     De-
fault |
| abled |
+-----+
+-----+
| MIG de-
vices:
+-----+
| GPU  GI  CI  MIG |      Memory-Usage |      Vol|
Shared  ID  ID  Dev |              BAR1-Usage |  SM      Unc|  CE   ENC  DEC  OFA
JPG|
| ECC |
+-----+
==| 0  1  0  0 |      0MiB / 20096MiB |  56      0 |  4  0  2  0
0  |
|           |      0MiB / 32767MiB |
+-----+
+-----+
| Process-
es:
| GPU  GI  CI      PID  Type  Process name          GPU Mem-
ory  |
| ID  ID      |          |          |          |
age  |
+-----+
==|
| No running processes
+-----+

```

5. Use the MIG instance. You can use the instance directly with the UUID—for example, assign it to a container or CUDA process.

You can also create a vGPU on top of it and assign it to a VM guest. The procedure is the same as for the vGPU with SR-IOV support. Refer to *the section called “Create a vGPU device with support for SR-IOV”*.

```
>sudo nvidia-smi -L
GPU 0: NVIDIA A100-PCIE-40GB (UUID: GPU-ee14e29d-dd5b-2e8e-
eeaf-9d3debd10788)
  MIG 4g.20gb   Device  0: (UUID: MIG-fed03f85-fd95-581b-837f-
d582496d0260)
```

## 4. Assign the vGPU device to a VM Guest

### 4.1. Assign by libvirt

1. Create a libvirt-based virtual machine (VM) with UEFI support and a normal VGA display.
2. Edit the VM's configuration by running **virsh edit VM-NAME**.
3. Add the new mdev device with the unique ID you used when creating the vGPU device to the <devices> section.

**Note**



If you are using Q-series, use `display='on'` instead.

```
<hostdev mode='subsystem' type='mdev' managed='no' model='vfio-pci' dis-
play='off'>
  <source>
    <address uuid='4f3b6e47-0baa-4900-b0b1-284c1ecc192f' />
  </source>
  <address type='pci' domain='0x0000' bus='0x00' slot='0x0a' func-
tion='0x0' />
</hostdev>
```

### 4.2. Assign by QEMU

Add the following device to the QEMU command line. Use the unique ID that you used when creating the vGPU device:

```
-device vfio-pci,sysfsdev=/sys/bus/mdev/devices/4f3b6e47-0baa-4900-
b0b1-284c1ecc192f
```

## 5. Configuring vGPU in VM Guest

### 5.1. Prepare the VM Guest

- During VM Guest installation, disable secure boot, enable the SSH service, and select wicked for networking.
- Disable the nouveau video driver. Edit the file `/etc/modprobe.d/50-blacklist.conf` and add the following line to its upper section:

```
blacklist nouveau
```

## Important



Disabling nouveau works after you re-generate the initrd image with dracut, and then reboot the VM Guest.

## 5.2. Install the vGPU driver in the VM Guest

1. Install the following packages and their dependencies:

```
>sudo zypper install kernel-default-devel libglvnd-devel
```

2. Download the vGPU software from the NVIDIA portal. Make the NVIDIA vGPU driver executable and run it:

```
>chmod +x NVIDIA-Linux-x86_64-470.63.01-grid.run  
>sudo ./NVIDIA-Linux-x86_64-470.63.01-grid.run
```

## Tip



To enable dynamic kernel module support to get the module rebuilt automatically when a new kernel is installed, add the `--dkms` option:

```
>sudo ./NVIDIA-Linux-x86_64-470.63.01-grid.run --dkms
```

3. During driver installation, select to run the **nvidia-xconfig** utility.
4. Verify the driver installation by checking the output of the **nvidia-smi** command:

```
>sudo nvidia-smi
+-----+
| NVIDIA-SMI 470.63.01      Driver Version: 470.63.01      CUDA Version:
11.4      |
+-----+
| GPU  Name      Persistence-M| Bus-Id      Disp.A | Volatile Uncorr.
ECC      |
| Fan  Temp  Perf  Pwr:Usage/Cap| Memory-Usage | GPU-Util  Compute
M.      |
|          |          |          |          |          |          |
M.      |
+-----+
==| 0  GRID A100-10C      On  | 00000000:07:00.0 Off  |
0  |
| N/A  N/A    P0    N/A /  N/A | 930MiB / 10235MiB | 0%      De-
fault  |
|          |          |          |          |          |          |
abled  |
+-----+
+-----+
| Process-
es:
| GPU  GI  CI      PID  Type  Process name          GPU Mem-
ory  |
|       ID  ID
age  |
|          |          |
+-----+
==|  No running processes
found  |
+-----+
-----+
```

## 6. Licensing vGPU in the VM Guest

1. Create the configuration file `/etc/nvidia/gridd.conf` based on `/etc/nvidia/gridd.conf.template`.
2. 1. For licenses that are served from the NVIDIA License System, update the following options:

### FeatureType

For GPU passthrough, set `FeatureType` to 4 for computing and 2 for graphic purposes. In case of a virtual GPU, whatever vGPU type is created via `mdev` determines the feature set that is enabled in VM Guest.

### ClientConfigTokenPath

Optional: To store the client configuration token in a custom location, add the `ClientConfigTokenPath` configuration parameter on a new line as `ClientConfigTokenPath="PATH_TO_TOKEN"`. By default, the client searches for

the client configuration token in the /etc/nvidia/ClientConfigToken/ directory.

Copy the client configuration token to the directory in which you want to store it.

2. For licenses that are served from the legacy NVIDIA vGPU software license server, update the following options:

#### **ServerAddress**

Add your license server IP address.

#### **ServerPort**

Use the default "7070" or the port configured during the server setup.

#### **FeatureType**

For GPU passthrough, set FeatureType to 4 for computing and 2 for graphic purposes. In case of a virtual GPU, whatever vGPU type is created via **mdev** determines the feature set that is enabled in VM Guest.

3. Restart the nvidia-gridd service:

```
>sudo systemctl restart nvidia-gridd.service
```

4. Inspect the log file for possible errors:

```
>sudo grep gridd /var/log/messages
[...]
Aug 5 15:40:06 localhost nvidia-gridd: Started (4293)
Aug 5 15:40:24 localhost nvidia-gridd: License acquired successfully.
```

## **7. Configuring a graphics mode**

### **7.1. Create or update the /etc/X11/xorg.conf file**

1. If there is no /etc/X11/xorg.conf on the VM Guest, run the **nvidia-xconfig** utility.
2. Query the GPU device for detailed information:

```
>nvidia-xconfig --query-gpu-info
Number of GPUs: 1

GPU #0:
Name      : GRID V100-160
UUID      : GPU-089f39ad-01cb-11ec-89dc-da10f5778138
PCI BusID : PCI:0:10:0

Number of Display Devices: 0
```

3. Add GPU's BusID to /etc/X11/xorg.conf, for example:

```
Section "Device"
Identifier "Device0"
Driver "nvidia"
BusID "PCI:0:10:0"
VendorName "NVIDIA Corporation"
EndSection
```

## 7.2. Verify the graphics mode

Verify the following:

- A graphic desktop is booted correctly.
- The 'X' process of a running X-server is running in GPU:

```
>nvidia-smi
+
---+
| NVIDIA-SMI 470.63.01      Driver Version: 470.63.01      CUDA Version:
11.4      |
|-----+
| GPU  Name      Persistence-M| Bus-Id      Disp.A  | Volatile Uncorr.
ECC  |
| Fan  Temp  Perf  Pwr:Usage/Cap|          Memory-Usage | GPU-Util  Compute
M.  |
|          |          |          |          |          |          |          |
M.  |
|-----+
==|
| 0  GRID V100-4C          On   | 00000000:00:0A.0 Off  |
N/A  |
| N/A  N/A    P0    N/A /  N/A | 468MiB / 4096MiB | 0%  De-
fault  |
|          |          |          |          |          |          |
N/A  |
+-----+
+-----+
+
---+
| Process-
es:
| GPU  GI  CI      PID  Type  Process name          GPU Mem-
ory  |
|       ID  ID
|          |          |
|          |          |
|-----+
==|
| 0  N/A  N/A      1921    G  /usr/bin/X
76MiB  |
| 0  N/A  N/A      1957    G  /usr/bin/gnome-shell
87MiB  |
+-----+
---+
```

### 7.3. Remote display

You need to install and configure the VNC server package `x11vnc` inside the VM Guest, and start it with the following command:

```
>sudo x11vnc -display :0 -auth /run/user/1000/gdm/Xauthority -forever -shared -ncache -bg -usepw -geometry 1900x1080
```

You can use **virt-manager** or **virt-viewer** to display the graphical output of a VM Guest.

## Important



For a libvirt-based VM Guest, verify that its XML configuration includes `display=on` as suggested in the section called "Assign by libvirt".

## 8. Configuring compute mode

1. Download and install the CUDA toolkit. You can find it at [https://developer.nvidia.com/cuda-downloads?tar-get\\_os=Linux&target\\_arch=x86\\_64&Distribution=SLES&target\\_version=15&target\\_type=runfile\\_id](https://developer.nvidia.com/cuda-downloads?tar-get_os=Linux&target_arch=x86_64&Distribution=SLES&target_version=15&target_type=runfile_id)
2. Download CUDA samples from <https://github.com/nvidia/cuda-samples>.
3. Run CUDA sampling example:

```
>cd YOUR_GIT_CLONE_LOCATION/cuda-samples/Samples/0_Introduction/clock
> make
/usr/local/cuda/bin/nvcc -ccbin g++ -I../../common/inc -m64 --threads 0
-gencode arch=compute_35,code=sm_35 -gencode arch=compute_37,code=sm_37 -
gencode
[...]
mkdir -p ../../bin/x86_64/linux/release
cp clock ../../bin/x86_64/linux/release
>./clock
CUDA Clock sample
GPU Device 0: "Volta" with compute capability 7.0
Average clocks/block = 2820.718750
```

## 9. Additional tasks

This section introduces additional procedures that may be helpful after you have configured your vGPU.

### 9.1. Disabling Frame Rate Limiter

Frame Rate Limiter (FRL) is enabled by default. It limits the vGPU to a fixed frame rate, for example, 60fps. If you experience a bad graphic display, you may need to disable FRL, for example:

```
>sudo echo "frame_rate_limiter=0" > /sys/bus/usb/devices/
86380ffb-8f13-4685-9c48-0e0f4e65fb87/nvidia/vgpu_params
```

### 9.2. Enabling/Disabling Error Correcting Code (ECC)

Since the NVIDIA Pascal architecture, NVIDIA GPU Cards support ECC memory to improve data integrity. ECC is also supported by software since NVIDIA vGPU 9.0.

To enable ECC:

```
>sudo nvidia-smi -e 1
>nvidia-smi -q
Ecc Mode
  Current : Enabled
  Pending : Enabled
```

To disable ECC:

```
>sudo nvidia-smi -e 0
```

### 9.3. Black screen in Virt-manager

If you see only a black screen in Virt-manager, press **Alt**–**Ctrl**–**2** from Virt-manager viewer. You should be able to get in the display again.

### 9.4. Black screen in VNC client when using a non-QEMU VNC server

Use the xvnc server.

### 9.5. Kernel panic occurs because the Nouveau and NVIDIA drivers compete on GPU resources

The boot messages looks as follows:

Make sure to run **mkintrd** and reboot after disabling the Nouveau driver. Refer to the section called “*Prepare the VM Guest*”.

## 9.6. Filing an NVIDIA vGPU bug

While filing an NVIDIA vGPU-related bug report to us, please attach the vGPU configuration data `nvidia-bug-report.log.gz` collected by the `nvidia-bug-report.sh` utility. Make sure you cover both VM Host Server and VM Guest.

## 9.7. Configuring a License Server

Refer to <https://docs.nvidia.com/grid/ls/latest/grid-license-server-user-guide/index.html>.

## 10. For more information

NVIDIA has an extensive documentation on vGPU. Refer to <https://docs.nvidia.com/grid/latest/grid-vgpu-user-guide/index.html> for details.

## 11. NVIDIA virtual GPU background


### 11.1. NVIDIA GPU architectures

There are two types of GPU architectures:

#### Time-sliced vGPU architecture

Introduced on GPUs that are based on the NVIDIA Ampere GPU architecture. Only Ampere GPU cards can support MIG-backed vGPU.

**Figure 1. Time-sliced architecture (source: <https://docs.nvidia.com/grid/latest/grid-vgpu-user-guide/index.html>)**




#### Multi-Instance GPU (MIG) vGPU architecture

All GPU cards support time-sliced vGPU. To do so, Ampere GPU cards use the Single Root I/O Virtualization (SR-IOV) mechanism, while Volta and the earlier GPU cards use the mediated device mechanism. Volta and the earlier architecture are based on mediated device

mechanism. These two mechanisms are transparent to a VM. However, they need different configurations from the host side.

**Figure 2. MIG-backed architecture** (source: <https://docs.nvidia.com/grid/latest/grid-vgpu-user-guide/index.html>)



## 11.2. vGPU types

Each physical GPU can support several different types of vGPUs. vGPU types have a fixed amount of framebuffer, the number of supported display heads, and maximum resolutions. NVIDIA has four types of vGPUs: A, B, C and Q-series. SUSE currently supports Q and C-series.

**Table 1. vGPU types**


vGPU series	Optimal workload
Q-series	Virtual workstations for creative and technical professionals who require the performance and features of the NVIDIA Quadro technology.
C-series	Compute-intensive server workloads, for example, artificial intelligence (AI), deep learning, or high-performance computing (HPC).
B-series	Virtual desktops for business professionals and knowledge workers.
A-series	Application streaming or session-based solutions for virtual applications users.

## 11.3. Valid vGPU configurations on a single GPU

### 11.3.1. Time-sliced vGPU configurations

For time-sliced vGPUs, all vGPUs types must be the same:


**Figure 3. Example time-sliced vGPU configurations on NVIDIA Tesla M60 (source: <https://docs.nvidia.com/grid/latest/grid-vgpu-user-guide/index.html>)**



### 11.3.2. MIG-backed vGPU configurations

For MIG-backed vGPUs, vGPUs can be both homogeneous and mixed-type:

**Figure 4. Example MIG-backed vGPU configurations on NVIDIA A100 PCIe 40 GB (source: <https://docs.nvidia.com/grid/latest/grid-vgpu-user-guide/index.html>)**



## 12. 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:

1. 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.
2. 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.
3. State on the Title page the name of the publisher of the Modified Version, as the publisher.
4. Preserve all the copyright notices of the Document.
5. Add an appropriate copyright notice for your modifications adjacent to the other copyright notices.
6. 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.
7. Preserve in that license notice the full lists of Invariant Sections and required Cover Texts given in the Document's license notice.
8. Include an unaltered copy of this License.
9. 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.
10. 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.

11. 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.
12. 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.
13. Delete any section Entitled "Endorsements". Such a section may not be included in the Modified Version.
14. Do not retitle any existing section to be Entitled "Endorsements" or to conflict in title with any Invariant Section.
15. 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.

# Raspberry Pi Quick Start

Fabian Vogt, Release Engineer, SUSE

Jay Kruemcke, Product Manager, SUSE

Andreas Faerber, Project Manager,  
SUSE



Koordinierungsstelle  
für IT-Standards

This guide contains an overview of SUSE Linux Enterprise Server for Arm on the Raspberry Pi\* platform and will guide you through the setup procedure.


## 1. Platform overview

To use SUSE Linux Enterprise Server for Arm on Raspberry Pi, a 64-bit Arm\* compatible Raspberry Pi\* is required. SUSE Linux Enterprise Server for Arm15 SP7 is tested to work on Raspberry Pi 3 Model A+, Model B, Model B+, and Raspberry Pi 4 Model B boards as well as on Raspberry Pi Compute Module 3 and 3+ on MyPi Industrial IoT Integrator Board.


### 1.1. Technical details of the Raspberry Pi 3 model B

Raspberry Pi is a series of small single-board computers based on a System-on-a-Chip (SoC) by Broadcom\*, featuring various peripherals on the board.

**Figure 1. Overview of Raspberry Pi 3 model B connectors, © Efa / English Wikipedia / CC BY-SA 3.0**



**Figure 2. Photo of Raspberry Pi 3 model B connectors, © Evan-Amos / own work / public domain**



### Selected features of Raspberry Pi 3 model B/B+

#### CPU

The Broadcom BCM2837 SoC includes a quad-core Arm\* Cortex\*-A53 Application Processor supporting the Armv8 32-bit and 64-bit instruction sets. With the default configuration, it is clocked up to 1.2 GHz.

#### RAM

1024 MiB DDR2 memory mounted on the back of the board.

#### Graphics

Broadcom\* VideoCore\* IV providing OpenGL\* ES 2.0 support. Displays can be connected over HDMI\*, composite (TRRS jack), or MIPI\* DSI\* (ribbon cable).

#### Ethernet

A USB Ethernet controller on the board provides 10/100 Mbit/s Ethernet (Model B) or 10/100/1000 Mbit/s Ethernet with a maximum throughput of 300 Mbit/s (Model B+).

#### WLAN

The BCM43438 chip on Model B supports IEEE-802.11b, IEEE-802.11g, and IEEE-802.11n in the 2.4 GHz band. It also provides Bluetooth 2.0 to 4.1 (Low Energy). The BCM43455 chip on Model B+ supports IEEE-802.11b, IEEE-802.11g, IEEE-802.11n, and IEEE-802.11ac in the 2.4 GHz and 5 GHz bands. It provides Bluetooth 4.2 (Low Energy).

## Storage

The microSDHC card slot allows for a memory card to be inserted as the primary boot medium.

## Power

Raspberry Pi 3's main power source is the Micro USB connector. If your Raspberry Pi comes with a power supply, it is recommended to use the bundled power supply only.

## USB

A total of four USB 2.0 ports are available.

## Connectors

A 0.1 inch multi-function pin header is also available. Note that not all functionality of this header is exposed in SUSE Linux Enterprise Server for Arm15 SP7.

## 1.2. SUSE Linux Enterprise Server for Arm 15 SP7

SUSE Linux Enterprise Server for Arm is the first fully supported commercial Linux operating system product available for Raspberry Pi. You can purchase subscriptions which entitle you to receive all released bug and security fixes, feature updates, and technical assistance from SUSE's worldwide support. Learn more about subscription and support options at [https://www.suse.com/support/programs/subscriptions/?id=SUSE\\_Linux\\_Enterprise\\_Server](https://www.suse.com/support/programs/subscriptions/?id=SUSE_Linux_Enterprise_Server).

### Trial version



If you want to try out SUSE Linux Enterprise Server for Arm15 SP7 on Raspberry Pi, SUSE will provide you with a trial version. This gives you access to free patches and updates for a period of 60 days. You must log in to the SUSE Customer Center at <https://scc.suse.com/> using your Customer Center account credentials to receive this free offer. If you do not have a Customer Center account, you must create one to take advantage of the trial version.

## Minimum System Requirements for Installation

- Raspberry Pi 3 Model A+, B, or B+, or Raspberry Pi 4 Model B
- microSD card with at least 8 GB capacity
- USB keyboard, mouse
- HDMI cable and monitor
- Power supply with at least 2.5 A capacity

### 1.2.1. Differences compared to Raspberry Pi OS

Raspberry Pi OS is the de-facto default distribution for Raspberry Pi. The following paragraphs provide a short overview of differences between SUSE Linux Enterprise Server for Arm on Raspberry Pi and Raspberry Pi OS.

#### Based on upstream kernel

Raspberry Pi OS uses a kernel with modifications especially for Raspberry Pi. SUSE Linux Enterprise Server for Arm uses the default SUSE Linux Enterprise kernel for AArch64, which is derived from the official mainline kernel.

#### AArch64 instruction set

SUSE Linux Enterprise Server for Arm on Raspberry Pi is the first distribution for Raspberry Pi using the AArch64 instruction set.

#### Boot process

In Raspberry Pi OS, the kernel is loaded directly. This is not supported by SUSE Linux Enterprise Server for Arm, where the U-Boot boot loader is used to provide an EFI boot environment. A GRUB2 EFI binary is chain-loaded to provide a graphical boot screen.


#### Root file system

SUSE Linux Enterprise Server for Arm for Raspberry Pi uses Btrfs as the file system for the root partition. Compression is enabled by default for better SD card performance.

### 1.2.2. YaST

YaST is the installation and configuration framework for SUSE Linux Enterprise. It is popular for its ease of use, flexible graphical interfaces and the capability to customize your system quickly during and after installation. YaST can be used to configure your entire system: You can configure hardware, set up networking, manage system services, and tune your security settings. All these tasks can be accessed from the YaST control center. To start it, choose YaST in the menu or run the command `xdg-su -c yast2`. You will be prompted to enter the password of the root user.

**Figure 3. The YaST control center**



When started, YaST shows an overview of available modules (*Figure 3, “The YaST control center”*). Simply click an icon to open a module.

### 1.2.3. Zypper

Zypper is the package manager for SUSE Linux Enterprise. It is the tool for installing, updating and removing packages and for managing repositories.

The general syntax for Zypper invocations is:

```
zypper [global-options]command[command-options][arguments] ...
```

#### Short command form



For most commands, there is both a short and a long form. An overview is available with **zypper --help**.

#### Installing a package

```
#zypper install mplayer
```

#### Removing a package

```
#zypper remove mplayer
```

## List available patches

```
>zypper list-patches
```

## Install available patches

```
#zypper patch
```

### Installing software updates



The recommended way to install available software updates is by using the YaST Online Updater. To start it, choose “Online Update” in “Settings” under “Desktop Apps” in the IceWM menu.

## 1.2.4. Limitations

### Graphics not hardware-accelerated

X.Org hardware acceleration is disabled to improve system stability and reliability.

For other limitations, refer to the online version of the Release Notes at <https://www.suse.com/releasenotes/aarch64/SUSE-SLES/15-SP7/>.

## 2. Installation

SUSE Linux Enterprise Server for Arm for Raspberry Pi is distributed as an XZ-compressed image file for microSD cards. This section will guide you through the process of preparing the card for the first boot. If you already have a microSD card containing the image, you can skip this section and go straight to *the section called “First boot”*.

### SD card space requirements



It is recommended to use a card with a capacity of at least 8 GB.

### All data on the card will be lost!



By following the procedure below, all data on the SD card will be overwritten and therefore irrevocably lost. Be very careful when choosing the destination device of the image writing process.

## 2.1. Preparing the card on Linux

Before and after you plug in the SD card, run the **lsblk** command. Between the two runs of **lsblk**, there should be a difference of one or more lines. The first column and first row is the name of

the node representing the SD card in your system. To write the image to the card, use the **dd** command:

```
xz -cd IMAGE | dd of=/dev/SDCARDDEV bs=4096
```

### Example 1. Writing the image to the card using dd

This command decompresses the image SLES15-SP7-Minimal-Image.aarch64-15.7-RaspberryPi-GM.raw.xz to the SD card mmcblk0:


```
>xz -cd SLES15-SP7-Minimal-Image.aarch64-15.7-RaspberryPi-GM.raw.xz | sudo dd of=/dev/mmcblk0 bs=4096 iflag=fullblock status=progress
```

## 2.2. Preparing the card on Microsoft Windows operating systems


The following steps will guide you through the installation of the image onto the SD card on Microsoft Windows operating systems. You need to meet these prerequisites:

- Latest available image of SUSE Linux Enterprise Server for Arm for Raspberry Pi
- Easy 7-Zip from <http://www.e7z.org/>
- Win32 Disk Imager from <https://sourceforge.net/projects/win32diskimager/>

1. Open the downloaded image using Easy 7-Zip. Extract the file content into a directory with sufficient free space.



2. Run Win32 Disk Imager and select the extracted file as *Image File*. Then, choose the correct drive letter as *Device*. Click *Write* to start the procedure.



### Finding the correct device




If you are unsure which drive letter to choose from in the list, remove the SD card and run Win32 Disk Imager again. The option that disappeared is the right target device.

## 2.3. Preparing the card on macOS

The following steps will guide you through the installation of the image onto the SD card on macOS. You need to meet these prerequisites:

- Latest available image of SUSE Linux Enterprise Server for Arm for Raspberry Pi
- The Unarchiver from <https://unarchiver.c3.cx/unarchiver>

1. Open Finder at the location where the downloaded image is stored. Use the *Open With* selection to choose **The Unarchiver** instead of the default **Archive** utility and extract the content into a directory with sufficient free space.



2. Open a terminal window and change into the destination directory of the uncompressed image. Enter **diskutil list** before and after the SD card is inserted to find out which device to use.
3. Run **diskutil unmountDisk /dev/diskX**, where X is the disk number from the previous step.
4. Run **sudo dd bs=4096 if=*imageFile.raw* of=/dev/diskX**, where X is the disk number and *imageFile.raw* is the name of the uncompressed image.

```
>sudo dd bs=4096 if=SLES15-SP7-Minimal-Image.aarch64-15.7-RaspberryPi-GM.raw.xz of=/dev/disk4
Password:
5550+0 records in
5550+0 records out
5819596800 bytes transferred in 1131.796649 secs (5141910 bytes/sec)
```

5. Now unmount the disk (which is now labelled *EFI*) as usual.

### 3. First boot

To boot Raspberry Pi from the prepared microSD, insert the card into the card slot, then connect a display, an Ethernet cable, a USB keyboard, and a mouse. Power Raspberry Pi. If everything works, you should see the text on the screen.



#### Operation without mouse

The JeOS First Boot Assistant can be controlled using the keyboard only. Each user interface element has an accelerator, shown as a highlighted character. To activate an accelerator, press **Alt** and the appropriate character.

During the first boot, the system expands to make use of the entire storage card. This operation may take a while.

#### 3.1. Initial system setup with JeOS First Boot Assistant

After the boot is completed, the JeOS First Boot Assistant leads you through the initial system set-up.


## JeOS First Boot Assistant via a serial connection




To run the JeOS First Boot Assistant wizard via a serial connection, you need to remove the `console=tty` from the kernel command line `console=ttyS0,115200` `console=tty` in GRUB 2 even if there is no HDMI cable connected.

The minimal JeOS image comes without language selection. After the assistant finishes, you will be able to install and select your system locale. You can start JeOS First Boot Assistant at any time using the **jeos-config** command.


1. First, the keyboard layout needs to be configured. Use the arrow keys or first letter and select **OK**.




2. Read the License Agreement. You need to accept it.




3. Select the time zone by using the arrow keys or first letter.




4. Select OK to proceed.
5. You will be asked to provide a root password now. Do not forget what you enter here, as you will need it for administration.



6. Confirm the root password you just provided by typing it in again.



7. Note that you should register your system after installation, as detailed in *the section called "Registration process"*.




### Deferring registration




You will not receive updates or patches until you register using *SUSEConnect*.

8. Select *OK* to continue.
9. Finally, you are offered the option to configure a wireless network. If you have a wired Ethernet connection, you can skip this step by selecting *No*.




If you select *Yes*, you will see the following screens:


1. Select a wireless network to connect to and proceed with *OK*.



2. Select the authentication method for the wireless network and proceed with *OK*.



3. Enter the passphrase for the wireless network and proceed with *OK*.



### WPA EAP authentication



WPA authentication supports just the mschapv2 EPA authentication, which is used as default.

If you chose to skip configuration of wireless networks or if you fully completed their configuration, the JeOS First Boot Assistant will exit. The system continues to boot to a login prompt.

You can now log in as `root` user with the password you chose earlier.

### Canceling the JeOS First Boot Assistant



If at any point you select to *Cancel* the assistant, the system will shut down. You can then disconnect the power supply, make any necessary changes, and start over.

## 3.2. Registration process

It is very important to register your SUSE Linux Enterprise Server for Arm subscription to ensure full functionality of your Raspberry Pi system. The SD card image provided by SUSE contains a minimal set of packages that are intended for the initial boot process and to get your Raspberry Pi onto the network.

When you have registered your SUSE Linux Enterprise Server for Arm subscription, you can download other packages you may need, such as compilers. The SUSE Linux Enterprise Server version that runs on your Raspberry Pi is the same version that runs on AMD64/Intel 64, POWER, IBM Z, or on other Arm-based systems.

## Setting the clock



Because Raspberry Pi does not have a persistent Real Time Clock, make sure that the clock is set to the current date and time before attempting to use Zypper or YaST to install additional packages.

You can initially register your system with the *SUSEConnect* tool. Afterwards, you can also make use of the *YaST Product Registration* module.

## Evaluation code




Sixty-day evaluation subscriptions may be requested at the following page: <https://www.suse.com/products/arm/>.

**After you have obtained a registration code from a subscription card, you need to activate your subscription on the SUSE Customer Center at:**


<https://scc.suse.com/>

To register your subscription, perform the following steps:


1. Log in to the SUSE Customer Center using a browser on another machine. Create an account if required:




2. Click *Manually Activate Subscriptions*:




3. Click the dialog field *Activate a single subscription*:



4. Enter the registration code and accept the terms and conditions:



5. Confirm the subscription activation and the organization assignment. Click *Activate*:



6. Your subscription is now active and ready to be used:



**Subscription successfully activated and associated with your organization.**

To register the system, use the *SUSEConnect* tool as instructed earlier:

```
#SUSEConnect -e company@example.com -r YOUR_CODE
Registering system to SUSE Customer Center
Using E-Mail: company@example.com

Announcing system to https://scc.suse.com ...

Activating SLES 15.7 aarch64 ...
-> Adding service to system ...

Activating sle-module-basesystem 15.7 aarch64 ...
-> Adding service to system ...
-> Installing release package ...

Activating sle-module-server-applications 15.7 aarch64 ...
-> Adding service to system ...
-> Installing release package ...

Successfully registered system
```

This may take a few moments. Afterwards, you will have access to online repositories for installing more packages, as described in *the section called “Zypper”*.

## Example 2. Installing YaST

Later sections may assume you have installed YaST for system configuration:

```
#zypper in -t pattern yast2_basis
```



### Module availability

Your subscription gives you access to more than the above *Basesystem Module* and *Server Applications Module*. To activate additional modules, you can use the *YaST System Extensions* module.

## 4. Initial system configuration

In this section, it is explained how to perform the initial system configuration for SUSE Linux Enterprise Server for Arm for Raspberry Pi 15 SP7.

## 4.1. Changing the language

Follow these steps to change the default language:

1. Install the `glibc-locale` package:

```
#zypper in glibc-locale
```

2. Set the language as desired:

```
#localectl set-locale LANG=de_DE.UTF-8
```



### Translation packages

Certain applications may require you to install an additional `-lang` package before you can see texts translated to the chosen language.

## 4.2. Changing the host name

Follow these steps to change the default host name:

1. Open the YaST network module either by running `yast2 lan` or clicking the network icon in the YaST control center.
2. Select *Hostname/DNS* in the top tab bar.
3. Type the new host name into the *Hostname* field.
4. Click *OK* to save the change. After YaST exits, you need to log out and in again.

## 4.3. Setting up networking

The default configuration has DHCP enabled on the Ethernet port. If that suits your network environment, you can skip this section. If you require the use of a static IP address, use YaST:

1. Open the YaST network module either by running `yast2 lan` or clicking the network icon in the YaST control center.
2. In YaST, you will see the network interface being selected. Select the built-in Ethernet and choose *Edit* to open the address configuration.
3. Select *Statically Assigned IP Address* and type in the desired values for *IP Address* and *Subnet Mask*. Click *Next*.
4. With a static network configuration, you will also need to specify a DNS server (if applicable) and a gateway. For the gateway, select *Routing* in the tab bar and enter the IPs of the gateways into the specific fields.
5. The DNS server is set in the *Hostname/DNS* tab. After choosing that tab, enter the IPs of the name servers into the respective *Name Server* fields.

Similarly to the procedure described above, YaST also lets you configure the built-in Wi-Fi network adapter.

For detailed information about the network configuration in SUSE Linux Enterprise Server for Arm, consult the respective sections of the *SUSE Linux Enterprise Server Deployment* and *Administration* guides at <https://documentation.suse.com/sles/>.

## 5. General system usage

After the initial configuration procedure and the first boot of the system, you can now use various components of the system.

### 5.1. Desktop

SUSE Linux Enterprise Server for Arm for Raspberry Pi 15 GA had packages for a minimal X11 desktop (IceWM) preinstalled.

SUSE Linux Enterprise Server for Arm for Raspberry Pi 15 SP7 comes as a tiny text-based appliance, allowing you to install any available desktop of your choice or none for headless usage. Should you want to re-create the package selection of the previous images, a pattern `x11_raspberrypi` is provided for convenience.

If not done already, you will need to enable the *Desktop Applications Module*. Assuming you have registered the base product already, as described in the section called “*Registration process*”, the fastest way to enable the module is:

```
#SUSEConnect -p sle-module-desktop-applications/15.7/aarch64
Registering system to SUSE Customer Center

Updating system details on https://scc.suse.com ...

Activating sle-module-desktop-applications 15.7 aarch64 ...
-> Adding service to system ...
-> Installing release package ...

Successfully registered system
```

You can then install the desktop packages pattern:

```
#zypper install -t pattern x11_raspberrypi
```

### 5.2. Bluetooth

Raspberry Pi 3 and later models have a Bluetooth\* controller onboard that can be used for various purposes, such as wireless keyboards, mice or audio devices.

To use the Bluetooth controller, install the `bluez` package:

```
#zypper install bluez
```

Then start and enable the `bluetooth` systemd service:

```
#systemctl enable --now bluetooth
```

You can use **bluetoothctl** to operate the Bluetooth device.

## 6. Product documentation

This introduction only covered the most basic tasks.

### 6.1. Product documentation

You can find the complete documentation for SUSE Linux Enterprise Server for Arm15 SP7 at <https://documentation.suse.com/>.

#### Applicability of product documentation



Not all content in the product documentation applies to SUSE Linux Enterprise Server for Arm on Raspberry Pi, because Raspberry Pi differs significantly from other hardware platforms.

### 6.2. SUSE forums

A valid and activated subscription entitles you to receive bug and security fixes, feature updates, and technical assistance from SUSE's support organization. Learn more at <https://www.suse.com/support/>. Via the SUSE Customer Center at <https://scc.suse.com/>, you can open an incident.

In addition, SUSE provides a community site where you can get answers to questions. Go to <https://community.suse.com/>.

#### Forums are not an official support channel



The SUSE forum is **not an official** support channel from SUSE. The individual articles or statements were contributed voluntarily by community members and users of the respective products. SUSE cannot verify either that the actions described in the articles do what they claim to do, or that they do not have unintended consequences. Therefore, neither SUSE LLC, its affiliates, nor the authors may be held liable for possible errors or the consequences thereof.

Visit <https://www.suse.com/support/> for more information about official support options.

## 7. Trademarks

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.

## 8. Legal notice

Copyright ©2006–2020 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”.

SUSE, the SUSE logo and YaST are registered trademarks of SUSE LLC in the United States and other countries. For SUSE trademarks, see <https://www.suse.com/company/legal/>. Linux is a registered trademark of Linus Torvalds. Other names or trademarks mentioned in this document may be trademarks or registered trademarks of their respective owners.

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.

## 9. 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:

1. 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.
2. 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.
3. State on the Title page the name of the publisher of the Modified Version, as the publisher.
4. Preserve all the copyright notices of the Document.

5. Add an appropriate copyright notice for your modifications adjacent to the other copyright notices.
6. 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.
7. Preserve in that license notice the full lists of Invariant Sections and required Cover Texts given in the Document's license notice.
8. Include an unaltered copy of this License.
9. 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.
10. 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.
11. 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.
12. 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.
13. Delete any section Entitled "Endorsements". Such a section may not be included in the Modified Version.
14. Do not retitle any existing section to be Entitled "Endorsements" or to conflict in title with any Invariant Section.
15. 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.

## Appendix A. GNU licenses

This appendix contains the GNU Free Documentation License version 1.2.

### A.1. 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:

1. 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.
2. 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.
3. State on the Title page the name of the publisher of the Modified Version, as the publisher.
4. Preserve all the copyright notices of the Document.
5. Add an appropriate copyright notice for your modifications adjacent to the other copyright notices.
6. 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.
7. Preserve in that license notice the full lists of Invariant Sections and required Cover Texts given in the Document's license notice.
8. Include an unaltered copy of this License.
9. 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.
10. 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.

11. 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.
12. 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.
13. Delete any section Entitled "Endorsements". Such a section may not be included in the Modified Version.
14. Do not retitle any existing section to be Entitled "Endorsements" or to conflict in title with any Invariant Section.
15. 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.