

Virtualization

WHAT?

By means of virtualization, you can run multiple virtual machines on a single bare-metal host.

WHY?

Sharing host hardware between multiple virtualized guests significantly saves resources.

EFFORT

It takes less than 15 minutes of your time to understand the concept of virtualization.

Publication Date: 27 Nov 2025

Contents

- 1 Introduction to virtualization 2
- 2 Installation of virtualization components 3
- 3 Virtualization modes 4
- 4 Virtualization limits and support 4
- 5 For more information 15
- 6 Legal Notice 16
- A GNU Free Documentation License 16

1 Introduction to virtualization

Virtualization is a technology that provides a way for a machine (VM Host Server) to run another operating system (VM Guest) on top of the host operating system.

1.1 How does virtualization work?

The primary component of VM Host Server that enables virtualization is a *hypervisor*. A hypervisor is a layer of software that runs directly on VM Host Server's hardware. It controls platform resources, sharing them among multiple VM Guests and their operating systems by presenting virtualized hardware interfaces to each VM Guest.

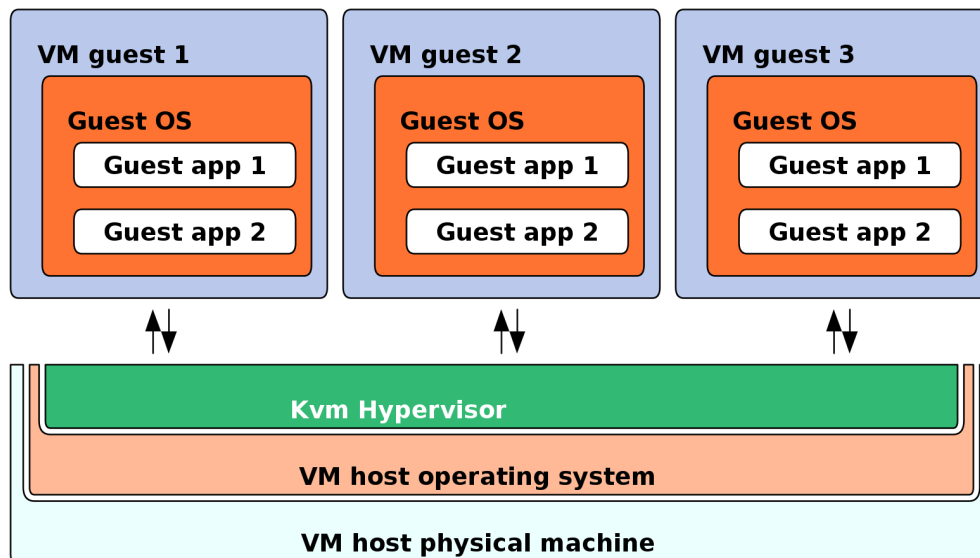


FIGURE 1: A GENERAL SCHEMA OF VIRTUALIZATION

1.2 Benefits of virtualization

Virtualization brings a lot of advantages while providing the same service as a hardware server.

Virtualization reduces the cost of your infrastructure. Servers are mainly used to provide a service to a customer. A virtualized operating system can provide the same service but with the following advantages:

- Less hardware: you can run several operating systems on one host, therefore all hardware maintenance is reduced.
- Less power/cooling: less hardware means you do not need to invest more in electric power, backup power, and cooling if you need more service.
- Save space: your data center space is saved because you do not need more hardware servers (fewer servers than services running).
- Less management: using a VM Guest simplifies the administration of your infrastructure.
- Agility and productivity: virtualization provides *migration* capabilities, *live migration* and *snapshots*. These features reduce downtime and bring an easy way to move your service from one place to another without any service interruption.

2 Installation of virtualization components

To run a virtualization server (VM Host Server) that can host multiple guest systems (VM Guests), you need to install required virtualization components on the server. These components vary depending on which virtualization technology you want to use.

You can install the virtualization tools required to run a VM Host Server either when installing the system (see the [manual installation \(https://documentation.suse.com/sles/16.0/html/SLES-x86-64-agama-based-installation/index.html#agama-manual-installation\)](https://documentation.suse.com/sles/16.0/html/SLES-x86-64-agama-based-installation/index.html#agama-manual-installation)), or from an already installed system by installing a virtualization pattern. The later option is described below:

```
> sudo zypper install -t pattern PATTERN_NAME
```

Replace the *PATTERN_NAME* with one of the following values:

kvm_server

Installs a basic VM Host Server with the KVM and QEMU environments.

kvm_tools

Installs libvirt tools for managing and monitoring VM Guests in the KVM environment.

3 Virtualization modes

Virtualization is a technology that provides a way for a machine (VM Host Server) to run another operating system (VM Guest) on top of the host operating system. There are two basic modes of hosting VM Guests on virtual machines—*full virtualization* mode and *paravirtual* mode.

Full virtualization (FV)

FV lets virtual machines run unmodified operating systems. It uses either Binary Translation or hardware-assisted virtualization technology, such as AMD* Virtualization or Intel* Virtualization Technology, to improve performance on processors that support it. In FV mode, VM Guest is also called the *Hardware Virtual Machine* (HVM).



Tip

Certain guest operating systems hosted in full virtualization mode can be configured to use drivers from the SUSE Virtual Machine Drivers Pack (VMDP) instead of drivers included in the operating system. Running virtual machine drivers improves performance on guest operating systems, such as Windows Server 2003.

Paravirtualization (PV)

PV normally requires that guest operating systems are modified for the virtualization environment. VM Guests running in paravirtual mode have better performance than those running under full virtualization. Operating systems currently modified to run in paravirtual mode are called *paravirtualized* operating systems and include SLES.

PV on HVM (PVHVM)

PVHVM enhances HVM (see [Full virtualization \(FV\)](#)) with paravirtualized drivers, and handling of paravirtualized interrupts and timers.

4 Virtualization limits and support

Virtualization limits and support

Important

QEMU is only supported when used for virtualization together with the KVM hypervisors. The TCG accelerator is not supported, even when it is distributed within SUSE products. Users must not rely on QEMU TCG to provide guest isolation, or for any security guarantees. See also <https://qemu-project.gitlab.io/qemu/system/security.html>.

4.1 Architecture support

4.1.1 KVM hardware requirements

SUSE supports KVM full virtualization on AMD64/Intel 64, AArch64, IBM Z and IBM LinuxONE hosts.

- On the AMD64/Intel 64 architecture, KVM is designed around hardware virtualization features included in AMD* (AMD-V) and Intel* (VT-x) CPUs. It supports virtualization features of chipsets and PCI devices, such as an I/O Memory Mapping Unit (IOMMU) and Single Root I/O Virtualization (SR-IOV). You can test whether your CPU supports hardware virtualization with the following command:

```
> egrep '(vmx|svm)' /proc/cpuinfo
```

If this command returns no output, your processor either does not support hardware virtualization, or this feature has been disabled in the BIOS or firmware.

The following Web sites identify AMD64/Intel 64 processors that support hardware virtualization: <https://ark.intel.com/Products/VirtualizationTechnology> (for Intel CPUs), and <https://products.amd.com/> (for AMD CPUs).

- On the Arm architecture, Armv8-A processors include support for virtualization.
- On the Arm architecture, we only support running QEMU/KVM via the CPU model host (it is named host-passthrough in Virtual Machine Manager or libvirt).

Note: KVM kernel modules not loading

The KVM kernel modules only load if the CPU hardware virtualization features are available.

However, additional RAM for each virtualized guest is needed. It should at least be the same amount that is needed for a physical installation. It is also strongly recommended having at least one processor core or hyper-thread for each running guest.



Note: AArch64


AArch64 is a continuously evolving platform. It does not have a traditional standards and compliance certification program to enable interoperability with operating systems and hypervisors. Ask your vendor for the support statement on SUSE Linux Enterprise Server.




Note: POWER

Running KVM hypervisor on the POWER platform is not supported.

4.2 Hypervisor limits

New features and virtualization limits for KVM are outlined in the [Release Notes \(https://www.suse.com/releasenotes/\)](https://www.suse.com/releasenotes/)  for each Service Pack (SP).

Only packages that are part of the official repositories for SUSE Linux Enterprise Server are supported. Conversely, all optional subpackages and plug-ins (for QEMU, `libvirt`) provided at [packagehub \(https://packagehub.suse.com/\)](https://packagehub.suse.com/)  are not supported.

For the maximum total virtual CPUs per host. The total number of virtual CPUs should be proportional to the number of available physical CPUs.

4.2.1 KVM limits

Supported (and tested) virtualization limits of a SUSE Linux Enterprise Server 16.0 host running Linux guests on AMD64/Intel 64. For other operating systems, refer to the specific vendor.

TABLE 1: KVM VM LIMITS

Maximum virtual CPUs per VM	768
Maximum memory per VM	4 TiB



Note

KVM host limits are identical to SUSE Linux Enterprise Server (see the corresponding section of release notes), except for:

- *Maximum virtual CPUs per VM*: see recommendations in the *Virtualization Best Practices Guide* regarding the overcommitment of physical CPUs. The total virtual CPUs should be proportional to the available physical CPUs.

4.3 Guest VM restrictions and limitations (secure VM)

Please be aware of the following functionalities and features that are not available or have limitations for guest VMs, especially when deployed within secure VM environments. These limitations are crucial for maintaining the enhanced security provided by the underlying hardware and software configurations.

- **Secure Boot (AMD side)**: Secure Boot functionality is not supported on the AMD platform for guest VMs within this secure environment. This means that guest VMs cannot leverage the UEFI Secure Boot mechanism to verify the digital signatures of boot components, which typically helps prevent the loading of unauthorized or malicious software during the boot process. Users should consider alternative methods for ensuring software integrity post-boot.
- **VM migration**: The live migration of virtual machines between hosts is currently not supported. This implies that planned maintenance, load balancing, or disaster recovery scenarios requiring VM movement without downtime will need to involve a full shutdown and restart of the guest VM on the new host. This limitation is often a consequence of maintaining the cryptographic isolation and attestation state of secure VMs.
- **Suspend/restore**: The ability to suspend a VM's execution state to disk and later restore it is not available. This impacts operational flexibility, as VMs cannot be paused and resumed seamlessly. Any interruption to a guest VM's operation will require a full shutdown and a fresh boot cycle, losing the immediate operational state.
- **Pass-through devices**: Direct pass-through of host devices (such as GPUs, network cards, or storage controllers) to the guest VM is not supported. This limitation restricts scenarios where guest VMs require exclusive, high-performance access to specific hardware com-

ponents. Workloads that heavily rely on direct hardware interaction, like certain graphical applications or specialized I/O operations, may experience reduced performance or incompatibility.

- **VM reboot:** The internal reboot functionality for guest VMs is not supported. If a guest VM requires a restart, it must be fully shut down and then started again from the host management interface. This ensures that the secure state of the VM is properly re-established upon each boot, rather than relying on an internal reset that might bypass certain security checks.
- **Memory ballooning:** Memory ballooning, which allows dynamic adjustment of VM memory by reclaiming unused guest memory back to the host, is not supported. This means that the allocated memory for a guest VM will remain fixed, regardless of its actual usage. Consequently, memory overcommitment strategies, where the sum of allocated VM memory exceeds the physical host memory, cannot be effectively utilized, potentially leading to less efficient memory utilization on the host.
- **Hotplug CPU/memory:** The hotplugging (adding or removing) of CPU cores or memory modules while the VM is running is not supported. Any changes to the vCPU or memory configuration of a guest VM will require a full shutdown and a restart of the VM for the changes to take effect. This affects the agility and flexibility in dynamically scaling resources for running workloads.
- **Virtio graphics:** Only Virtio block devices (for storage) and network devices are supported. Virtio graphics are not available for guest VMs in this environment. This implies that guest VMs will rely on basic graphics emulation, which may not provide optimal performance for graphically intensive applications, user interfaces, or remote desktop protocols requiring accelerated graphics.
- **Huge pages:** The use of huge pages for memory allocation within the guest VM is not supported. Huge pages can improve performance by reducing Translation Lookaside Buffer (TLB) misses, especially for applications with large memory footprints. Without huge page support, memory management might incur slightly higher overhead, which could subtly impact the performance of memory-intensive applications.
- **vCPU limit (AMD SNP):** The number of virtual CPUs (vCPUs) that can be assigned to a guest VM is limited to 255 when utilizing AMD Secure Nested Paging (SNP). This specific limitation is imposed by the AMD SNP architecture to maintain the integrity and performance characteristics of the secure execution environment. Workloads requiring more than 255 vCPUs cannot be deployed on these secure VMs.



4.4 Supported host environments (hypervisors)

This section describes the support status of SUSE Linux Enterprise Server 16.0 running as a guest operating system on top of different virtualization hosts (hypervisors).

TABLE 2: THE FOLLOWING SUSE HOST ENVIRONMENTS ARE SUPPORTED:


SUSE Linux Enterprise Server	Hypervisors
SUSE Linux Enterprise Server 12 SP5	KVM (SUSE Linux Enterprise Server 15 SP6 guest must use UEFI boot)
SUSE Linux Enterprise Server 15 SP3 to SP7	KVM

THE FOLLOWING THIRD-PARTY HOST ENVIRONMENTS ARE SUPPORTED:

- Nutanix Acropolis Hypervisor with AOS (<https://portal.nutanix.com/page/documents/compatibility-matrix/guestos>) 
- Oracle VM Server 3.X (<https://www.oracle.com/fr/virtualization/virtualbox/>) 
- Oracle Linux KVM 7, 8 (<https://www.oracle.com/linux/>) 
- VMware ESXi 6.7, 7.0, 8.0 (<https://www.vmware.com/products/esxi-and-esx.html>) 
- Windows Server 2019, 2022, 2025

You can also search in the SUSE YES certification database (<https://www.suse.com/yessearch/Search.jsp>) .


THE LEVEL OF SUPPORT IS AS FOLLOWS:

- Support for SUSE host operating systems is full L3 (both for the guest and host), according to the respective [product lifecycle](https://www.suse.com/lifecycle/) (<https://www.suse.com/lifecycle/>) .
- SUSE provides full L3 support for SUSE Linux Enterprise Server guests within third-party host environments.
- Support for the host and cooperation with SUSE Linux Enterprise Server guests must be provided by the host system's vendor.

4.5 Supported guest operating systems

This section lists the support status for guest operating systems virtualized on top of SUSE Linux Enterprise Server 16.0 for KVM hypervisors.

Important

Microsoft Windows guests can be rebooted by `libvirt/virsh` only if paravirtualized drivers are installed in the guest. Refer to <https://www.suse.com/products/vmdriverpack/>  for more details on downloading and installing PV drivers.


THE FOLLOWING GUEST OPERATING SYSTEMS ARE FULLY SUPPORTED (L3):

- SUSE Linux Enterprise Server 12 SP5
- SUSE Linux Enterprise Server 15 SP3, 15 SP4, 15 SP5, 15 SP6, 15 SP6
- SUSE Linux Enterprise Micro 6.0, 6.1, 6.2
- Windows Server 2022, 2025
- Oracle Linux 7, 8


THE FOLLOWING GUEST OPERATING SYSTEMS ARE SUPPORTED AS A TECHNOLOGY PREVIEW (L2, FIXES IF REASONABLE):

- SLED 15 SP3
- Windows 10 / 11

RED HAT AND CENTOS GUEST OPERATING SYSTEMS ARE FULLY SUPPORTED (L3) IF THE CUSTOMER HAS PURCHASED SUSE MULTI-LINUX SUPPORT.

- Refer to the SUSE Multi-Linux Support documentation at <https://documentation.suse.com/liberty>  for the list of available combinations and supported releases. In other cases, they are supported on a limited basis (L2, fixes if reasonable).

ALL OTHER GUEST OPERATING SYSTEMS

- In other combinations, L2 support is provided but fixes are available only if feasible. SUSE fully supports the host OS (hypervisor). The guest OS issues need to be supported by the respective OS vendor. If an issue fix involves both the host and guest environments, the customer needs to approach both SUSE and the guest VM OS vendor.
- All guest operating systems are supported both fully virtualized and paravirtualized. The exception is Windows systems, which are only supported fully virtualized (but they can use PV drivers: <https://www.suse.com/products/vmdriverpack/> ) , and OES operating systems, which are supported only paravirtualized.
- All guest operating systems are supported both in 32-bit and 64-bit environments, unless stated otherwise.

4.5.1 Availability of paravirtualized drivers

To improve the performance of the guest operating system, paravirtualized drivers are provided when available. Although they are not required, it is strongly recommended to use them.

The paravirtualized drivers are available as follows:

Red Hat

Available since Red Hat Enterprise Linux 5.4. Starting from Red Hat Enterprise Linux 7.2, Red Hat removed the PV drivers.

Windows

SUSE has developed Virtio-based drivers for Windows, which are available in the Virtual Machine Driver Pack (VMDP). For more information, see <https://www.suse.com/products/vmdriverpack/>.

4.6 Supported VM migration scenarios

SUSE Linux Enterprise Server supports migrating a virtual machine from one physical host to another.

4.6.1 Offline migration scenarios

SUSE supports offline migration, powering off a guest VM, then moving it to a host running a different SLE product, from SLE 12 to SLE 15 SPX. The following host operating system combinations are fully supported (L3) for migrating guests from one host to another:

TABLE 3: SUPPORTED OFFLINE MIGRATION GUESTS

Target SLES host	12	12	12	15	15	15	15	15	15	15
Source SLES host	SP3	SP4	SP5	GA	SP1	SP2	SP3	SP4	SP5	SP6
12 SP3	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗
12 SP4	✗	✓	✓	✓ ¹	✓	✗	✗	✗	✗	✗
12 SP5	✗	✗	✓	✗	✓	✓	✗	✗	✗	✗
15 GA	✗	✗	✗	✗	✓	✓	✓	✗	✗	✗

Target SLES host	12	12	12	15	15	15	15	15	15	15
Source SLES host	SP3	SP4	SP5	GA	SP1	SP2	SP3	SP4	SP5	SP6
15 SP1	✗	✗	✗	✗	✓	✓	✓	✗	✗	✗
15 SP2	✗	✗	✗	✗	✗	✓	✓	✓	✗	✗
15 SP3	✗	✗	✗	✗	✗	✗	✓	✓	✓	✓
15 SP4	✗	✗	✗	✗	✗	✗	✗	✓	✓	✓
15 SP5	✗	✗	✗	✗	✗	✗	✗	✗	✓	✓
15 SP6	✗	✗	✗	✗	✗	✗	✗	✗	✗	✓

✓ Fully compatible and fully supported

✓¹ Supported for KVM hypervisor only

✗ Not supported

4.6.2 Live migration scenarios

This section lists support status of live migration scenarios when running virtualized on top of SLES. The following host operating system combinations are fully supported (L3 according to the respective [product life cycle \(https://www.suse.com/lifecycle\)](https://www.suse.com/lifecycle)).



Note: Live migration

- SUSE always supports live migration of virtual machines between hosts running SLES with successive service pack numbers. For example, from SLES 16 to 16.1.
- SUSE strives to support live migration of virtual machines from a host running a service pack under LTSS to a host running a newer service pack, within the same major version of SUSE Linux Enterprise Server. SUSE only performs minimal testing of LTSS-to-newer migration scenarios and recommends thorough on-site testing before attempting to migrate critical virtual machines.

! Important: Confidential Computing

SLES 15 SP6 includes kernel patches and tooling to enable Intel TDX Confidential Computing technology in the product. As this technology is not yet fully ready for a production environment, it is provided as a technology preview.

TABLE 4: SUPPORTED LIVE MIGRATION GUESTS

Source SLES host	Target SLES host	15 SP7	16
15 SP7		✓	✗
16		✗	✓ ²

✓ Fully compatible and fully supported

✓² When available

✗ Not supported

4.7 Feature support

! Important: Nested virtualization: technology preview

Nested virtualization allows you to run a virtual machine inside another VM while still using hardware acceleration from the host. It has low performance and adds more complexity while debugging. Nested virtualization is normally used for testing purposes. In SUSE Linux Enterprise Server, nested virtualization is a technology preview. It is only provided for testing and is not supported. Bugs can be reported, but they are treated with low priority. Any attempt to live migrate or to save or restore VMs in the presence of nested virtualization is also explicitly unsupported.

! Important: Post-copy live migration: technology preview

Post-copy is a method to live migrate virtual machines that is intended to get VMs running as soon as possible on the destination host, and have the VM RAM transferred gradually in the background over time as needed. Under certain conditions, this can be an

optimization compared to the traditional pre-copy method. However, this comes with a major drawback: An error occurring during the migration (especially a network failure) can cause the whole VM RAM contents to be lost. Therefore, we recommend using pre-copy only in production, while post-copy can be used for testing and experimentation in case losing the VM state is not a major concern.

4.7.1 Guest feature support



Note: Windows guest

Hotplugging of virtual network and virtual block devices, and resizing, shrinking and restoring dynamic virtual memory are supported in KVM only when PV drivers are used (VMDP (<https://www.suse.com/products/vmdriverpack/>)).



Note: Intel FlexMigration

For machines that support Intel FlexMigration, CPU-ID masking and faulting allow for more flexibility in cross-CPU migration.



Tip

For KVM, a detailed description of supported limits, features, recommended settings and scenarios, and other useful information is maintained in `kvm-supported.txt`. This file is part of the KVM package and can be found in `/usr/share/doc/packages/qemu-kvm`.

TABLE 5: GUEST FEATURE SUPPORT FOR KVM

Features	KVM FV guest		
Virtual network and virtual block device hotplugging	✓		
Virtual CPU Hotplug	✗		
Virtual CPU Overcommit	✓		
Dynamic virtual memory resize	✓		

Features	KVM FV guest		
VM save and restore	✓		
VM Live Migration	✓		
VM snapshot	✓		
Advanced debugging with GDB	✓		
Memory ballooning	✗		
PCI Pass-Through	✓		
AMD SEV and SEV-SNP	✓ [3]		

✓ Fully compatible and fully supported

✗ Not supported

[1] NetWare guests are excluded.

[2] See <https://documentation.suse.com/sles/html/SLES-amd-sev/article-amd-sev.html>.

5 For more information


For further steps in virtualization, refer to the following sources:

- Managing virtual machines with libvirt (<https://documentation.suse.com/smart/virtualization-cloud/html/concept-manage-virtual-machines-libvirt/concept-manage-virtual-machines-libvirt.html>)
- Configuring Virtual Machines with Virtual Machine Manager (<https://documentation.suse.com/smart/virtualization-cloud/html/task-configure-virtual-machine-manager/task-configure-virtual-machine-manager.html>)
- Assigning Host Devices to Virtual Machines (<https://documentation.suse.com/smart/virtualization-cloud/html/vm-assign-pci-device/vm-assign-pci-device.html>)
- Configuring a Virtual Disk Cache Mode (<https://documentation.suse.com/smart/virtualization-cloud/html/virtual-disk-cache-mode-configure/virtual-disk-cache-mode-configure.html>)

6 Legal Notice

Copyright© 2006–2025 SUSE LLC and contributors. All rights reserved.

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

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

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

A GNU Free Documentation License

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

0. PREAMBLE

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

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

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

1. APPLICABILITY AND DEFINITIONS

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

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

A "Secondary Section" is a named appendix or a front-matter section of the Document that deals exclusively with the relationship of the publishers or authors of the Document to the Document's overall subject (or to related matters) and contains nothing that could fall directly within that overall subject. (Thus, if the Document is in part a textbook of mathematics, a Secondary Section may not explain any mathematics.) The relationship could be a matter of historical connection with the subject or with related matters, or of legal, commercial, philosophical, ethical or political position regarding them.

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

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

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

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

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

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

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

2. VERBATIM COPYING

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

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

3. COPYING IN QUANTITY

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

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

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

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

4. MODIFICATIONS

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

- A. Use in the Title Page (and on the covers, if any) a title distinct from that of the Document, and from those of previous versions (which should, if there were any, be listed in the History section of the Document). You may use the same title as a previous version if the original publisher of that version gives permission.
- B. List on the Title Page, as authors, one or more persons or entities responsible for authorship of the modifications in the Modified Version, together with at least five of the principal authors of the Document (all of its principal authors, if it has fewer than five), unless they release you from this requirement.
- C. State on the Title page the name of the publisher of the Modified Version, as the publisher.
- D. Preserve all the copyright notices of the Document.
- E. Add an appropriate copyright notice for your modifications adjacent to the other copyright notices.
- F. Include, immediately after the copyright notices, a license notice giving the public permission to use the Modified Version under the terms of this License, in the form shown in the Addendum below.
- G. Preserve in that license notice the full lists of Invariant Sections and required Cover Texts given in the Document's license notice.
- H. Include an unaltered copy of this License.
- I. Preserve the section Entitled "History", Preserve its Title, and add to it an item stating at least the title, year, new authors, and publisher of the Modified Version as given on the Title Page. If there is no section Entitled "History" in the Document, create one stating the title, year, authors, and publisher of the Document as given on its Title Page, then add an item describing the Modified Version as stated in the previous sentence.

- J. Preserve the network location, if any, given in the Document for public access to a Transparent copy of the Document, and likewise the network locations given in the Document for previous versions it was based on. These may be placed in the "History" section. You may omit a network location for a work that was published at least four years before the Document itself, or if the original publisher of the version it refers to gives permission.
- K. For any section Entitled "Acknowledgements" or "Dedications", Preserve the Title of the section, and preserve in the section all the substance and tone of each of the contributor acknowledgements and/or dedications given therein.
- L. Preserve all the Invariant Sections of the Document, unaltered in their text and in their titles. Section numbers or the equivalent are not considered part of the section titles.
- M. Delete any section Entitled "Endorsements". Such a section may not be included in the Modified Version.
- N. Do not retitle any existing section to be Entitled "Endorsements" or to conflict in title with any Invariant Section.
- O. Preserve any Warranty Disclaimers.

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

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

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

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

5. COMBINING DOCUMENTS

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

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

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

6. COLLECTIONS OF DOCUMENTS

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

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

7. AGGREGATION WITH INDEPENDENT WORKS

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

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

8. TRANSLATION


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

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

9. TERMINATION

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

10. FUTURE REVISIONS OF THIS LICENSE

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

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

ADDENDUM: How to use this License for your documents

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

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

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

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

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