Configuring GPU pass-through for NVIDIA Cards

GPU pass-through assigns a physical NVIDIA* GPU directly to a virtual machine, giving it near-native graphics and compute performance instead of a shared, virtualized device. Use it for machine learning, AI workloads, 3D rendering, and other compute-intensive tasks that need direct hardware access. After configuring pass-through on the host and installing the guest driver, the virtual machine can use the GPU as if it were installed locally.

Prerequisites

- GPU pass-through is supported on the AMD64/Intel 64 architecture only.
- The host operating system needs to be SUSE Linux Enterprise Server 16 or newer.
- This article covers GPU pass-through for NVIDIA Tesla, Data Center, and RTX series cards, for GPU computation purposes only.
- For optimal CUDA compute performance with System-Allocated Memory (SAM) migration, kernel 6.12 or newer is recommended to include required HMM fixes.
- To manage the host system after configuring GPU pass-through, you need either an additional display card on the host or a functional SSH connection.

Publication Date: 05 Nov 2026

Contents

1 Configuring the host 3

2 Configuring the guest 7

3 For more information 11

4 Legal Notice 11

A GNU Free Documentation License 12

Acronyms 19

Glossary 25

1 Configuring the host

Complete the following steps on the host to prepare an NVIDIA GPU for pass-through to a virtualized guest.

1.1 Verify the host environment

1. Verify that the host operating system is SUSE Linux Enterprise Server 16 or newer:

```
> cat /etc/issue

Welcome to SUSE Linux Enterprise Server 16.1  (x86_64) - Kernel \r (\l).
```

2. Verify that the host supports VT-d technology and that it is already enabled in the firmware settings:

```
> dmesg | grep -e "Directed I/O"

[   12.819760] DMAR: Intel(R) Virtualization Technology for Directed I/O
```

   a. If VT-d is not enabled in the firmware, enable it.

   b. Reboot the host.

3. Verify that the host has an extra GPU or VGA card:

   - Check for a generic VGA-compatible controller:

```
> lspci | grep -i "vga"

07:00.0 VGA compatible controller: Matrox Electronics Systems Ltd. \
  MGA G200e [Pilot] ServerEngines (SEP1) (rev 05)
```

   - Check for an NVIDIA card directly, for example with a Tesla V100 card:

```
> lspci | grep -i nvidia

03:00.0 3D controller: NVIDIA Corporation GV100 [Tesla V100 PCIe] (rev a1)
```

The output varies by card model, for example:

- Tesla V100: `NVIDIA Corporation GV100 [Tesla V100 PCIe]`
- A100: `NVIDIA Corporation A100-SXM4-40GB`
- L4: `NVIDIA Corporation L4`

1.2 Enable IOMMU

IOMMU is disabled by default. You need to enable it at boot time in the `/etc/default/grub` configuration file.

1. For Intel-based hosts:

   ```
   GRUB_CMDLINE_LINUX="intel_iommu=on iommu=pt rd.driver.pre=vfio-pci"
   ```

   For AMD-based hosts:

   ```
   GRUB_CMDLINE_LINUX="iommu=pt amd_iommu=on rd.driver.pre=vfio-pci"
   ```

2. When you save the modified `/etc/default/grub` file, re-generate the main GRUB 2 configuration file `/boot/grub2/grub.cfg`:

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

3. Reboot the host and verify that IOMMU is enabled:

   ```
   > dmesg |  grep -e DMAR -e IOMMU
   ```

1.3 Blacklist the Nouveau driver

To assign the NVIDIA card to a VM guest, prevent the host OS from loading the built-in `nouveau` driver for NVIDIA GPUs. Create the file `/etc/modprobe.d/60-blacklist-nouveau.conf` with the following content:

```
blacklist nouveau
```

1.4 Configure VFIO and isolate the GPU used for pass-through

1. Find the card vendor and model IDs. Use the bus number identified in the host verification step, for example, `03:00.0`:

```
> lspci -nn | grep 03:00.0

03:00.0 3D controller [0302]: NVIDIA Corporation GV100 [Tesla V100 PCIe] [10de:1db4]
 (rev a1)
```

2. Create the file `/etc/modprobe.d/vfio.conf` with the following content:

```
options vfio-pci ids=10de:1db4
```



Note: Cards requiring an additional device ID

Some cards expose a separate audio function on the same bus and slot, for example, `03:00.1`. Check for one with **`lspci -nn | grep 03:00`**: if an audio device shows up alongside the GPU, add its ID to the `ids=` parameter too, otherwise you cannot use the card.

1.5 Load the VFIO driver

There are three ways you can load the VFIO driver.



Note: IOMMUFD in kernel 6.12

Kernel 6.12 introduces IOMMUFD as the new user API for VFIO. The legacy `vfio_iommu_type1` container/group model is still supported but gradually being superseded. For libvirt-based passthrough (virt-manager, virt-install), no changes are required as libvirt abstracts the VFIO interface.

1.5.1 Including the driver in the initrd file

1. Create the file `/etc/dracut.conf.d/gpu-passthrough.conf` with the following content (mind the leading whitespace):

```
add_drivers+=" vfio vfio_iommu_type1 vfio_pci vfio_virqfd"
```

2. Re-generate the initrd file:

```
> sudo dracut --force /boot/initrd $(uname -r)
```

1.5.2 Adding the driver to the list of auto-loaded modules

Create the file `/etc/modules-load.d/vfio-pci.conf` and add the following content:

```
vfio
vfio_iommu_type1
vfio_pci
kvm
kvm_intel
```

1.5.3 Loading the driver manually

To load the driver manually at runtime, execute the following command:

```
> sudo modprobe vfio-pci
```

1.6 Disable MSR for Microsoft Windows guests

For Microsoft Windows guests, we recommend disabling MSR (model-specific register) to avoid the guest crashing. Create the file `/etc/modprobe.d/kvm.conf` and add the following content:

```
options kvm ignore_msrs=1
```

1.7 Install UEFI firmware

For proper GPU pass-through functionality, the host needs to boot using UEFI firmware, not a legacy-style BIOS boot sequence. Install the `qemu-ovmf` package if not already installed:

```
> sudo zypper install qemu-ovmf
```

1.8 Reboot the host machine

For most of the changes in the above steps to take effect, you need to reboot the host machine:

```
> sudo shutdown -r now
```

2 Configuring the guest

This section describes how to configure the guest virtual machine so that it can use the host's NVIDIA GPU. Use virt-manager or **`virt-install`** to install the guest VM. Find more details in Guest installation (https://documentation.suse.com/sles/16.1/html/SLES-virtualization-libvirt/index.html#libvirt-guest-installation).

2.1 Requirements for the guest configuration

During the guest VM installation, select *Customize configuration before install* and configure the following devices:

- Use Q35 chipset for proper PCIe passthrough with modern GPUs.
- Install the guest VM using UEFI firmware.
- Add the following emulated devices:
  - Graphic: VNC
  - Device: virtio-gpu, VGA or Virtio

  Find more information in 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).
- Add the host PCI device (for example, `03:00.0`) to the guest. Find more information in Assigning Host Devices to Virtual Machines (https://documentation.suse.com/smart/virtualization-cloud/html/vm-assign-pci-device/vm-assign-pci-device.html).
- For the best performance, we recommend using virtio drivers for the network card and storage.

2.2 Install the graphic card driver

2.2.1 RPM-based Linux distributions

1. Download the driver RPM package from the official NVIDIA driver download page at https://www.nvidia.com/Download/index.aspx.
2. Install the downloaded RPM package:

```
> sudo rpm -i nvidia-diag-driver-local-repo-sles*.rpm
```

3. Refresh repositories and install `cuda-drivers`:

```
> sudo zypper refresh && zypper install cuda-drivers
```

4. Reboot the guest VM:

```
> sudo shutdown -r now
```



Note: Minimum driver version for CUDA

CUDA 13.3 requires driver version 610.43.02 or newer. The minimum driver version for CUDA 13.x is 580.

2.2.2 Generic Linux installer

1. Because the installer needs to compile the NVIDIA driver modules, install the `gcc-c++` and `kernel-devel` packages.
2. NVIDIA driver modules are unsigned, so disable Secure Boot on the guest. Either disable it in the UEFI firmware setup, or use **`mokutil`** to enroll a Machine Owner Key. Find more information about UEFI Secure Boot in Using UEFI Secure Boot (https://documentation.suse.com/sles/16.1/html/SLES-virtualization-libvirt/index.html#vle-libvirt-inst-virt-install-ovmf).
3. Download the driver installation script from https://www.nvidia.com/Download/index.aspx?lang=en-us and make it executable:

```
> chmod +x NVIDIA-Linux-x86_64-*.run
```

4. Run the script to complete the driver installation:

```
> sudo ./NVIDIA-Linux-x86_64-*.run
```

5. Download CUDA drivers from https://developer.nvidia.com/cuda-downloads and install following the on-screen instructions.

2.2.3 Testing the Linux driver installation

Regardless of which installation method you used, verify the driver installation as follows.



Note: Display issues

After you have installed the NVIDIA drivers, the virt-manager display loses its connection to the guest OS. To access the guest VM, you must either login via **ssh**, change to the console interface, or install a dedicated VNC server in the guest. To avoid a flickering screen, stop and disable the display manager:

```
> sudo systemctl stop display-manager && systemctl disable display-manager
```

1. Verify the installation with the NVIDIA System Management Interface:

```
> nvidia-smi
```

2. Change to the directory of a CUDA sample. For example, the deviceQuery sample:

```
> cd /usr/local/cuda/samples/1_Utilities/deviceQuery
```

3. Compile the sample:

```
> sudo make
```

4. Run the compiled sample:

```
> ./deviceQuery
```

2.2.4 Microsoft Windows guest

Important: Hide the hypervisor from the guest drivers

Before you install the NVIDIA drivers, you need to hide the hypervisor from the drivers by using the `<hidden state='on'/>` directive in the guest's `libvirt` definition, for example:

```
<features>
 <acpi/>
 <apic/>
 <kvm>
  <hidden state='on'/>
 </kvm>
</features>
```

1. Download and install the NVIDIA driver from https://www.nvidia.com/Download/index.aspx.
2. Download and install the CUDA toolkit from https://developer.nvidia.com/cuda-downloads?target_os=Windows&target_arch=x86_64.
3. Find several NVIDIA demo samples in the directory `Program Files\NVIDIA GPU Computing Toolkit\CUDA\v12.x\extras\demo_suite` on the guest.

Note: Microsoft Windows driver no longer bundled with the CUDA Toolkit

Starting with CUDA 13.1, the Microsoft Windows display driver is no longer bundled with the CUDA Toolkit. You must download and install the appropriate NVIDIA driver separately from the official NVIDIA driver download page.

3 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)
- Converting VMware Virtual Machines to KVM with virt-v2v (#virt-v2v-vmware-migration)

4 Legal Notice

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

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

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

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

A GNU Free Documentation License

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

0. PREAMBLE

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

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

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

1. APPLICABILITY AND DEFINITIONS

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

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

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

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

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

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

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

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

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

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

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

2. VERBATIM COPYING

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

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

3. COPYING IN QUANTITY

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

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

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

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

4. MODIFICATIONS

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

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

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

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

D. Preserve all the copyright notices of the Document.

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

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

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

H. Include an unaltered copy of this License.

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

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

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

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

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

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

O. Preserve any Warranty Disclaimers.

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

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

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

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

5. COMBINING DOCUMENTS

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

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

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

6. COLLECTIONS OF DOCUMENTS

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

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

7. AGGREGATION WITH INDEPENDENT WORKS

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

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

8. TRANSLATION

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

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

9. TERMINATION

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

10. FUTURE REVISIONS OF THIS LICENSE

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

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

ADDENDUM: How to use this License for your documents

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

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

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

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

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

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

Acronyms

ACPI

Advanced Configuration and Power Interface

An open standard that provides uniform device configuration and power management capabilities directly to the operating system.

AER

Advanced Error Reporting

A capability within the PCI Express specification that enables advanced reporting, logging and recovery from PCI bus errors.

APIC

Advanced Programmable Interrupt Controller

An architecture of interrupt controllers designed to efficiently manage and route hardware interrupts across multiprocessing systems.

BDF

Bus:Device:Function

A standard notation used to uniquely address and identify individual physical or virtual devices on a PCI/PCIe bus.

CG

Control Groups (cgroups)

A Linux kernel feature used to limit, account for, isolate and prioritize the resource usage (such as CPU, memory and disk I/O) of collections of processes.

EDF

Earliest Deadline First

A dynamic priority scheduling algorithm that provides intuitive, real-time CPU resource sharing based on strict time guarantees.

EPT

Extended Page Tables

Performance in a virtualized environment is close to that in a native environment. However, virtualization does create overhead. It comes from the virtualization of the CPU, the *MMU*, and the I/O devices. In some recent x86 processors, AMD and Intel have begun to provide hardware extensions to help bridge this performance gap. In 2006, both vendors introduced their first-generation hardware support for x86 virtualization with AMD-Virtualization (AMD-V) and Intel® VT-x technologies. Intel introduced its second generation of hardware support that incorporates MMU-virtualization, called Extended Page Tables (EPT). EPT-enabled systems can improve performance compared to using shadow paging for *MMU* virtualization. EPT increases memory access latencies for a few workloads. This cost can be reduced by effectively using large pages in the guest and the hypervisor.

HVM

Hardware Virtual Machine

A fully virtualized operating system instance running on hardware that natively supports virtualization extensions (such as Intel VT-x or AMD-V).

IOMMU

Input/Output Memory Management Unit

IOMMU (AMD* technology) is a memory management unit (*MMU*) that connects a direct memory access-capable (DMA-capable) I/O bus to the main memory.

KSM

Kernel Same-page Merging

A Linux kernel feature that allows for automatic sharing of identical memory pages between guests to save host memory. KVM is optimized to use KSM if enabled on the VM Host Server.

MMU

Memory Management Unit

A hardware component responsible for handling accesses to memory requested by the CPU. Its functions include translation of virtual addresses to physical addresses (that is, virtual memory management), memory protection, cache control, bus arbitration and, in simpler computer architectures (especially 8-bit systems), bank switching.

PAE

Physical Address Extension

A memory management feature that allows 32-bit x86 operating systems to address more than 4 GB of physical memory by expanding page table entries to 64 bits.

PCID

Process-Context Identifiers

A processor facility that caches information for multiple linear-address spaces so that the processor may retain cached information when software switches to a different linear address space. The INVPCID instruction is used for fine-grained *TLB* flush, which benefits the kernel.

PCIe

Peripheral Component Interconnect Express

A high-speed serial expansion bus standard designed to replace older PCI, PCI-X and AGP bus standards. PCIe has numerous improvements, including a higher maximum system bus throughput, a lower I/O pin count and a smaller physical footprint. Moreover, it also has a more detailed error detection and reporting mechanism (*AER*) and a native hotplug functionality. It is also backward compatible with PCI.

PSE / PSE36

Page Size Extension / Page Size Extension 36-Bit

Hardware features in x86 processors that allow for pages larger than the traditional 4 KiB size. PSE-36 capability offers 4 more bits, in addition to the normal 10 bits, which are used inside a page directory entry pointing to a large page. This allows a large page to be located in a 36-bit address space.

PT

Page Table

The data structure used by a virtual memory system in a computer operating system to store the mapping between virtual addresses and physical addresses. Virtual addresses are those unique to the accessing process. Physical addresses are those unique to the hardware (RAM).

QXL

QXL Display Device

QXL is a cirrus VGA framebuffer (8M) driver for virtualized environment.

RVI / NPT

Rapid Virtualization Indexing / Nested Page Tables

AMD's second-generation hardware-assisted virtualization technology for the processor memory management unit (*MMU*).

SATA

Serial ATA

A computer bus interface that connects host bus adapters to mass storage devices, such as hard disks and optical drives.

Seccomp2-based sandboxing

Secure Computing Mode sandboxing

A sandboxed environment where only predetermined system calls are permitted for added protection against malicious behavior.

SPICE

Simple Protocol for Independent Computing Environments

An open-source remote display protocol designed for virtual environments to deliver high-quality graphics and audio streaming over networks.

TCG

Tiny Code Generator

The core component of the QEMU emulator that acts as a just-in-time compiler, translating target CPU instructions into host system architecture instructions.

THP

Transparent Huge Pages

A Linux kernel feature that allows CPUs to address memory using pages larger than the default 4 KB. This helps reduce memory consumption and CPU cache usage. KVM is optimized to use THP (via madvise and opportunistic methods) if enabled on the VM Host Server.

TLB

Translation Lookaside Buffer

A hardware cache that memory management hardware uses to improve virtual address translation speed. All current desktop, notebook and server processors use a TLB to map virtual and physical address spaces, and it is nearly always present in any hardware that uses virtual memory.

VCPU

Virtual Central Processing Unit

A software-defined scheduling entity managed by a hypervisor that represents a single physical CPU core allocated to a virtual machine.

VDI

Virtual Desktop Infrastructure

An IT architecture that hosts user desktop environments within virtual machines running on centralized data center servers.

VFIO

Virtual Function I/O

A Linux kernel framework that provides secure, direct user-space access to PCI devices.

VHS

Virtualization Host Server

A physical server hardware platform explicitly configured with a hypervisor to host, manage and run multiple virtual machines.

VMCS

Virtual Machine Control Structure

VMX non-root operation and VMX transitions are controlled by a data structure called a virtual-machine control structure (VMCS). Access to the VMCS is managed through a component of processor state called the VMCS pointer (one per logical processor). The value of the VMCS pointer is the 64-bit address of the VMCS. The VMCS pointer is read and written using the instructions VMPTRST and VMPTRLD. The *VMM* configures a VMCS using the VMREAD, VMWRITE, and VMCLEAR instructions. A *VMM* could use a different VMCS for each virtual machine that it supports. For a virtual machine with multiple logical processors (virtual processors), the *VMM* could use a different VMCS for each virtual processor.

VMDq

Virtual Machine Device Queues

Multi-queue network adapters that support multiple VMs at the hardware level, having separate packet queues associated with the different hosted VMs (by means of the IP addresses of the VMs).

VMM

Virtual Machine Monitor (Hypervisor)

Software, firmware or hardware that creates, runs and manages virtual machines, mediating access to the underlying physical resources.

VMX

Virtual Machine Extensions

Intel's hardware-assisted virtualization technology architecture added to x86 processors to support native hypervisor execution.

VMX Root / Non-Root Operation

Processor execution modes introduced by *VMX*. The *VMM* runs in VMX root operation, while guest software runs in VMX non-root operation. Transitions between the two modes are called VMX transitions.

VPID

Virtual Processor Identifiers

New support for software control of *TLB* (VPID improves *TLB* performance with small *VMM* development effort).

VT-d

Virtualization Technology for Directed I/O

Intel's (https://software.intel.com/en-us/articles/intel-virtualization-technology-for-directed-io-vt-d-enhancing-intel-platforms-for-efficient-virtualization-of-io-devices) hardware implementation of an *IOMMU*, providing device isolation and direct memory access (DMA) translation for virtualized environments.

vTPM

Virtual Trusted Platform Module

Component to establish end-to-end integrity for guests via Trusted Computing.

Glossary

General

Create Virtual Machine Wizard

Virtual Machine Manager provides a graphical interface to guide you through the steps to create virtual machines. It can also be run in text mode by entering **`virt-install`** at a command prompt in the host environment.

Hardware-assisted

Intel* and AMD* provide virtualization hardware-assisted technology. This reduces the frequency of VM IN/OUT (fewer VM traps), because software is a major source of overhead, and increases the efficiency (the execution is done by the hardware). Moreover, this reduces the memory footprint, provides better resource control, and allows secure assignment of specific I/O devices.

Host Environment

The desktop or command line environment that allows interaction with the host computer's environment. It provides a command line environment and can also include a graphical desktop, such as GNOME or IceWM. The host environment runs as a special type of virtual machine that has privileges to control and manage other virtual machines.

Hypervisor

The software that coordinates the low-level interaction between virtual machines and the underlying physical computer hardware.

Paravirtualized Frame Buffer

The video output device that drives a video display from a memory buffer containing a complete frame of data for virtual machine displays running in paravirtual mode.

VHS

A Virtualization Host Server (VHS) is the physical computer running SUSE virtualization platform software. The virtualization environment consists of the hypervisor, the host environment, virtual machines and associated tools, commands and configuration files. Other commonly used terms include host, Host Computer, Host Machine (HM), Virtual Server (VS), Virtual Machine Host (VMH) and VM Host Server (VHS).

VirtFS

VirtFS is a new paravirtualized file system interface designed for improving pass-through technologies in the KVM environment. It is based on the VirtIO framework.

Virtual Machine

A software-defined computer environment (VM) capable of hosting an independent guest operating system and its associated applications. Also commonly referred to as a VM Guest.

Virtual Machine Manager

A software program that provides a graphical user interface for creating and managing virtual machines.

Virtualized

The status or condition of a guest operating system or application executing within an isolated virtual machine environment.

CPU

CPU capping

The resource restriction process that limits a virtual CPU's allocation window to a defined percentage (from 1 to 100 percent) of available physical CPU capacity.

CPU hotplugging

CPU hotplugging is used to describe the functions of replacing/adding/removing a CPU without shutting down the system.

CPU over-commitment

Virtual CPU over-commitment is the ability to assign more virtual CPUs to VMs than the actual number of physical CPUs present in the physical system. This procedure does not increase the overall performance of the system, but may be useful for testing purposes.

CPU pinning

Also referred to as processor affinity; CPU pinning enables the binding and unbinding of a process or a thread to a central processing unit (CPU) or a range of CPUs.

Network

Bridged Networking

A type of network connection that lets a virtual machine be identified on an external network as a unique identity that is separate from and unrelated to its host computer.

Empty Bridge

A type of network bridge that has no physical network device or virtual network device provided by the host. This lets virtual machines communicate with other virtual machines on the same host, but not with the host or on an external network.

External Network

The network outside a host's internal network environment.

Internal Network

A type of network configuration that restricts virtual machines to their host environment.

Local Bridge

A type of network bridge that has a virtual network device but no physical network device provided by the host. This lets virtual machines communicate with the host and other virtual machines on the host. Virtual machines can communicate on an external network through the host.

Network Address Translation (NAT)

A type of network connection that lets a virtual machine use the IP address and MAC address of the host.

No Host Bridge

A type of network bridge that has a physical network device but no virtual network device provided by the host. This lets virtual machines communicate on an external network but not with the host. This lets you separate virtual machine network communications from the host environment.

Traditional Bridge

A type of network bridge that has both a physical network device and a virtual network device provided by the host.

Storage

AHCI

The Advanced Host Controller Interface (AHCI) is a technical standard defined by Intel* that specifies the operation of Serial ATA (SATA) host bus adapters in a non-implementation-specific manner.

Block Device

Data storage components, such as CD-ROM drives or disk drives, that move data in the form of blocks. Partitions and volumes are also considered block devices.

File-Backed Virtual Disk

A virtual disk based on a file, also called a disk image file.

Raw Disk

A method of accessing data on a disk at the individual byte level instead of through its file system.

Sparse image file

A disk image file that does not reserve its entire amount of disk space but expands as data is written to it.

xvda

The drive designation given to the first virtual disk on a paravirtual machine.