



NVIDIA Grace Software with SUSE Linux Enterprise Server 15

Installation Guide

Document History

DI-11466-001_05

Version	Date	Description of Change
01	August 2, 2023	Initial Release.
02	August 31, 2023	<ul style="list-style-type: none">• Added Appendix E: Known Issues.• Updated D.2 with a minimum kernel version that contains the fix for that issue.
03	October 5, 2023	Added post-install steps.
04	December 18, 2023	<ul style="list-style-type: none">• Updated installation media requirements.• Removed unnecessary workarounds.• Added Appendix A.1, and removed Appendix E.
05	March 8, 2024	Added a note to Appendix A.1 about explicitly requiring the open-source GPU driver.

Table of Contents

Introduction	4
Related Documentation	4
Prerequisites	4
SUSE Subscription	4
SUSE Linux Enterprise Server Installation Media	4
Access to Repositories	5
SUSE Repositories	5
Installing SUSE Linux Enterprise Server	5
Obtaining SUSE Linux Enterprise Server 15	5
Remotely Booting the SUSE Linux Enterprise Server 15 ISO File	6
Installing the SUSE Linux Enterprise Server	9
Appendix A: Installing Software	24
A.1 NVIDIA GPU Driver and CUDA Toolkit	24
Appendix B: Changing the BMC Login	25
Appendix C: Automated Installation	25
C.1 Installing with AutoYaST	25
Appendix D: Platform-Dependent Workarounds	26
D.1 All Grace Platforms	26

Introduction

NVIDIA® Grace systems can run SUSE Linux Enterprise Server (SLES) and take advantage of the advanced Grace features.

This document explains how to install and configure Grace systems with SUSE Linux Enterprise Server 15 SP5.



Attention: Although it might be possible to use other Linux distributions that are related to (or inspired by) SLES, for example openSUSE Leap, only SLES has been pre-validated with the steps in this guide for installation on Grace platforms.

Related Documentation

Refer to the [SUSE Linux Enterprise Server 15 Documentation Portal](#) for more information about SLES.

Prerequisites

This section lists the required (or recommended) prerequisites.

SUSE Subscription

A SUSE subscription is required to install and use SLES 15 on the Grace platform. A subscription allows you to obtain update packages and additional packages for SLES. To purchase a subscription or obtain a free evaluation subscription, go to the [SUSE Customer Center](#) (SCC).

SUSE Linux Enterprise Server Installation

Media

The installation method in this guide requires the SLES 15 SP5 Quarterly Update 1 (or later) installation media. To obtain these Quarterly Update (QU) ISO images:

1. Go to the [SLES download page](#).
2. Select **Stable Release 15 SP5 for the Arm Architecture**.
3. For the latest installer images link, navigate to the corresponding SCC product page from the login page to SUSE Customer Center.

Access to Repositories

The repositories can be accessed from the internet. If you are using a proxy server, follow the instructions in the [SLES Deployment Guide](#) to ensure that the system can access the necessary URLs.

SUSE Repositories

To install software for the Grace platform over SLES 15, you need access to the following [repositories](#):

- Basesystem Module: `sle-module-basesystem`
- Server Applications Module: `sle-module-server-applications`

Installing SUSE Linux Enterprise Server

SUSE provides several methods to install SLES (refer to the SLES [Deployment Guide](#) for more information). **Before you install**, review the [platform-dependent workarounds](#) section in this document to determine whether there are any modifications that are required for your environment.

This section describes how to install SLES using the Quick Install method and reclaim the disk space used by an existing installation in the process. It describes a minimal installation. If you have a preferred method to install SLES, you can skip this section but ensure that you reclaim the disk space that is used by an existing OS installation.

This method installs SUSE Linux Enterprise Server on the Grace system remotely through facilities hosted on the BMC.

Obtaining SUSE Linux Enterprise Server 15

Obtain the SLES ISO for Arm® image (aarch64) and store it on your local disk. Refer to [SUSE Linux Enterprise Server Installation Media](#) or [Downloading SUSE Linux Enterprise Server](#) for the instructions.

Remotely Booting the SUSE Linux Enterprise Server 15 ISO File

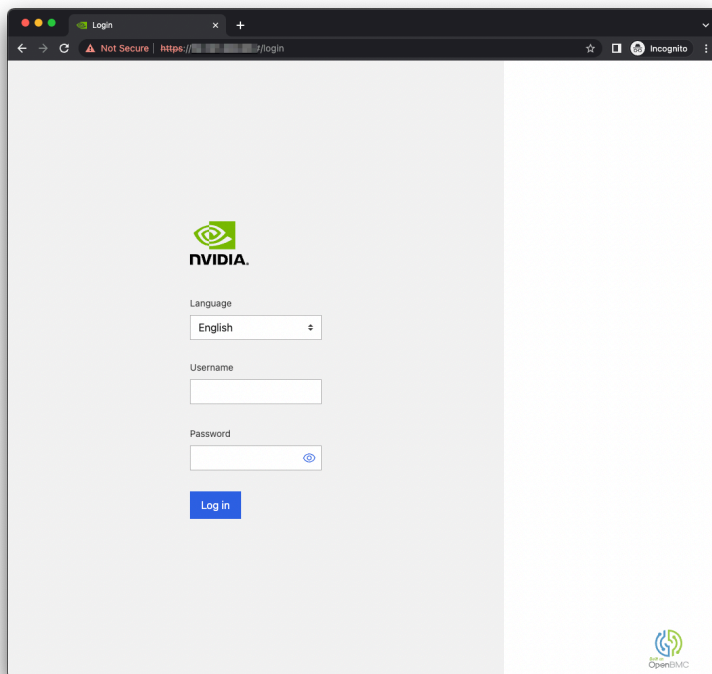


Attention: Here is some important information to know **before** you begin:

- The Grace BMC comes with default login credentials, and NVIDIA recommends that you create a unique user ID and password. Contact the system vendor if you have issues logging into the BMC.
- As a performance consideration, for networks that are distributed across a wide geographical area, such as a corporate VPN, we recommend that you use a browser from a host near the target server.
This can be accomplished by using a remote application software solution, for example VNC, or by exporting the application X session and tunneling over SSH.
- This sequence is intended for the NVIDIA reference BMC, and your results might vary when using an IBV BMC.

1. Connect to the BMC.
 - a. Open a browser in your LAN, navigate to **https://<BMC-IP-address>/**, and log in.

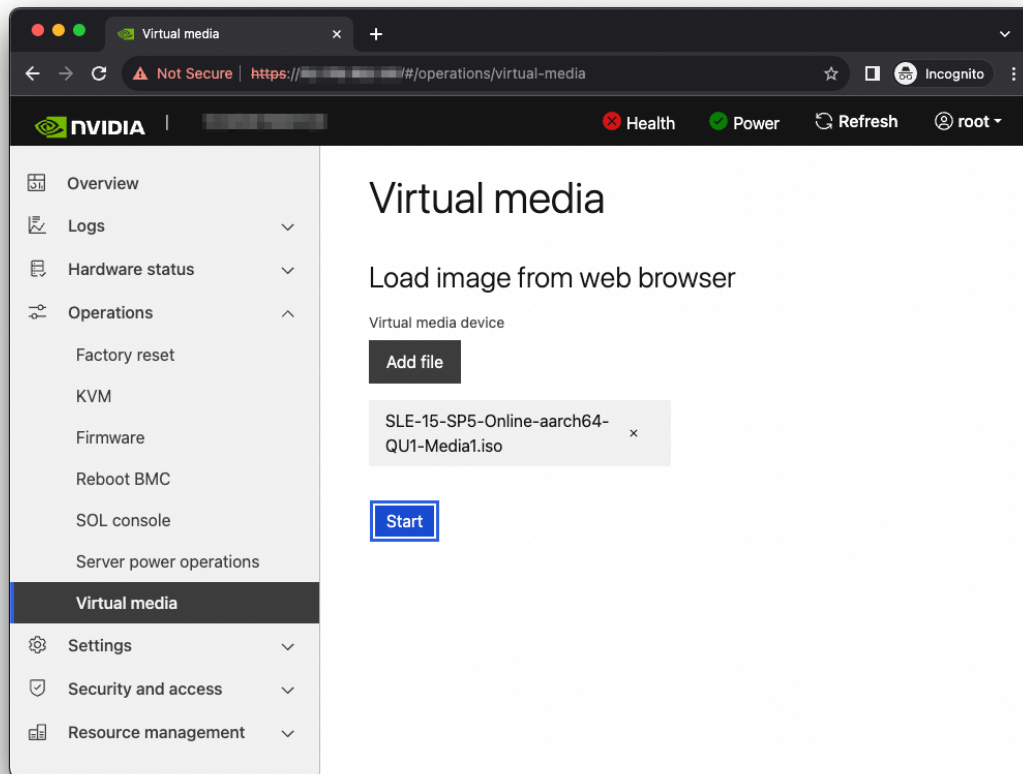
Figure 1. BMC Login



2. Download the ISO image to a location that can be accessed by the browser.
3. Set up the ISO image as virtual media.

- a. From the left hand menu, expand **Operations**.
- b. Select **Virtual media**.
- c. Click **Add file**, navigate to the downloaded ISO image, and select it.
- d. To begin serving the ISO image to the target server, click **Start**.

Figure 2. BMC Virtual Media

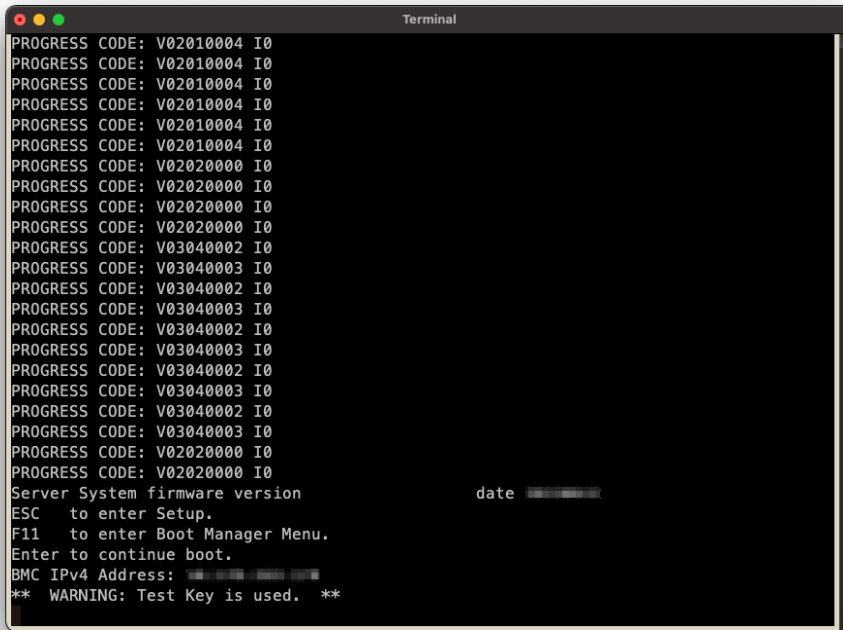


4. Boot from the virtual media.

Typically, the default boot order does not boot the CD-ROM image. You can change this in the BIOS or as a one-time option in the boot menu.

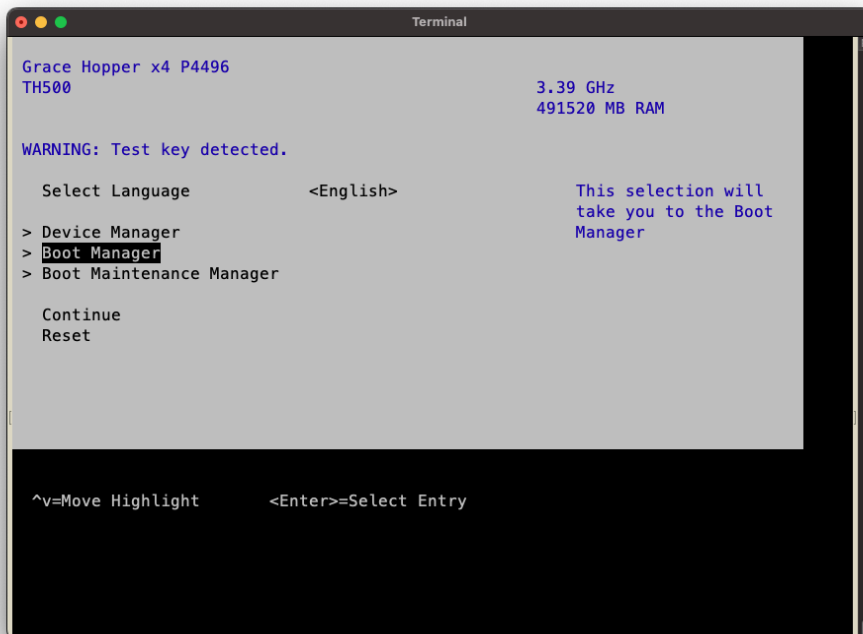
 - a. Connect to the console.
 - i. From the left hand menu, expand **Operations**, and select the SOL **console**.
 - ii. Alternatively, SSH to the BMC from a terminal emulator, log in, and at the prompt, run the `obmc-console-client` command.
 - b. To bring up the boot menu, press **Escape** or **F11** at the beginning of the boot process.

Figure 3. Console Splash Screen



```
Terminal
PROGRESS CODE: V02010004 I0
PROGRESS CODE: V02010004 I0
PROGRESS CODE: V02010004 I0
PROGRESS CODE: V02010004 I0
PROGRESS CODE: V02010004 I0
PROGRESS CODE: V02010004 I0
PROGRESS CODE: V02020000 I0
PROGRESS CODE: V02020000 I0
PROGRESS CODE: V02020000 I0
PROGRESS CODE: V02020000 I0
PROGRESS CODE: V02020000 I0
PROGRESS CODE: V03040002 I0
PROGRESS CODE: V03040003 I0
PROGRESS CODE: V03040002 I0
PROGRESS CODE: V03040003 I0
PROGRESS CODE: V03040002 I0
PROGRESS CODE: V03040003 I0
PROGRESS CODE: V03040002 I0
PROGRESS CODE: V03040003 I0
PROGRESS CODE: V03040002 I0
PROGRESS CODE: V03040003 I0
PROGRESS CODE: V03040002 I0
PROGRESS CODE: V03040003 I0
PROGRESS CODE: V02020000 I0
PROGRESS CODE: V02020000 I0
Server System firmware version      date ██████████
ESC   to enter Setup.
F11  to enter Boot Manager Menu.
Enter to continue boot.
BMC IPv4 Address: ██████████
** WARNING: Test Key is used. **
```

Figure 4. Boot Manager Menu



```
Terminal
Grace Hopper x4 P4496
TH500                               3.39 GHz
                                      491520 MB RAM

WARNING: Test key detected.

  Select Language      <English>          This selection will
                                      take you to the Boot
                                      Manager
> Device Manager
> Boot Manager
> Boot Maintenance Manager

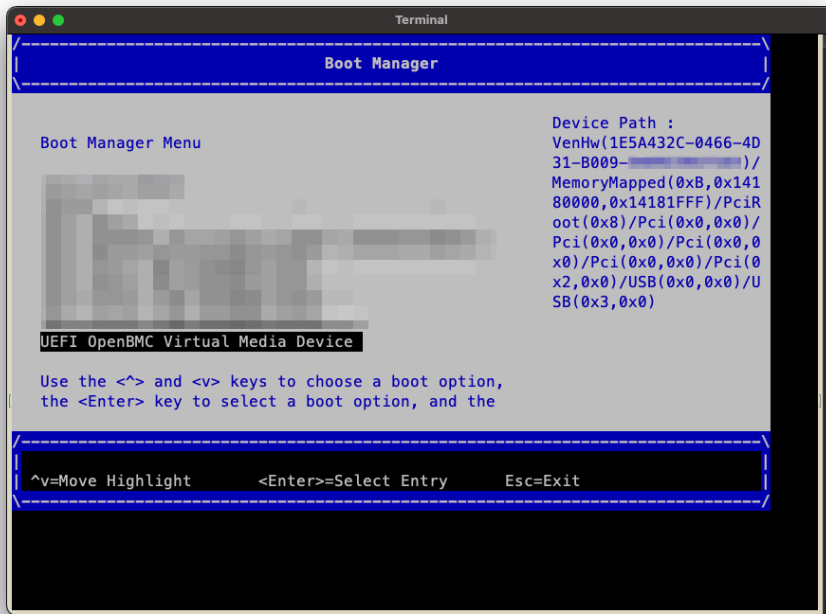
Continue
Reset

^v=Move Highlight      <Enter>=Select Entry
```

- c. In the boot menu, select **UEFI OpenBMC Virtual Media Device** as the boot device and press **Enter**.

d. Follow the instructions in [Installing SUSE Linux Enterprise Server](#).

Figure 5. Boot Manager



Installing the SUSE Linux Enterprise Server

Prerequisites: This section assumes you have already booted the SLES ISO image.

1. After booting the ISO image, the GRUB splash menu will appear.
2. Press **t** to display the GRUB boot menu for the installer.
The menu will time out after 60 seconds. If no action is taken, it will proceed with booting the default selection (**Installation**).

Figure 6. GRUB Splash Menu


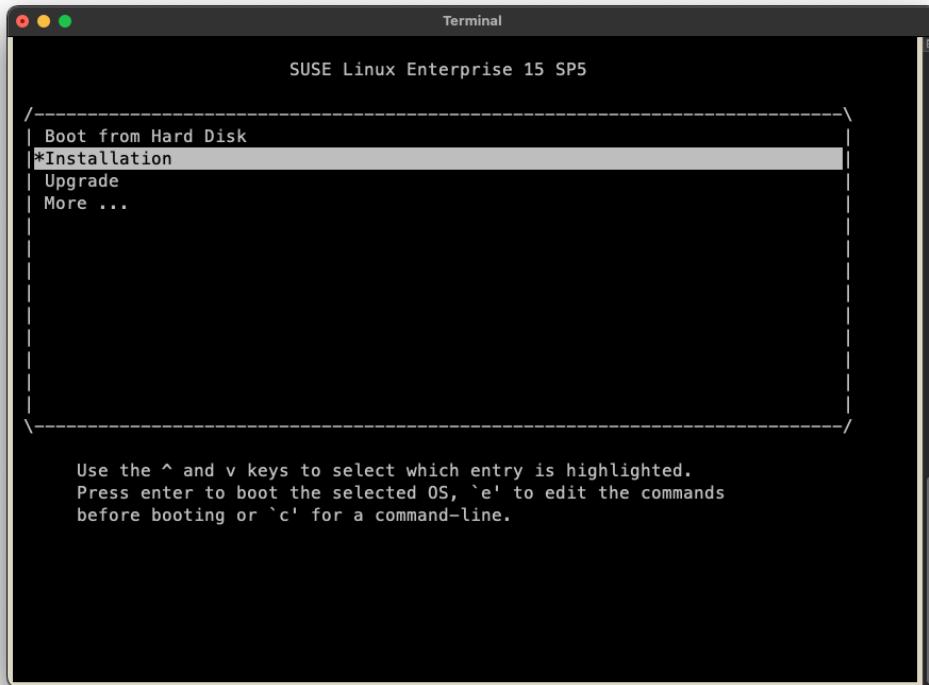
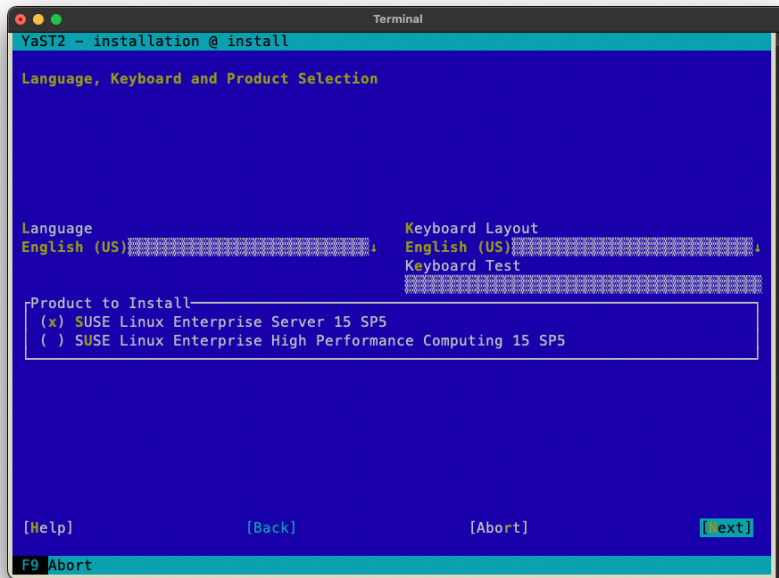


Figure 7. GRUB Boot Menu



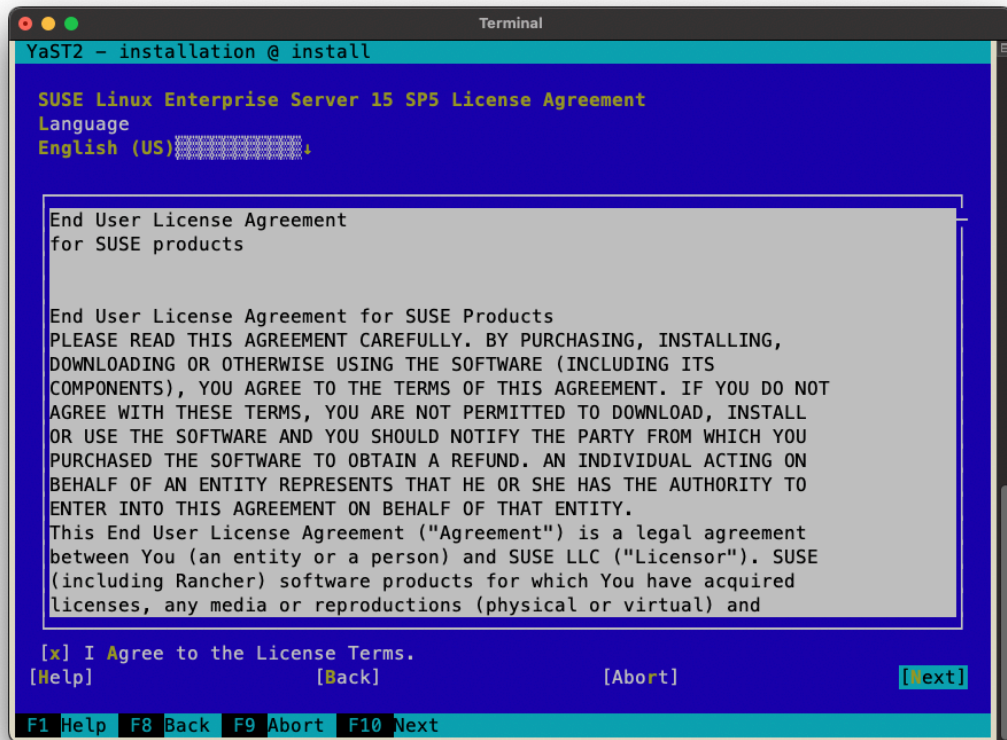
3. To automate the installation with an AutoYaST configuration file, refer to [Installing with AutoYaST](#).
4. Select **Installation** or wait for the timeout.
Refer to the [SUSE Linux Enterprise Server Installation Quick Start](#) for guidance on using the installer.
5. Follow the installer prompts to configure the manual installation.
6. Select the **Language**, **Keyboard**, and **Product**.

Figure 8. Language, Keyboard, and Product Selection



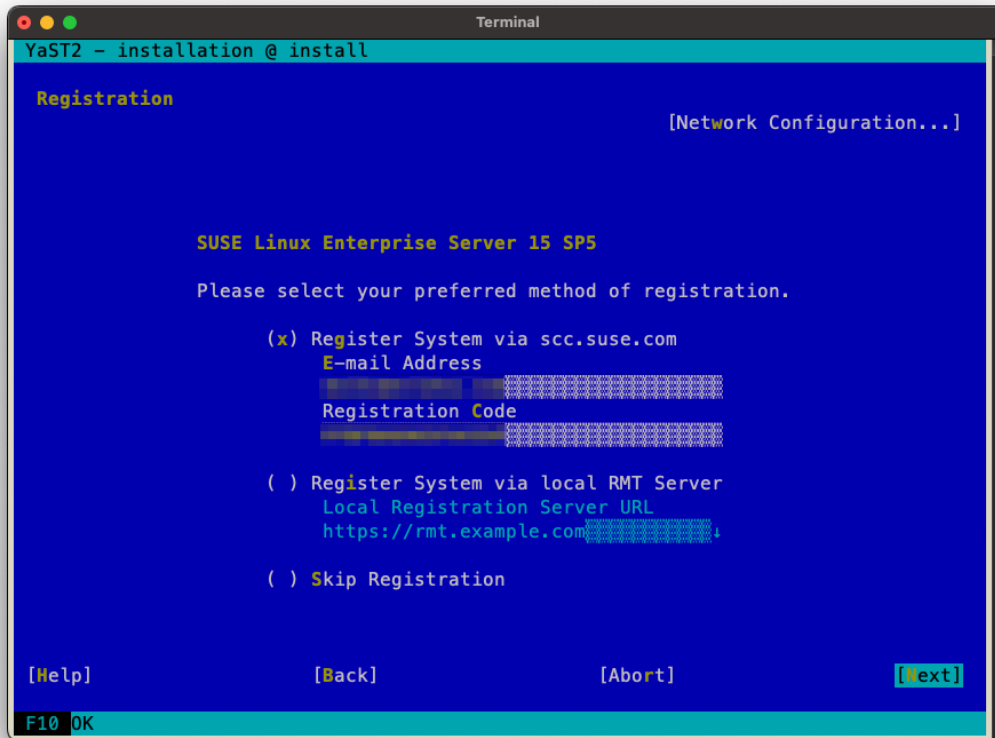
7. Review the End User License Agreement (EULA) and indicate whether you agree to the License Terms.

Figure 9. EULA Prompt



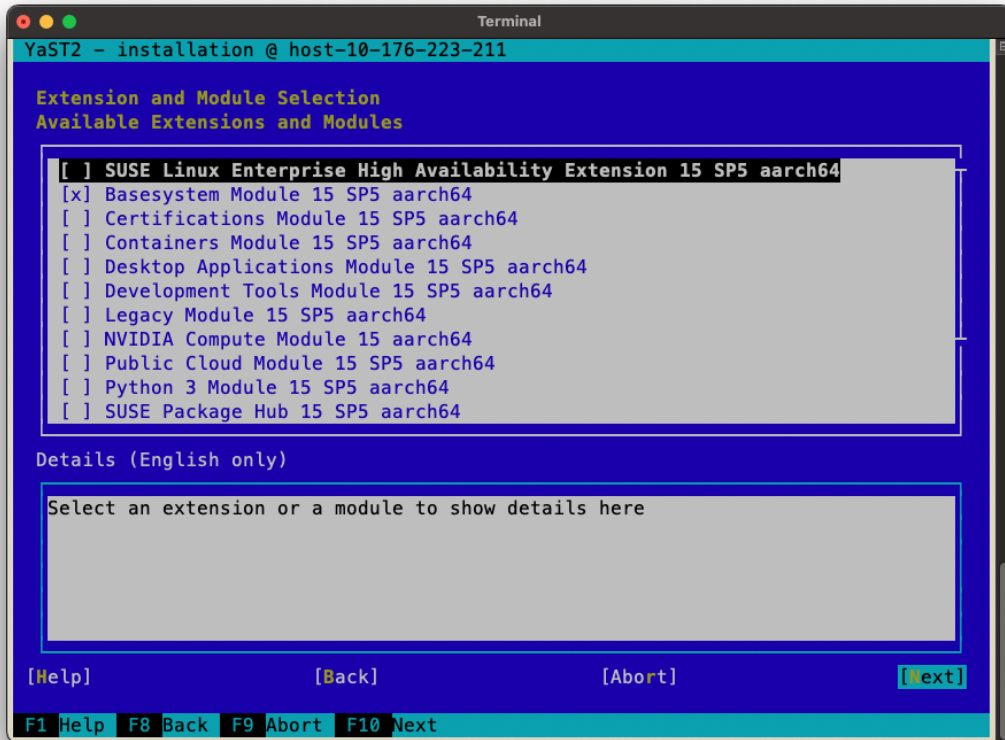
8. Register the system with SUSE through the SUSE Customer Center or a local RMT server.

Figure 10. System Registration



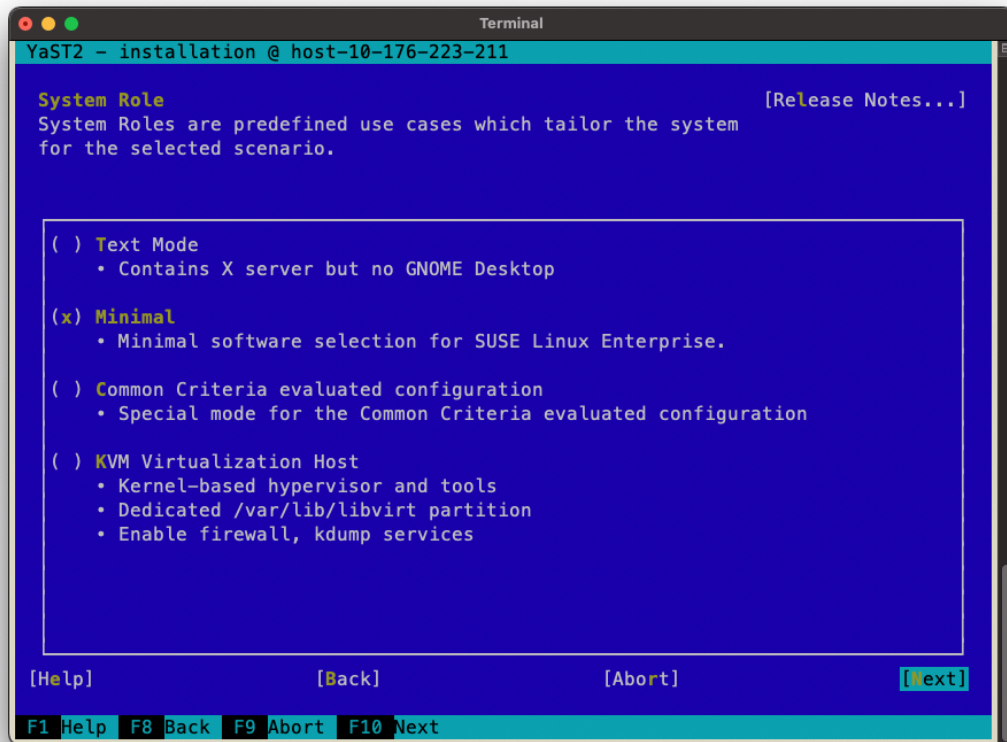
9. Select the modules and add-on products you want to install

Figure 11. Module Selection



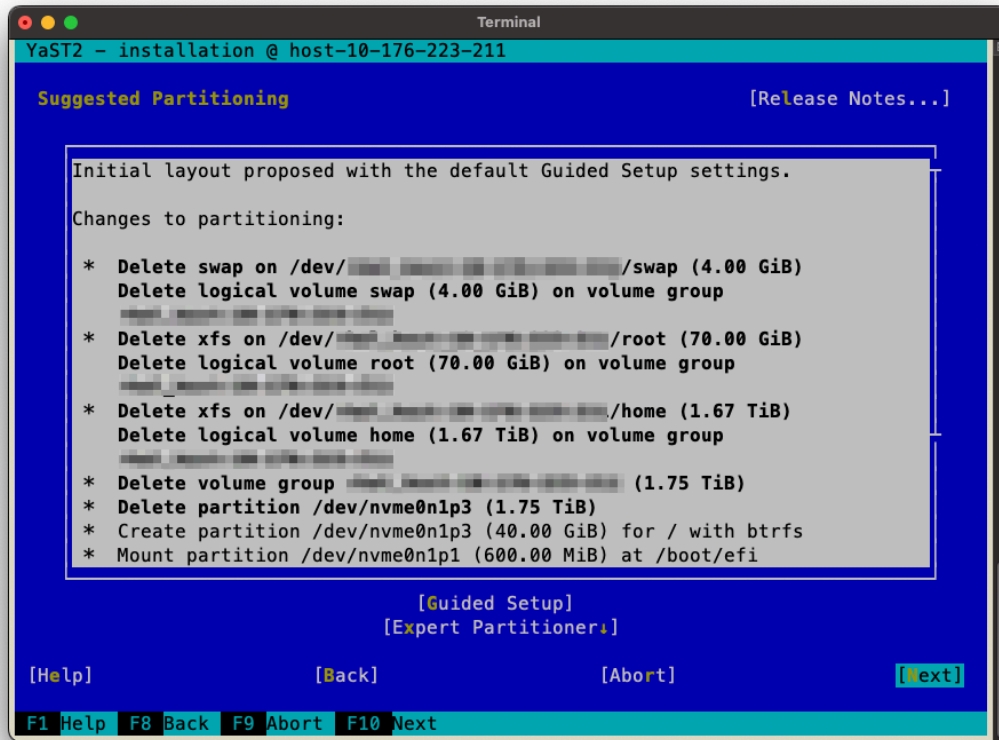
10. Select the system role, which dictates which software packages are installed by default.

Figure 12. System Role Selection



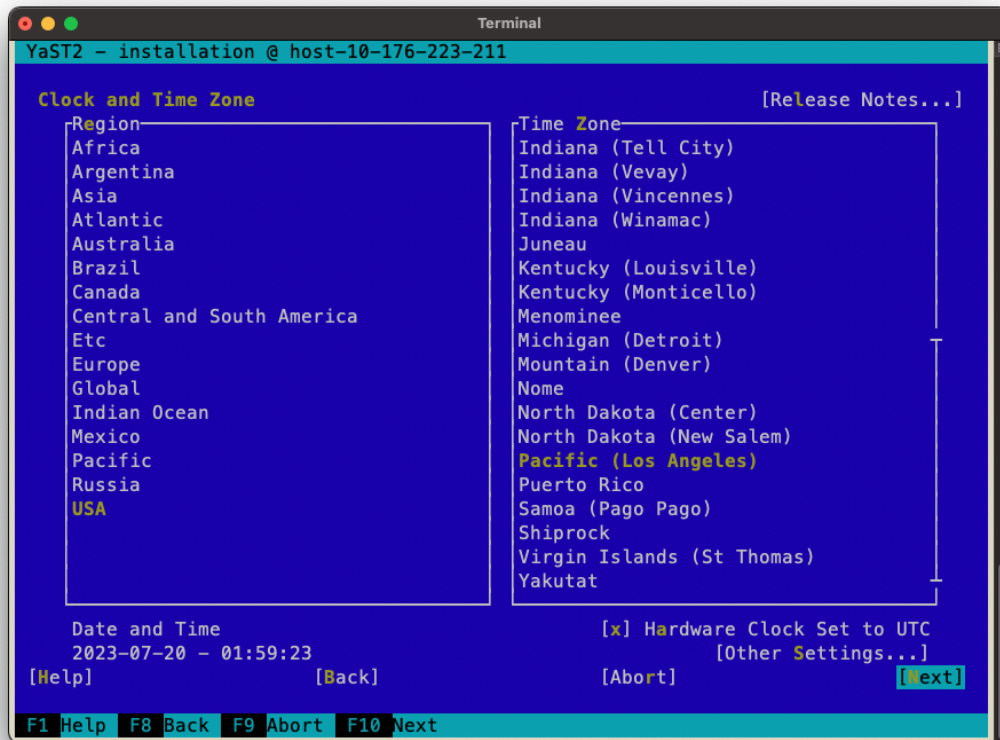
11. Select a disk partitioning scheme.

Figure 13. Partition Scheme Selection



12. Configure the clock and time zone.

Figure 14. Clock and Timezone Selection



13. Create the users and set the root password.
14. Review the installation settings.

Figure 15. User and Password Selection

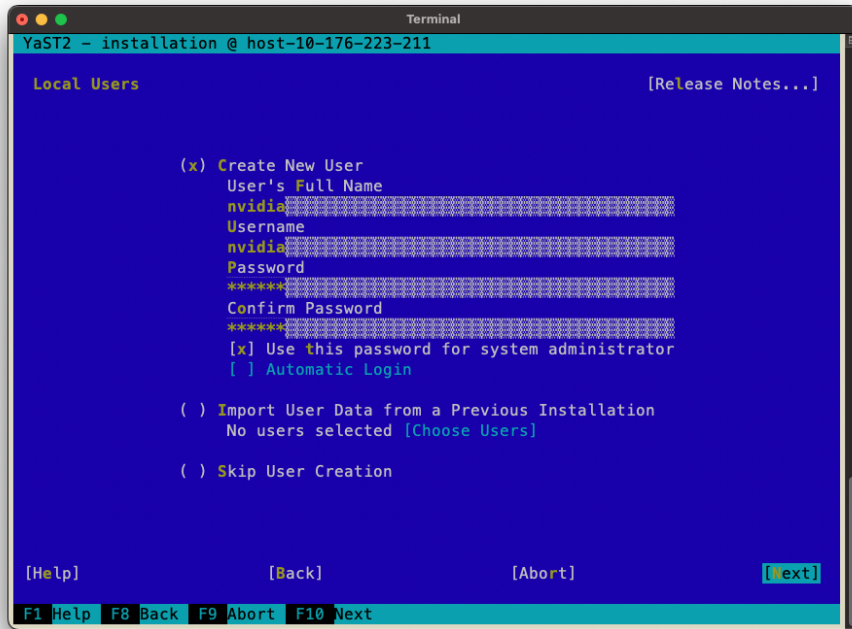
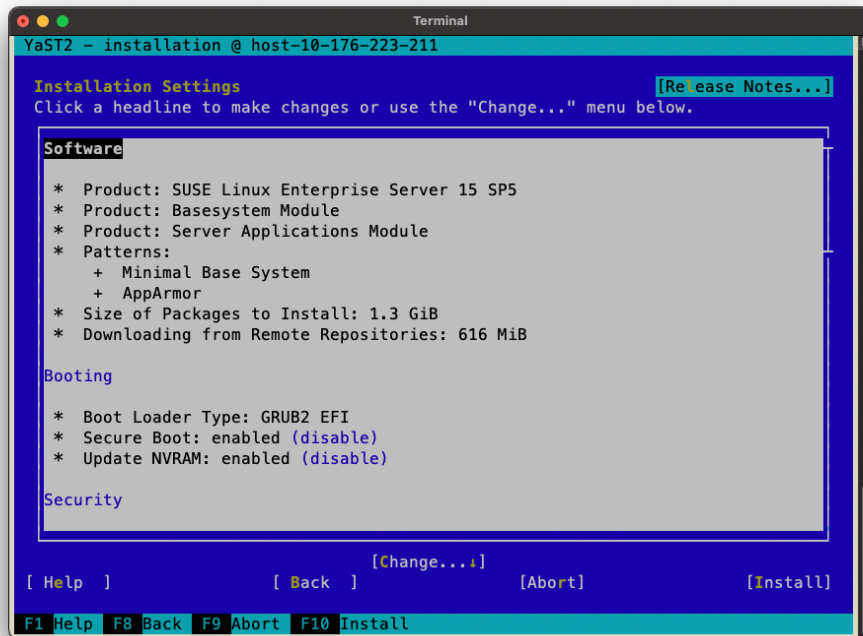
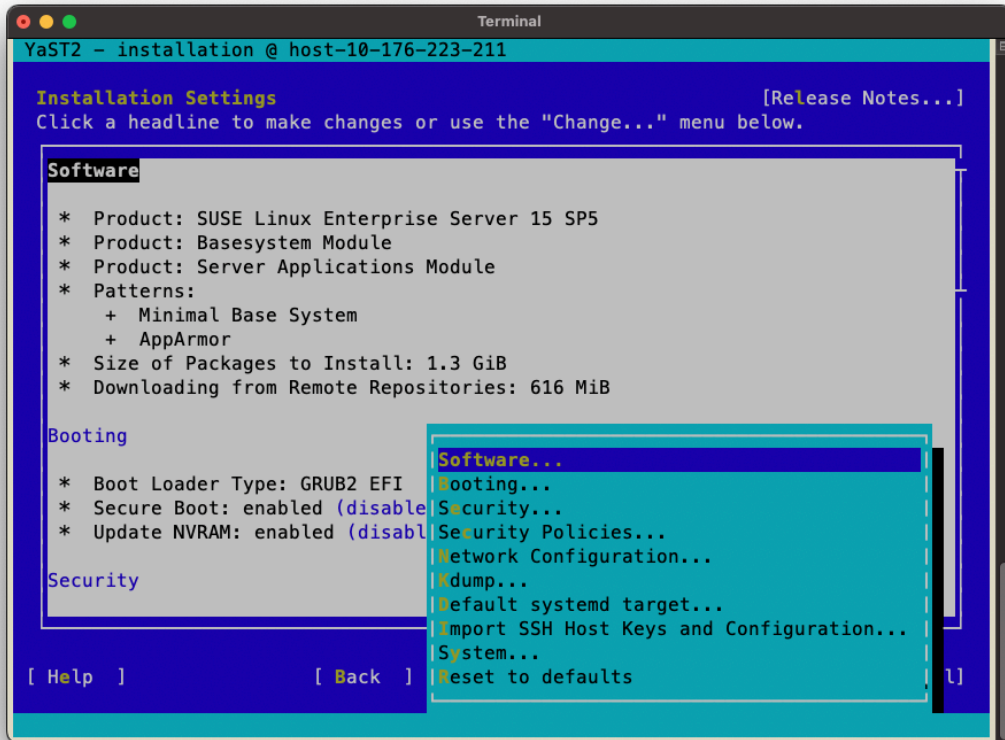


Figure 16. Installation Settings Summary




15. From the **Installation Settings** summary, select the **Change** menu and then select the **Software** submenu to alter the installation to install the 64k kernel as the sole, default kernel for the system.

Figure 17. Software Change Menu




16. From the software change pane navigate to the **Search Phrase** box, and search for **kernel-64k**.
17. Select the **kernel-64k** package.

Figure 18. 64k Kernel Selection



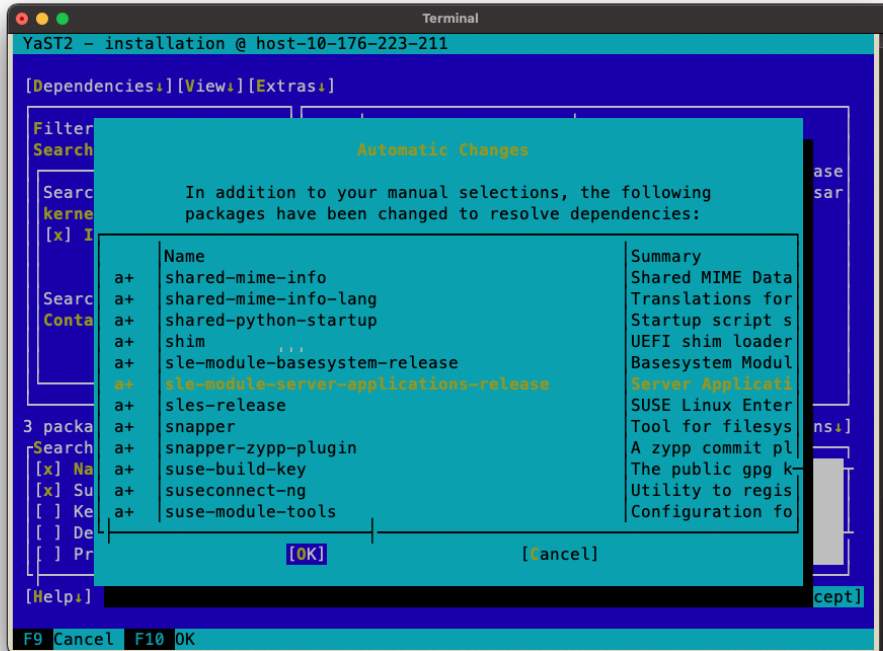
- 18. Return to the **Search Phrase** box and search for **kernel-default**.
- 19. Deselect the **kernel-default** package.

Figure 19. Default Kernel Deselection



20. Navigate to and select **Accept** to update the planned installation.
Additional changes will automatically be made due to package dependencies.

Figure 20. Dependency Updates Review



21. Select **Install** and then confirm to begin the installation.

Figure 21. Install Selection

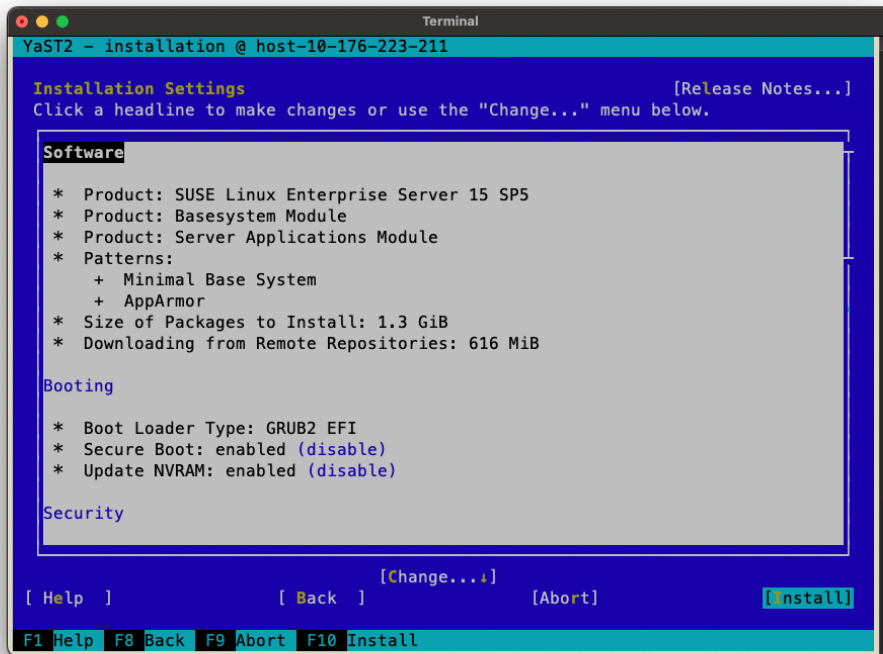
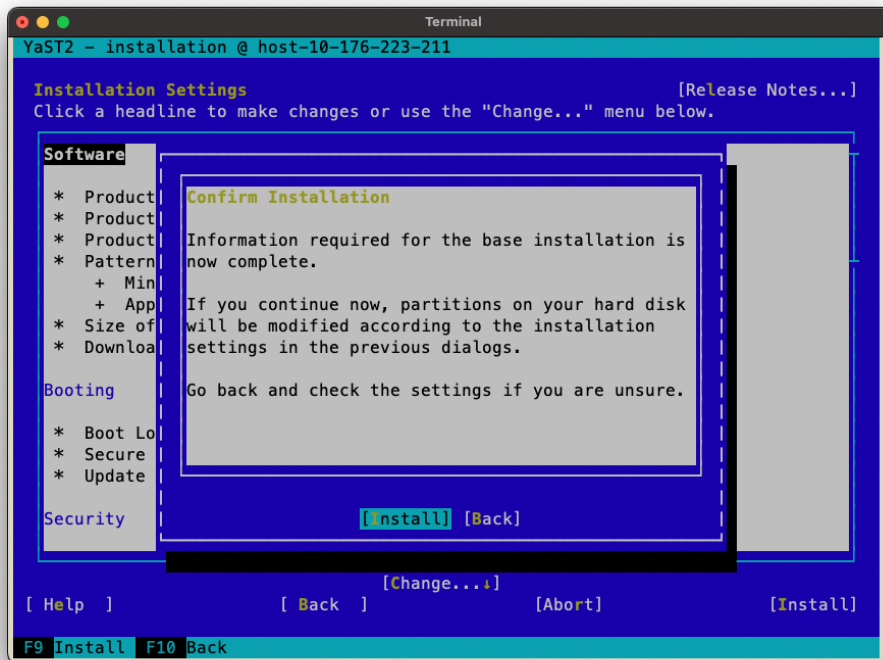


Figure 22. Install Confirmation



22. The system will reboot.

23. After the system is rebooted, NVIDIA recommends the following post-installation settings:
- a. SLES 15 SP5 does not enable the network time protocol (NTP) by default. This can create inconsistent time settings on the system, which can lead to unintended side effects. For example, SSL certificates might fail verification. NVIDIA recommends that you enable NTP on SLES 15 SP5 with the chrony daemon by using the following commands:

```
sudo systemctl start chronyd.service
sudo systemctl enable chronyd.service
```
 - b. By default, SLES 15 SP5 sets the CPU scaling governor to `ondemand`, and NVIDIA tools like `nvidia-persistenced` can take longer to run with this setting. To permanently set the CPU scaling governor to `performance`, we recommend that you add the `cpufreq.default_governor=performance` kernel parameter. With administrative privileges, edit the `/etc/default/grub` file.
 - I. Append the parameter `cpufreq.default_governor=performance` to the end of the list of kernel boot parameters specified in `GRUB_CMDLINE_LINUX_DEFAULT`.
 - II. Save the file and exit the editor.
 - III. Run the following command to update GRUB:

```
sudo update-bootloader -refresh
```
 - c. Reboot the system.

The installation on Grace is now complete.

Appendix A: Installing Software

SUSE Linux Enterprise Server uses the Zypper package manager to install, update, and remove packages. The utility can also be used to manage repositories. For more information about using Zypper, refer to the [SUSE Zypper Package Manager](#) article.

A.1 NVIDIA GPU Driver and CUDA Toolkit

Refer to the [NVIDIA CUDA Installation Guide for Linux](#) for instructions on how to install the NVIDIA GPU driver and CUDA support for SLES. The `R535.129.03` driver is the minimum level required for the Hopper GPU.

The following commands can be used to install the minimum levels required for the Hopper GPU:

```
sudo zypper install -y kernel-64kb-devel
sudo usermod -a -G video <username>
sudo zypper addrepo
https://developer.download.nvidia.com/compute/cuda/repos/sles15/sbsa/cuda-sles15.repo
sudo zypper refresh
```



```
sudo zypper install -y cuda-toolkit-12-2
sudo zypper install -y nvidia-open-gfxG05-kmp-64kb=535.129.03_k5.14.21_150500.55.28-0
sudo sed -i "s/NVreg_DeviceFileGID\=.* /NVreg_DeviceFileGID\=$(grep video /etc/group |
cut -d ':' -f3) /g" /etc/modprobe.d/50-nvidia-64kb.conf
sudo reboot now
```



Note: The open-source GPU driver is required for Hopper GPUs.

Appendix B: Changing the BMC Login

The NVIDIA Grace servers include a base management controller (BMC) for out-of-band management of the Grace system. NVIDIA recommends that you create a unique username and password as soon as possible.

Appendix C: Automated Installation

This section provides information about automated installations.

C.1 Installing with AutoYaST

AutoYaST provides a way to automate the installation process by providing a configuration file with the answers to commonly asked installation questions. NVIDIA provides an AutoYaST template for Grace platforms:

<https://repo.download.nvidia.com/baseos/sles/sles-files/15/sles15-ay/grace-autoyast.xml>

For more information about using AutoYaST files with SUSE Linux Enterprise Server 15, refer to the [SUSE AutoYaST Guide](#).

In these files, there are tags that you must replace with site-specific information including the following:

- Language
- Keyboard
- Timezone
- Hostname

Each tag is in the <CHANGE_YOUR_xxxx> form and must be replaced with your information.

1. After you have an autoyast file customized for your installation, place that file in a location that can be accessed by NFS, FTP, HTTP, or HTTPS.
2. After booting from the installation medium, when the grub menu appears, press **e** to edit the grub entry and append `autoyast=<URL>` to the list of kernel boot parameters.

For example:

```
autoyast=https://192.168.1.2/autoyast/grace-autoyast.xml
```

Appendix D: Platform-Dependent Workarounds

Some Grace platforms require temporary (or permanent) alterations to their configurations to work around known issues, such as hardware errata. These workarounds are described in the following sections by the corresponding Grace platform.

D.1 All Grace Platforms

CUDA Application Workaround

CUDA applications on the Grace-Hopper platform require ATS support. Currently, ATS is not enabled on the arm64 platform when IOMMU passthrough is enabled. NVIDIA is working with the Linux kernel community and SUSE to resolve this issue.

Linux provides the `iommu.passthrough` kernel parameter to configure the DMA to use (or not use) the IOMMU to access the memory for addressing. By default, SLES 15 SP5 sets the IOMMU in passthrough mode, which prevents CUDA applications from running. We recommend that you add the `iommu.passthrough=0` kernel parameter until this issue is resolved.

- To permanently deploy this workaround so that it is **always** active upon boot:
 1. With administrative privileges, edit the `/etc/default/grub` file.
 - a. Append the parameter `iommu.passthrough=0` to the end of the list of kernel boot parameters specified in `GRUB_CMDLINE_LINUX_DEFAULT`.
 - b. Save the file and exit the editor.
 2. Run the following command to update GRUB:

```
sudo update-bootloader -refresh
```
 3. Reboot the system.
 4. To verify the workaround, run the following commands:

```
sudo dmesg | grep "iommu: Default"
iommu: Default domain type: Translated (set via kernel command line)
```

5. If iommu passthrough is required for a device or iommu group, refer to the instructions in https://www.kernel.org/doc/Documentation/ABI/testing/sysfs-kernel-iommu_group.



Note: When this workaround is applied from the SLES Installer, it is automatically included in the installed system.

Refer to [SUSE Modifying Kernel Boot Parameters](#) for additional guidance about modifying kernel boot parameters.

Notice

This document is provided for information purposes only and shall not be regarded as a warranty of a certain functionality, condition, or quality of a product. NVIDIA Corporation ("NVIDIA") makes no representations or warranties, expressed or implied, as to the accuracy or completeness of the information contained in this document and assumes no responsibility for any errors contained herein. NVIDIA shall have no liability for the consequences or use of such information or for any infringement of patents or other rights of third parties that may result from its use. This document is not a commitment to develop, release, or deliver any Material (defined below), code, or functionality.

NVIDIA reserves the right to make corrections, modifications, enhancements, improvements, and any other changes to this document, at any time without notice.

Customer should obtain the latest relevant information before placing orders and should verify that such information is current and complete.

NVIDIA products are sold subject to the NVIDIA standard terms and conditions of sale supplied at the time of order acknowledgement, unless otherwise agreed in an individual sales agreement signed by authorized representatives of NVIDIA and customer ("Terms of Sale"). NVIDIA hereby expressly objects to applying any customer general terms and conditions with regards to the purchase of the NVIDIA product referenced in this document. No contractual obligations are formed either directly or indirectly by this document.

NVIDIA products are not designed, authorized, or warranted to be suitable for use in medical, military, aircraft, space, or life support equipment, nor in applications where failure or malfunction of the NVIDIA product can reasonably be expected to result in personal injury, death, or property or environmental damage. NVIDIA accepts no liability for inclusion and/or use of NVIDIA products in such equipment or applications and therefore such inclusion and/or use is at customer's own risk.

NVIDIA makes no representation or warranty that products based on this document will be suitable for any specified use. Testing of all parameters of each product is not necessarily performed by NVIDIA. It is customer's sole responsibility to evaluate and determine the applicability of any information contained in this document, ensure the product is suitable and fit for the application planned by customer, and perform the necessary testing for the application in order to avoid a default of the application or the product. Weaknesses in customer's product designs may affect the quality and reliability of the NVIDIA product and may result in additional or different conditions and/or requirements beyond those contained in this document. NVIDIA accepts no liability related to any default, damage, costs, or problem which may be based on or attributable to: (i) the use of the NVIDIA product in any manner that is contrary to this document or (ii) customer product designs.

No license, either expressed or implied, is granted under any NVIDIA patent right, copyright, or other NVIDIA intellectual property right under this document. Information published by NVIDIA regarding third-party products or services does not constitute a license from NVIDIA to use such products or services or a warranty or endorsement thereof. Use of such information may require a license from a third party under the patents or other intellectual property rights of the third party, or a license from NVIDIA under the patents or other intellectual property rights of NVIDIA.

Reproduction of information in this document is permissible only if approved in advance by NVIDIA in writing, reproduced without alteration and in full compliance with all applicable export laws and regulations, and accompanied by all associated conditions, limitations, and notices.

THIS DOCUMENT AND ALL NVIDIA DESIGN SPECIFICATIONS, REFERENCE BOARDS, FILES, DRAWINGS, DIAGNOSTICS, LISTS, AND OTHER DOCUMENTS (TOGETHER AND SEPARATELY, "MATERIALS") ARE BEING PROVIDED "AS IS." NVIDIA MAKES NO WARRANTIES, EXPRESSED, IMPLIED, STATUTORY, OR OTHERWISE WITH RESPECT TO THE MATERIALS, AND EXPRESSLY DISCLAIMS ALL IMPLIED WARRANTIES OF NONINFRINGEMENT, MERCHANTABILITY, AND FITNESS FOR A PARTICULAR PURPOSE. TO THE EXTENT NOT PROHIBITED BY LAW, IN NO EVENT WILL NVIDIA BE LIABLE FOR ANY DAMAGES, INCLUDING WITHOUT LIMITATION ANY DIRECT, INDIRECT, SPECIAL, INCIDENTAL, PUNITIVE, OR CONSEQUENTIAL DAMAGES, HOWEVER CAUSED AND REGARDLESS OF THE THEORY OF LIABILITY, ARISING OUT OF ANY USE OF THIS DOCUMENT, EVEN IF NVIDIA HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. Notwithstanding any damages that customer might incur for any reason whatsoever, NVIDIA's aggregate and cumulative liability towards customer for the products described herein shall be limited in accordance with the Terms of Sale for the product.

Trademarks

NVIDIA, the NVIDIA logo, Grace, are trademarks and/or registered trademarks of NVIDIA Corporation in the U.S. and other countries. Other company and product names may be trademarks of the respective companies with which they are associated.

VESA DisplayPort

DisplayPort and DisplayPort Compliance Logo, DisplayPort Compliance Logo for Dual-mode Sources, and DisplayPort Compliance Logo for Active Cables are trademarks owned by the Video Electronics Standards Association in the United States and other countries.

HDMI

HDMI, the HDMI logo, and High-Definition Multimedia Interface are trademarks or registered trademarks of HDMI Licensing LLC.

Arm

Arm, AMBA, and ARM Powered are registered trademarks of Arm Limited. Cortex, MPCore, and Mali are trademarks of Arm Limited. All other brands or product names are the property of their respective holders. "Arm" is used to represent ARM Holdings plc; its operating company Arm Limited; and the regional subsidiaries Arm Inc.; Arm KK; Arm Korea Limited.; Arm Taiwan Limited; Arm France SAS; Arm Consulting (Shanghai) Co. Ltd.; Arm Germany GmbH; Arm Embedded Technologies Pvt. Ltd.; Arm Norway, AS, and Arm Sweden AB.

OpenCL

OpenCL is a trademark of Apple Inc. used under license to the Khronos Group Inc.

Copyright

© 2023 NVIDIA Corporation & Affiliates. All rights reserved.