



NVIDIA Grace Software with Ubuntu 22.04

Installation Guide

Document History

DI-11485-001_05

Version	Date	Description of Change
01	August 22, 2023	Initial Release
02	December 18, 2023	<ul style="list-style-type: none">• Added Appendix A.1.• Updated Appendix D.1
03	January 22, 2024	Updated post-install sequence
04	March 8, 2024	Added a note to Appendix A.1 about explicitly requiring the open-source GPU driver.
05	April 30, 2024	<ul style="list-style-type: none">• Updated Appendix A.1 with new driver level.• Added Appendix A.2.• Updated Appendix D.1 with an additional workaround for the boot issue.

Table of Contents

Introduction	4
Related Documentation	4
Prerequisites	4
Access to Repositories	4
Installing the Ubuntu Server	4
Obtaining the Ubuntu Server 22.04	5
Remotely Booting the Ubuntu 22.04 Server 22.04 ISO	5
Installing Ubuntu Server	9
Appendix A: Installing Software	27
A.1 NVIDIA GPU Driver and CUDA Toolkit	27
A.2 NVIDIA MLNX_OFED	28
Appendix B: Changing the BMC Login	28
Appendix C: Automated Installation	29
C.1 Installing with Autoinstall	29
Appendix D: Platform-Dependent Workarounds	30
D.1 All Grace Platforms	30

Introduction

NVIDIA® Grace systems can run the version of Ubuntu Server that is distributed by Canonical and take advantage of the advanced Grace features.

This document explains how to install and configure Grace systems with Ubuntu Server 22.04.3.



Attention: Although it might be possible to use other Linux distributions that are related to (or inspired by) an Ubuntu Server, for example Kubuntu, **only** the Ubuntu Server 22.04 has been pre-validated with the steps in this guide for installation on Grace platforms.

Related Documentation

Refer to the [Ubuntu Server documentation](#) for more information about Ubuntu.

Prerequisites

This section lists the required (or recommended) prerequisites.

Access to Repositories

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

Installing the Ubuntu Server

Canonical provides several methods to install Ubuntu (refer to the [Ubuntu Server Documentation](#) for more information). **Before you install**, review [platform-dependent workarounds](#) to determine whether modifications are required for your environment.

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

This method installs Ubuntu Server on the Grace system remotely through facilities hosted on the BMC.

Obtaining the Ubuntu Server 22.04

Obtain the Ubuntu Server ISO for Arm® image (aarch64) and store it on your local disk. Refer to [Downloading Ubuntu Server](#) for the instructions.

If you are deploying Ubuntu on Grace-based platforms at-scale, contact Canonical for information about ongoing bug fixes, critical security patching, and long-term support. For information about using Canonical's Ubuntu images on a commercial project, read [Canonical's IP policy](#) or contact Canonical.

Remotely Booting the Ubuntu 22.04 Server 22.04 ISO

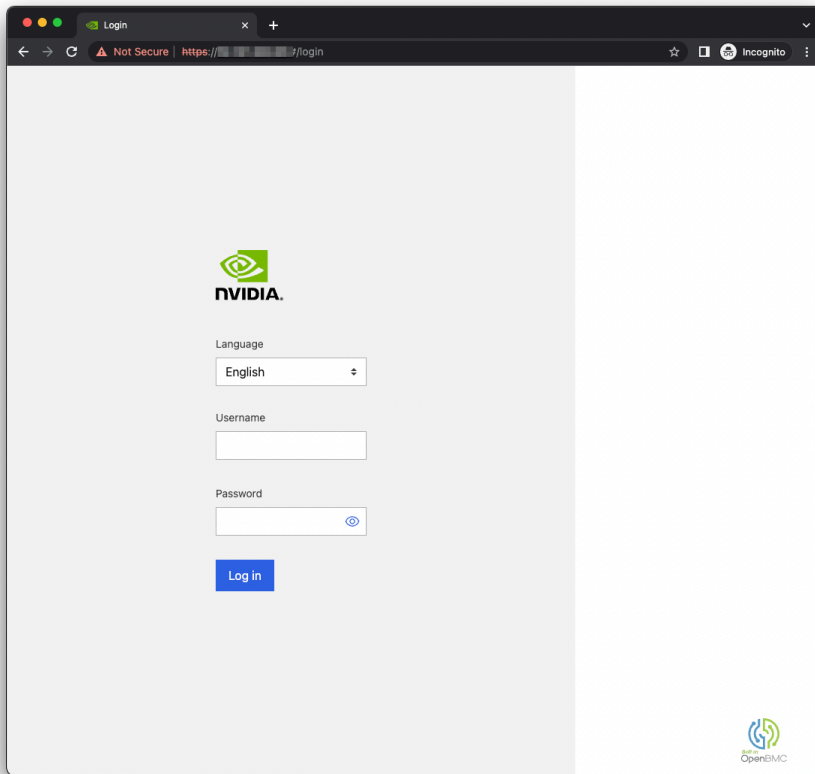


Attention: Here is some important information to know **before** you begin:

- 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.

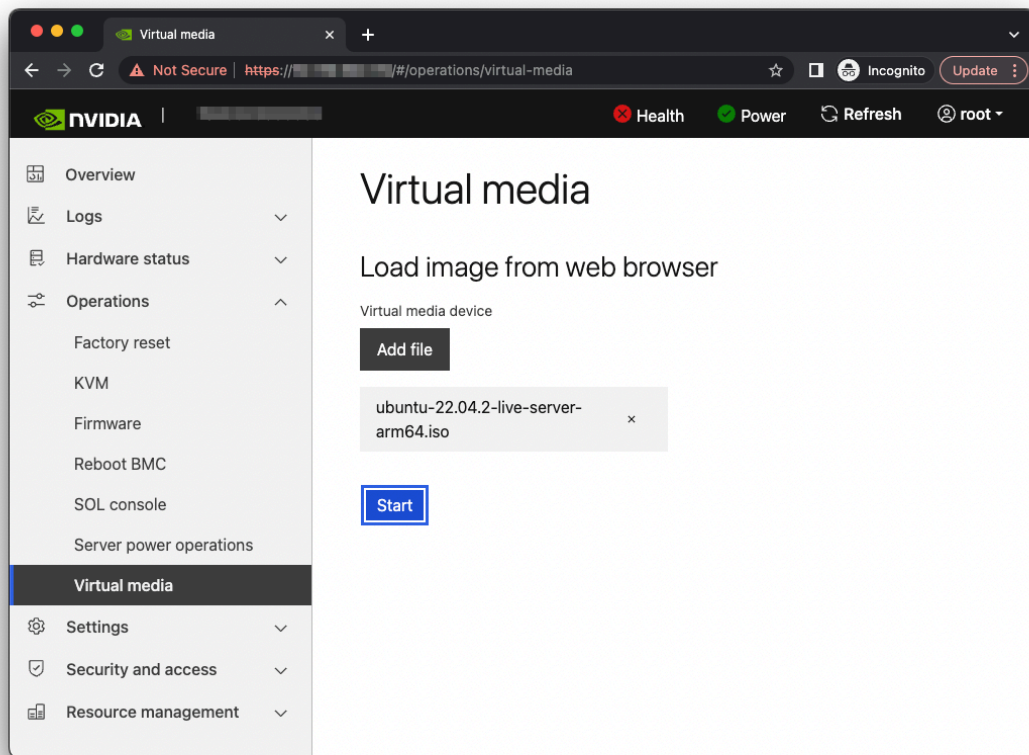
1. Connect to the BMC.
 - a. Open a browser in your LAN, navigate to `https://<BMC-IP-address>/`, and log in.

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

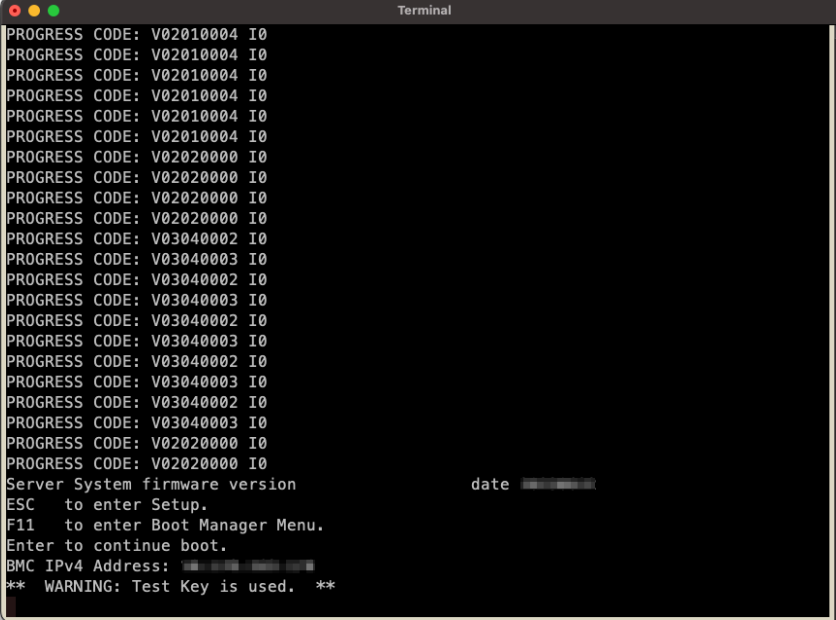


4. Boot from the virtual media.

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.

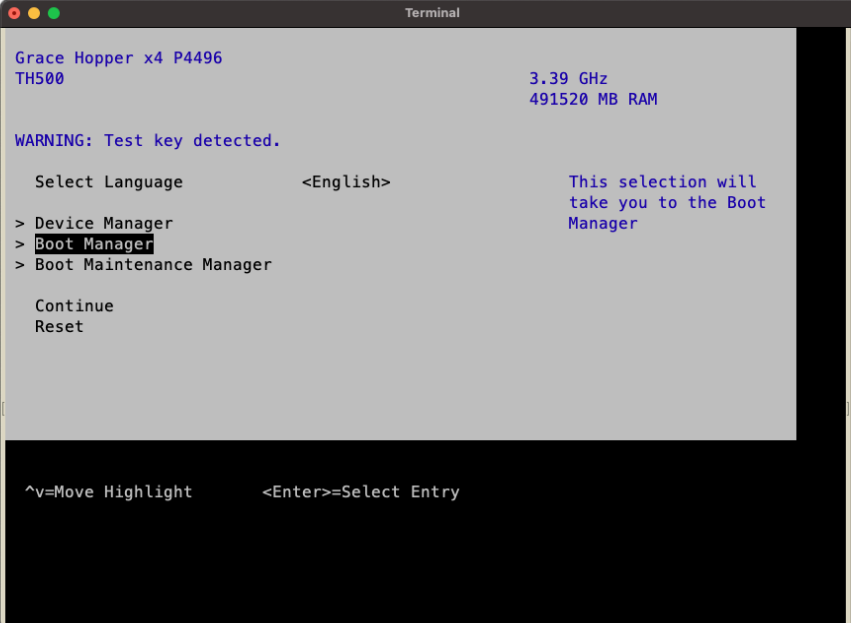
 - a. Connect to the console.
 - i. From the left hand menu, expand **Operations**, and select the SOL **console**.
 - ii. Alternatively, SSH to the BMC from a terminal emulator, log in, and at the prompt, run the `obmc-console-client` command.
 - b. To bring up the boot menu, press **Escape** or **F11**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

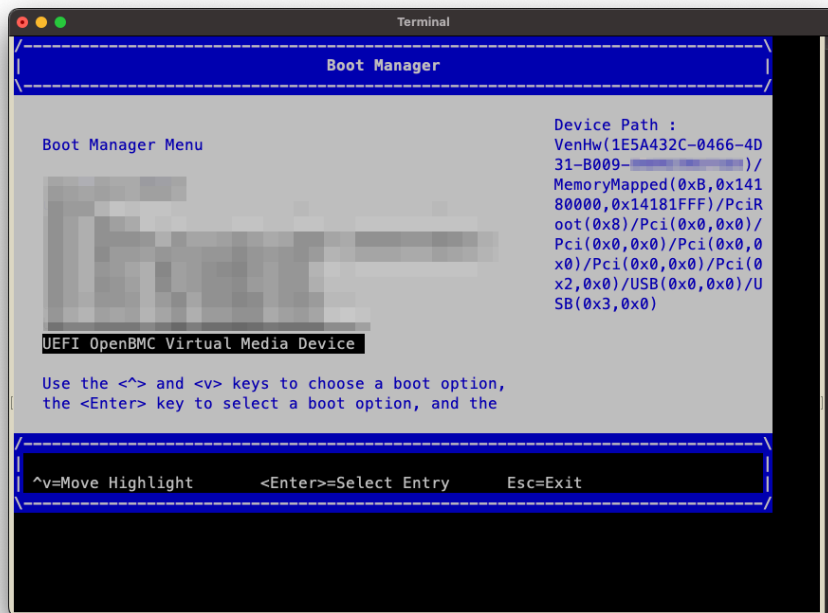
  Select Language      <English>      This selection will
                                      take you to the Boot
                                      Manager
> Device Manager
> Boot Manager
> Boot Maintenance Manager

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 Ubuntu Server](#).

Figure 5. Boot Manager




Installing Ubuntu Server

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

1. After booting the ISO image, the GRUB boot menu will appear. The menu will time out after 60 seconds. If no action is taken, it will proceed with booting the default selection (**Try or Install Ubuntu Server**).

Figure 6. GRUB Boot Menu



2. To automate the installation with an autoinstall configuration file, refer to [Installing with Autoinstall](#).
3. Select **Ubuntu Server with the HWE kernel**.
Refer to the [Ubuntu Server Installer Operations](#) for guidance on using the installer.

The performance of the Virtual Media device on Grace systems has been known to interfere with the Ubuntu installer, and this issue can cause services and cloud-init to time out.

When this is experienced, a hybrid approach to load the kernel and initial ramdisk and use the network to load the rest of the ISO is required:

- a. Host the Ubuntu Server ISO from a web server.
This can be an existing web server or a [spontaneous web server](#).
- b. Ensure that **Ubuntu Server with the HWE kernel** selected and press **e** from the GRUB boot menu.
- c. Append `ip=dhcp` and `url=http://webserver:port/path/to/UbuntuServer.iso` to the list of kernel boot parameters.

When the system has multiple network interfaces, you might be required to configure an interface with DHCP. This can be accomplished by specifying

`ip=:::::<interface_name>:dhcp::` instead of `ip=dhcp`, and this format can also be used to configure the interface with a static IP address configuration. Refer to the [nfsroot documentation](#) for more information about how to use this kernel boot parameter.

- d. To boot, click **Ctrl+X** or press **F10**.
The kernel will display output messages to the console after it begins to boot.
- e. Stop the Virtual Media device from the BMC web browser console.

The kernel will load the Ubuntu Server ISO from the network and start the installer.

4. Follow the installer prompts to configure the manual installation.
5. Select the text-based installer mode and update the installer if necessary.

Figure 7. Installer Mode Selection

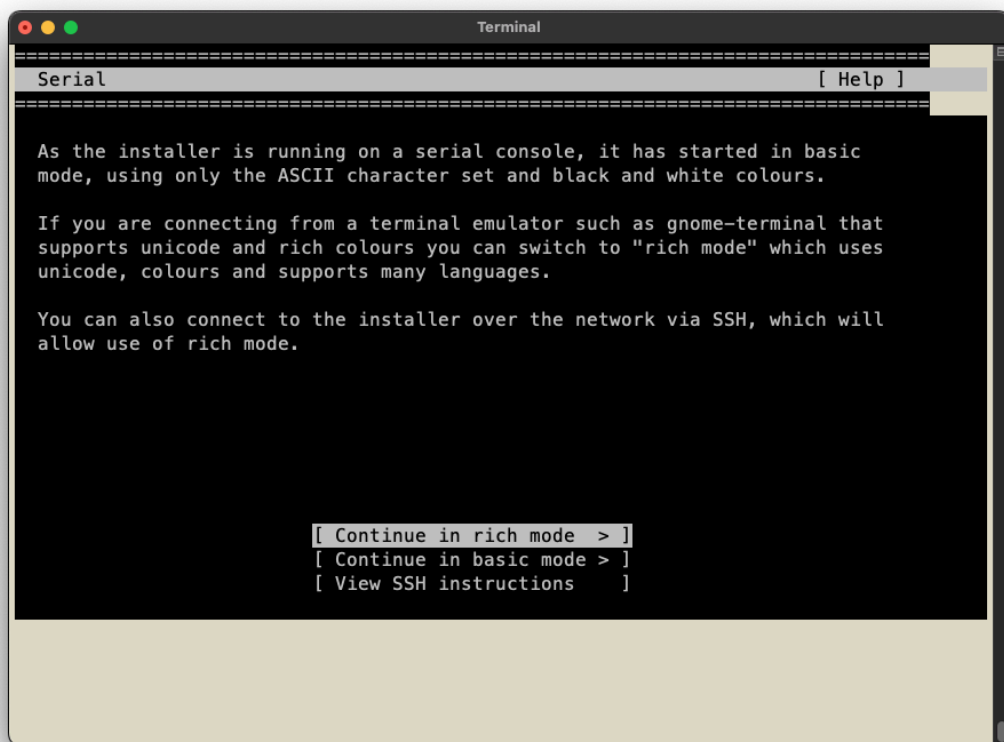

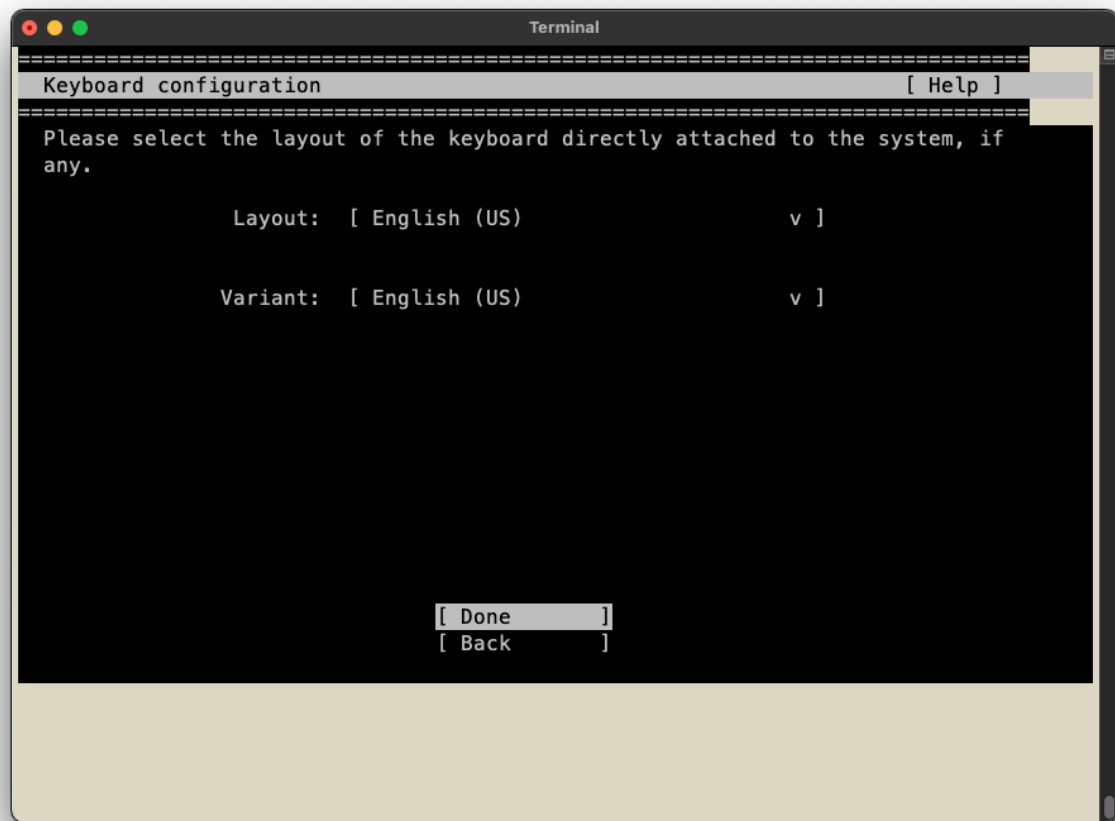


Figure 8. Installer Updates



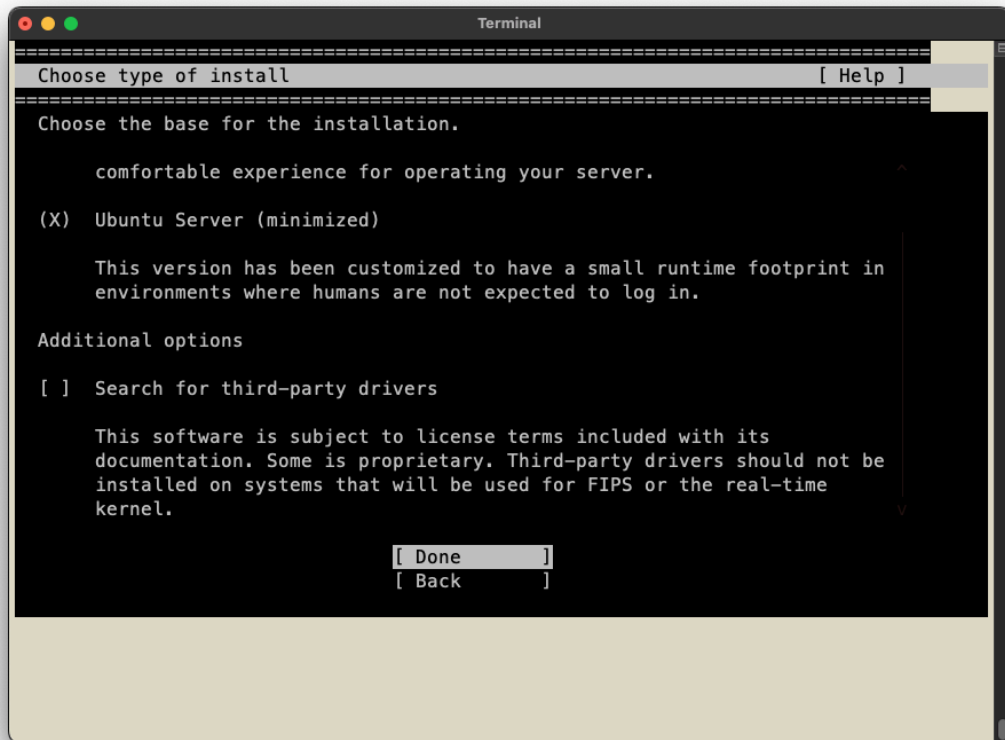
6. Select the **Keyboard** configuration.

Figure 9. Keyboard Selection



7. Select the base for the installation.

Figure 10. Installation Selection



8. Configure the network connections, proxy settings, and Ubuntu archive mirror.

Figure 11. Network Connections

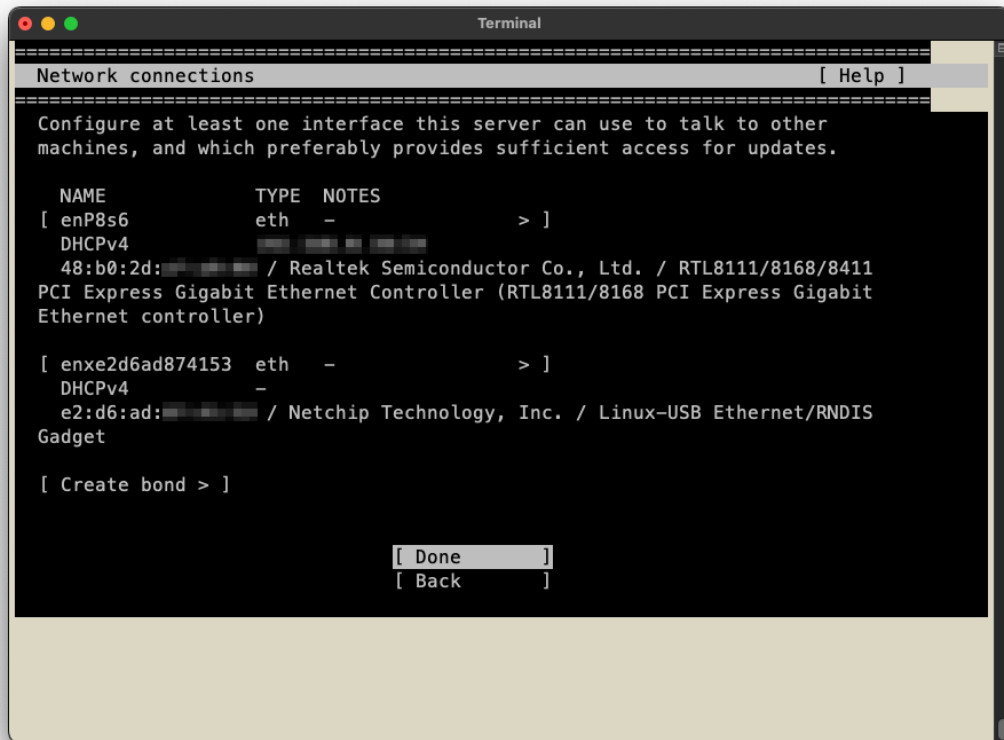


Figure 12. Proxy Configuration

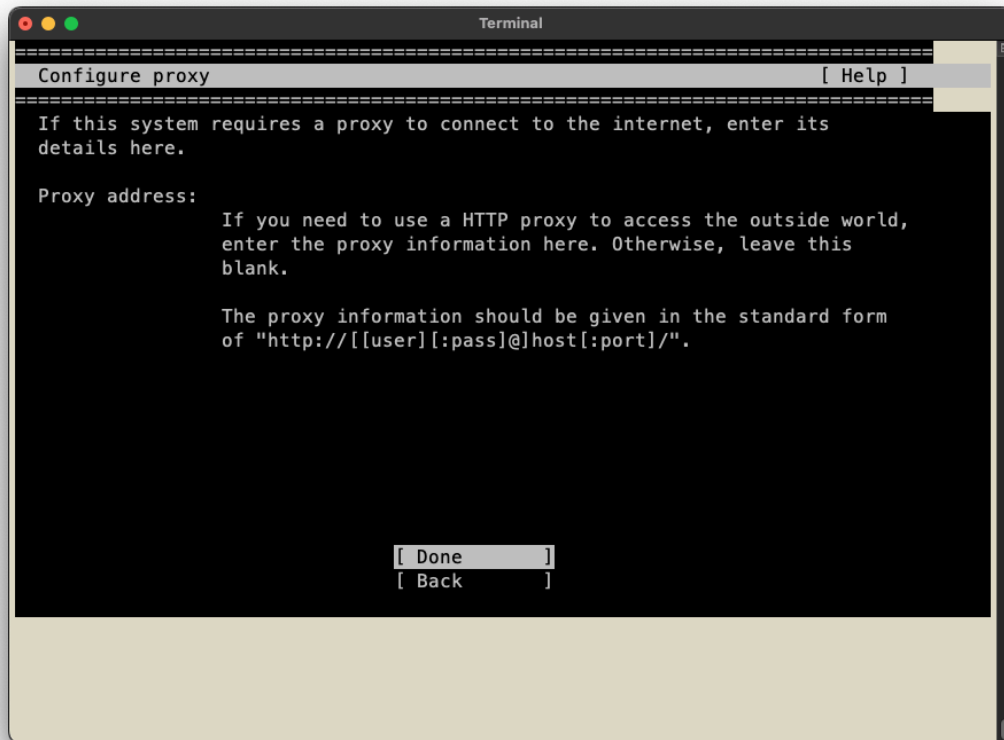



Figure 13. Ubuntu Mirror Selection



9. Select a storage configuration and confirm the destructive action.

Figure 14. Guided Storage Configuration

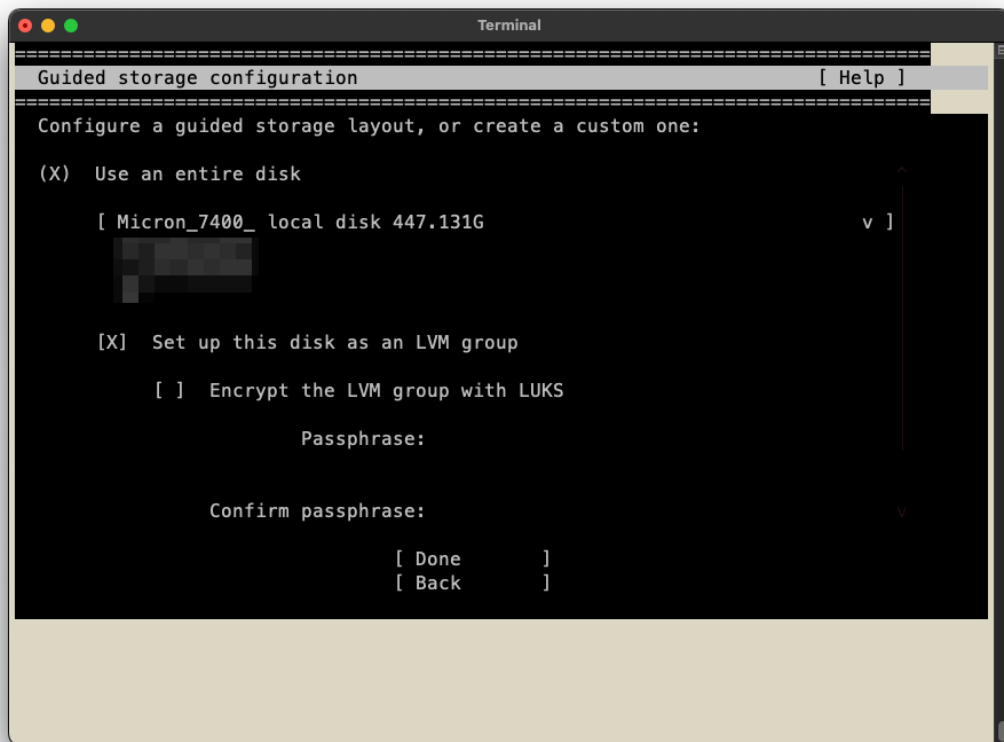


Figure 15. Storage Configuration Summary

A terminal window titled "Terminal" displays the "Storage configuration" screen. The screen has a dark background with white text. At the top, there's a header "Storage configuration" with a "[Help]" link on the right. Below this is a section titled "FILE SYSTEM SUMMARY". It contains a table with four columns: "MOUNT POINT", "SIZE", "TYPE", and "DEVICE TYPE". The table lists three entries: "/" (100.000G, new ext4, new LVM logical volume), "/boot" (2.000G, new ext4, new partition of local disk), and "/boot/efi" (1.049G, new fat32, new partition of local disk). Below the table is a section titled "AVAILABLE DEVICES". It contains a table with three columns: "DEVICE", "TYPE", and "SIZE". The table lists two entries: "ubuntu-vg (new)" (LVM volume group, 444.078G) and "free space" (344.078G). At the bottom, there are three menu items: "[Create software RAID (md) >]", "[Create volume group (LVM) >]", and a list of "[Done]", "[Reset]", and "[Back]".

```
=====
Storage configuration                                     [ Help ]
=====
FILE SYSTEM SUMMARY

MOUNT POINT      SIZE      TYPE      DEVICE TYPE
[ /              100.000G  new ext4  new LVM logical volume  > ]
[ /boot          2.000G   new ext4  new partition of local disk > ]
[ /boot/efi      1.049G   new fat32  new partition of local disk > ]

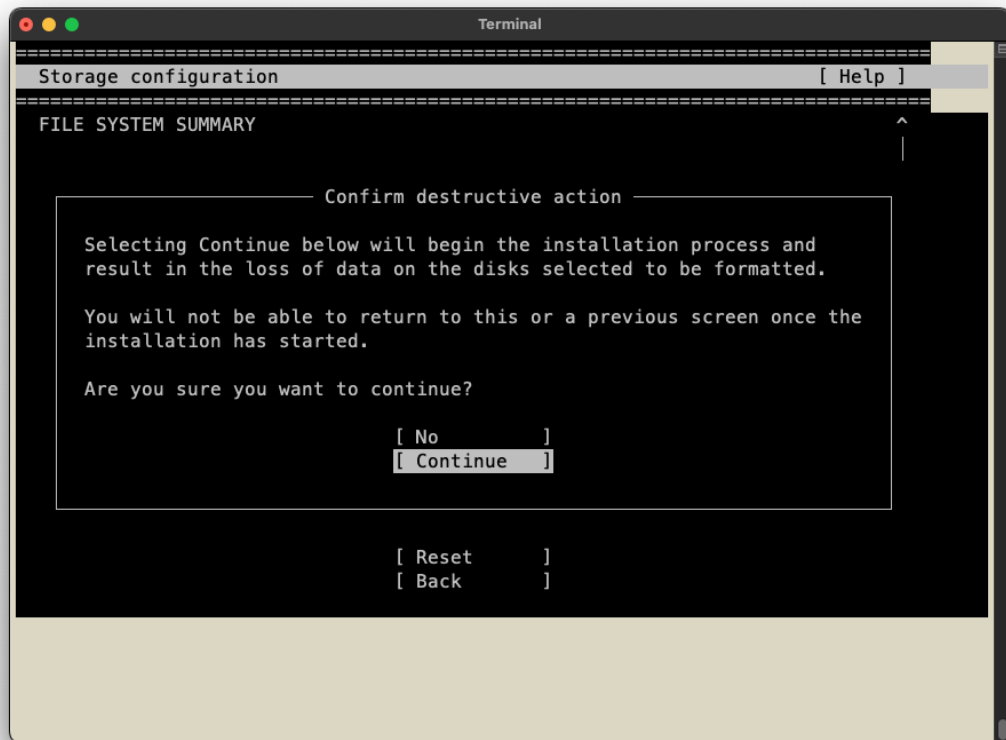
AVAILABLE DEVICES

DEVICE                                TYPE                                SIZE
[ ubuntu-vg (new)                     LVM volume group                   444.078G > ]
[ free space                           >                                344.078G > ]

[ Create software RAID (md) > ]
[ Create volume group (LVM) > ]

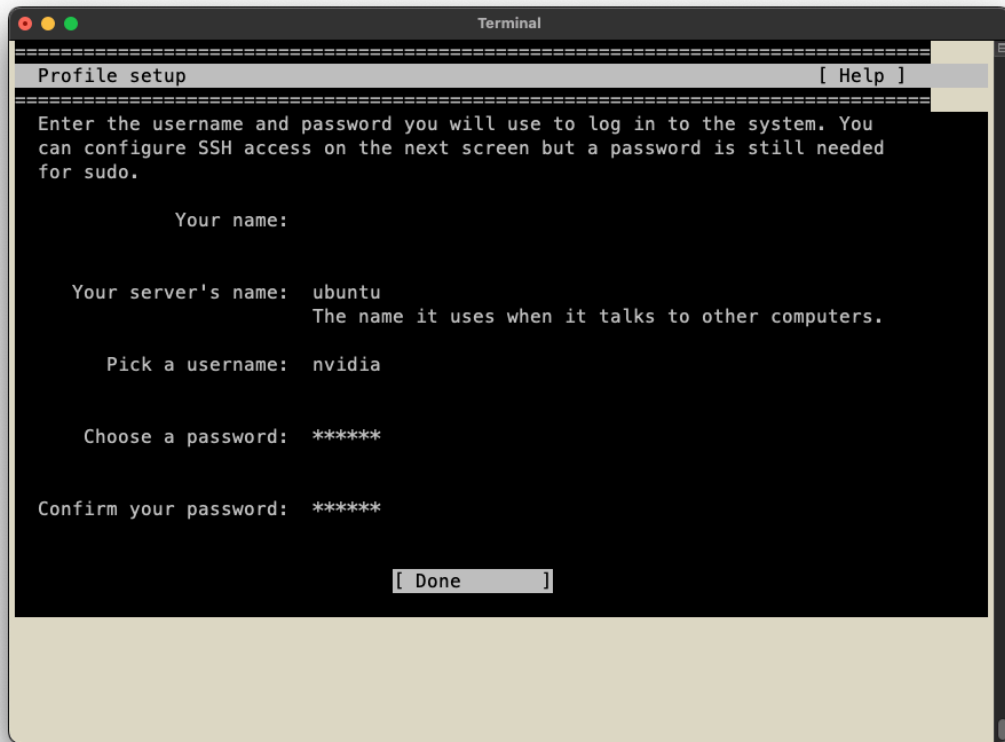
[ Done ]
[ Reset ]
[ Back ]
```

Figure 16. Storage Configuration Confirmation



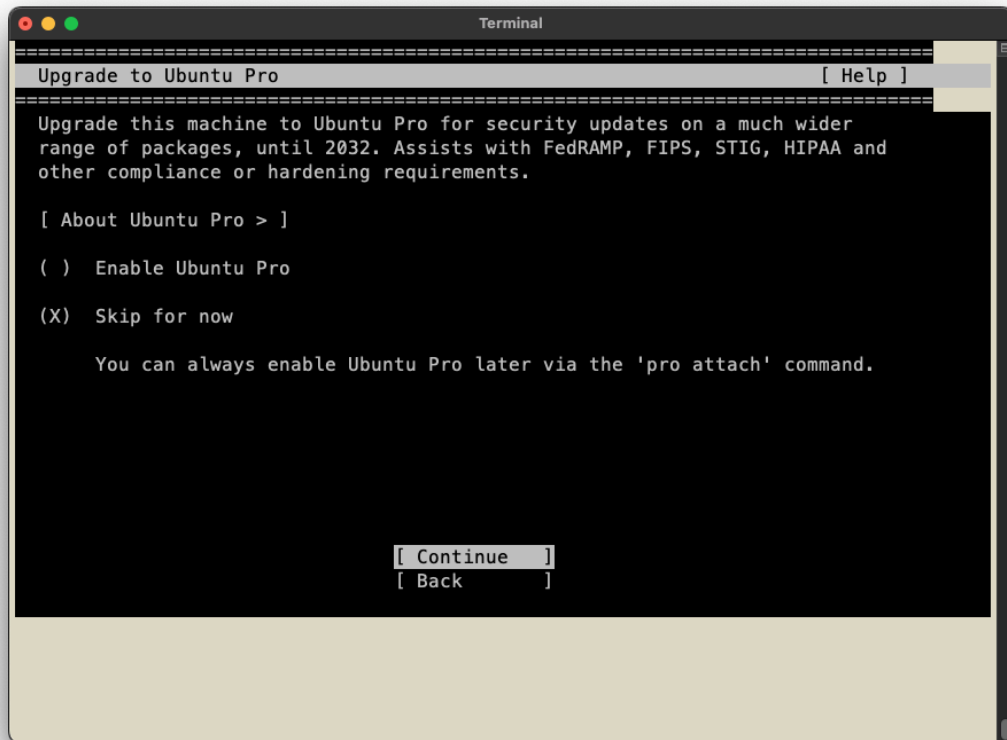
10. Select the system hostname and create a username and password.

Figure 17. System Profile Selection



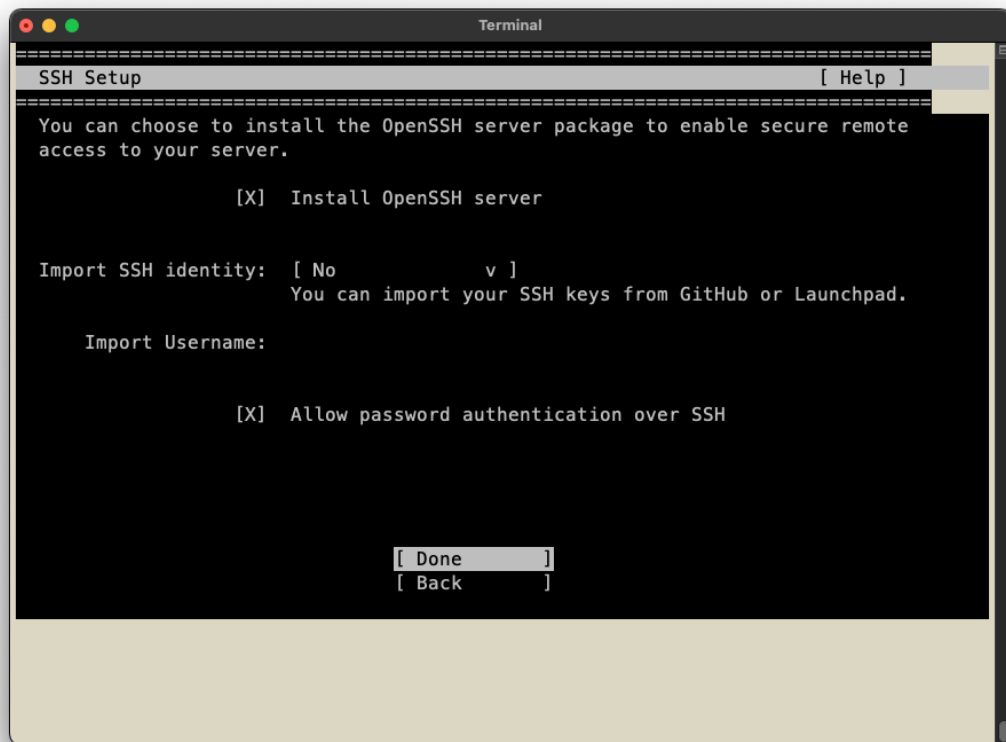
11. Select or decline the Ubuntu Pro option (refer to [Ubuntu Pro](#) for more information).

Figure 18. Ubuntu Pro Selection




12. Select the system SSH server and settings.

Figure 19. SSH Selection



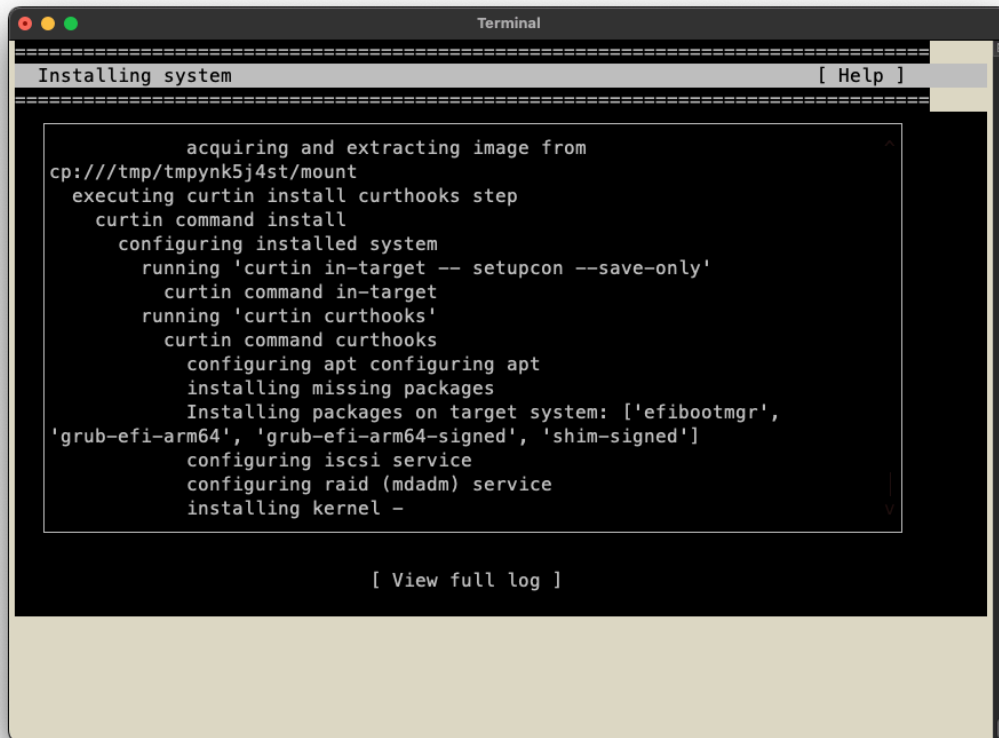
13. Select the featured Ubuntu server snaps.

Figure 20. Server Snap Selection



14. System installation begins, wait for the **Install Complete** banner and the **Reboot Now** option before proceeding to the next step.

Figure 21. Installing System

A terminal window titled "Terminal" with a dark background. The window has a title bar with red, yellow, and green window control buttons. The main content area shows a log of system installation steps. At the top, there's a header "Installing system" followed by a "[Help]" link. The log text is as follows:

```
=====
acquiring and extracting image from
cp:///tmp/tmpynk5j4st/mount
executing curtin install curthooks step
curtin command install
configuring installed system
  running 'curtin in-target -- setupcon --save-only'
  curtin command in-target
  running 'curtin curthooks'
  curtin command curthooks
  configuring apt configuring apt
  installing missing packages
  Installing packages on target system: ['efibootmgr',
'grub-efi-arm64', 'grub-efi-arm64-signed', 'shim-signed']
  configuring iscsi service
  configuring raid (mdadm) service
  installing kernel -
=====
```

At the bottom of the log area, there is a "[View full log]" link. The terminal window has a light-colored footer bar.

Figure 22. Installation Complete



15. Once the installation has completed, select **Reboot Now**.
16. After the system reboots, log in using the credentials that were specified during the installation.
17. Run the following commands to update the system and install the NVIDIA optimized Ubuntu kernel variant and reboot:

```
sudo DEBIAN_FRONTEND=noninteractive apt purge linux-image-$(uname -r)
linux-headers-$(uname -r) linux-modules-$(uname -r) -y
sudo apt update
sudo apt install linux-nvidia-64k-hwe-22.04 -y
sudo reboot now
```



Note: This sequence installs the latest available kernel. If you encounter issues after installing this kernel, it might be a known issue with a workaround described in [Appendix D](#).

18. The system will reboot.
The installation on Grace is now complete.

Appendix A: Installing Software

The Ubuntu Server uses the apt package manager to install, update, and remove packages. The utility can also be used to manage repositories. For more information about using apt, refer to [Ubuntu Package Management](#).

A.1 NVIDIA GPU Driver and CUDA Toolkit

Refer to the [NVIDIA CUDA Installation Guide for Linux](#) for detailed instructions about how to install the NVIDIA GPU driver and CUDA support for Ubuntu. 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 apt-get install linux-headers-$(uname -r)
wget
https://developer.download.nvidia.com/compute/cuda/repos/ubuntu2204/sbsa/cuda-keyring_
1.1-1_all.deb
sudo dpkg -i cuda-keyring*.deb
sudo apt-get update
sudo apt-get install cuda-toolkit-12-4 -y
sudo apt-get install nvidia-kernel-open-550 cuda-drivers-550 -y
sudo mkdir /lib/systemd/system/nvidia-persistenced.service.d
sudo dd status=none of=/lib/systemd/system/nvidia-persistenced.service.d/override.conf
<< EOF
[Service]
ExecStart=
ExecStart=/usr/bin/nvidia-persistenced --persistence-mode --verbose
[Install]
WantedBy=multi-user.target
EOF
sudo reboot now
```



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

A.2 NVIDIA MLNX_OFED

Refer to the [NVIDIA MLNX OFED Linux Drivers](#) for the current version, release notes, and the user manual. The Grace platform requires one of the following NVIDIA MLNX_OFED versions:

- MLNX_OFED v24.01-0.3.3.1 (or later)
- MLNX_OFED LTS v23.10-0.5.5.0-LTS (or later)

To install NVIDIA MLNX OFED on Ubuntu 22.04:

1. Add the NVIDIA MLNX OFED repository to your system.

```
wget -qO - http://www.mellanox.com/downloads/ofed/RPM-GPG-KEY-Mellanox | sudo gpg
--dearmor -o /usr/share/keyrings/GPG-KEY-Mellanox.gpg
echo "deb [signed-by=/usr/share/keyrings/GPG-KEY-Mellanox.gpg]
https://linux.mellanox.com/public/repo/mlnx_ofed/latest/ubuntu22.04/arm64 /" | sudo
tee /etc/apt/sources.list.d/mlnx.list
sudo apt-get update
```



Note: This sequence configures the repository for the latest available version of NVIDIA MLNX OFED. If you need a specific version, replace <latest> with the version. For example, to get the latest 23.10 LTS version, use the following URI:
https://linux.mellanox.com/public/repo/mlnx_ofed/latest-23.10/ubuntu22.04/arm64/

2. Install the NVIDIA MLNX OFED software.

```
sudo apt-get install mlnx-ofed-all -y
```

3. Update the device firmware with the version in the MLNX_OFED package.

```
sudo apt-get install mlnx-fw-updater
```

4. Reboot the system.

```
sudo reboot now
```

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 Autoinstall

Autoinstall allows you to automate the installation process by providing a configuration file with the answers to commonly asked installation questions. For Grace platforms, NVIDIA provides the following Autoinstall template:

<https://repo.download.nvidia.com/baseos/ubuntu/ubuntu-files/2204/arm64/ai/user-data.grace>

For more information about using Autoinstall files with an Ubuntu Server, refer to the [Ubuntu Autoinstall Guide](#) and the [Curtin documentation](#).

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

- User
- Timezone
- Hostname

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

1. After you customize the autoinstall file for your installation, place that file in a location that can be accessed by NFS, FTP, HTTP, or HTTPS.



Note: The Ubuntu autoinstall expects the following files to be in the location:

- user-data
- meta-data
- vendor-data.

You must rename user-data.grace to user-data and create the other two files. These files can be empty.

2. After booting from the installation medium, when the grub menu appears, press **e** to edit the grub entry and append the following to the list of kernel boot parameters:

```
autoinstall ds=nocloud-net;s=http://url.to.autoinstall.files/
```

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

The Ubuntu 22.04.3 installation media does **not** carry a patch that is required to resolve an issue with the `ast` driver that is used to interface with the AST2600 BMC. The absence of this patch can manifest a variety of issues, including kernel hangs and distorted output from the on-board VGA port. Instead of trying to workaround this issue with the Ubuntu 22.04.3 installation media, to resolve the issue, NVIDIA recommends that you use the pre-tested [22.04 LTS daily installer image](#) that contains a later kernel, which includes the `ast` driver patch. NVIDIA has worked with Canonical to pre-validate this daily installer ISO and ensure that its URL remains active until 22.04.4 is released in early 2024.

- Kernels version 6.5.0-1014 and later introduce a new `coresight_etm4x` module that might be incompatible with earlier firmware versions and ultimately prevent the system from booting. NVIDIA recommends that you update the system firmware to the latest version.

If the system still cannot boot after updating the firmware, to avoid loading the `coresight_etm4x` module, append the following to the list of kernel boot parameters:

```
module_blacklist=coresight_etm4x
```

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

© 2024 NVIDIA Corporation & Affiliates. All rights reserved.

