



NVIDIA GH200 SuperChip

MGX System Quick Start Guide

Document History

DU-11592-001_04

Version	Date	Description of Change
01	November 9, 2023	Initial release
02	December 6, 2023	<ul style="list-style-type: none">• Updated the OS and the GPU/CUDA installation steps for Ubuntu.• Added the SLES GPU/CUDA installation steps.
03	January 22, 2024	Updated the post-install sequence for Ubuntu.
04	May 17, 2024	Updated the firmware versions.

Introduction	4
Prerequisites	4
Physical Installation	5
Installing the Cable	5
First Boot	6
Locating the BMC on the Automatic DHCP Network	7
Accessing BMC Information through the System BIOS	7
Setting a Static IP for the BMC through the System BIOS	8
Accessing the BMC Network Interface	9
Installing the Operating System.....	9
Virtually Mounting the ISO File	10
Installing Ubuntu on the Grace Server	11
Installing Red Hat Enterprise Linux 9.3 on a Grace Server	12
Installing SUSE Linux Enterprise Server 15 SP5 on a Grace Server	13
Installing the NVIDIA GPU Driver and the CUDA Toolkit	15
Post-CUDA Installation Checks	16
Checking the CUDA Installation.....	16
Installing or Upgrading the ConnectX-7 Driver and Software	18
Checking the Driver/Firmware on ConnectX-7	19
Changing the ConnectX-7 Port Mode from IB to Ethernet	20
Performance Benchmarking Guide	20

Introduction

The information in this quick start guide helps you unbox and set up a Quanta Grace Hopper Unit:

Figure 1. Front and Back View of the NVIDIA MGX GH200




Table 1. Details About NVIDIA® MGX GH200

Item Number	Part Name
1	Front Panel Power
2	Front Left USB 3 (FL-USB)
3	NVMe Drive Bay
4	BMC Information Tag
5	Front Left USB 3 (FR-USB)
6	Front Panel Serial Connection
7	ConnectX-7 QSFP Ports
8	Rear Panel USB
9	Micro Display Port
10	BMC RJ45 Ethernet
11	UART/Serial
12	Rear Panel Power
13	Power Supply Cable Inserts

Prerequisites

If you plan to use your NVIDIA® MGX GH200 node outside an at-scale or planned deployment, such a lab environment or a pilot environment, we strongly recommend that you purchase the accessories in Table 2.



Note: NVIDIA does not provide these accessories.

Table 2. Required Accessories

Item	Required / Optional	Description
2x IEC C13 to NMEA 5-15P Power Cables (Standard 15A rated)	Required	Needed for system power-up.
Gigabit Ethernet Adapter (USB or PCIe)	Optional, recommended	Provides a network adapter that can be used by the operating system without configuring the bundled high-speed CX7 card. This accessory is necessary in benchtop or lab without needing a high-speed network, or during the operating system installation before you install the CX7 configuration tools. Any device supported by the Linux kernel will work such as the following: <ul style="list-style-type: none"> > USB: Devices powered by the ASIX AX88179 chipset, such as the solutions from Plugable. > PCIe: Single-port gigabit adapters built around the I210AT chipset.
Mini DisplayPort to HDMI Adapter & HDMI monitor, or VGA monitor; USB Keyboard	Optional	If the DHCP lease information is unavailable, this accessory is used to access the local system video for BIOS menus, particularly the BMC IP configuration.
MicroUSB to USB-A Cable	Optional	Can be used to access the BMC UART from a client device.
Hearing Protection	Review environmental requirements	Datacenter-class server hardware might exceed the permissible noise levels for laboratory and benchtop environments. Review the server manufacturer specification and relevant health and safety regulations for compliance information.

Physical Installation

Depending on location of install, install the server rail kit by following the manufacturer's instructions that are located inside the shipping container.



Note: Install the server rail kit into server rack by following the rack manufacturer's recommended instructions.

Installing the Cable



Note: In this section, the numbers in parenthesis refer to the callout numbers in Figure 1

1. The server comes with a ConnectX-7 networking card that is configured with dual QSFP+ ports.
 - > If your network is already QSFP compatible, plug the cables for general server networking into this port.
 - > If the final configuration of the server is going to be for copper ethernet, you need a set of the following adapters:
 - Mellanox QSFP+ to SFP+ Adapter, such as [MAM1Q00A-QSA](#).
 - SFP+ to Copper RJ45, such as [Transceiver Module for FS Switches](#)

*The links are for the examples and are **not** specific product recommendations.

2. Changing the ConnectX-7 card from QSFP mode to Ethernet Mode requires a change in the firmware at the OS level.
 - If you are converting from QSFP+ to RJ45, complete the initial setup of the machine by using a USB-to-Ethernet adapter.
 - To issue a change to the Mellanox firmware, the OFED toolset must be in a fully installed server.
 - To use Ethernet during OS installation, plug a USB-to-Ethernet converter into one of the following:
 - (2) Front Left USB
 - (8) Rear Panel USB



*When this guide was written, the Front Right USB was not functional during the system setup.

3. Connect the (10) BMC RJ45 Ethernet to the network switch that will be used to manage the Baseboard Management Console (BMC).
4. (Optional) Connect the Keyboard and Video Monitor to ports (8) and (9).

This step is optional because the BMC interface provides virtualized platforms that provide the same functionality as physical devices.

5. Connect the C13 power cables into both power supplies on the rear of the server.
6. Insert Power Supply (Cable (13)).

First Boot



Note: In this section, the numbers in parenthesis refer to the callout numbers in Figure 1

You can now boot the server using (1) or (12) power buttons. After several minutes, the server will complete different types of cycles on its own and eventually “come to life”. Approximately five minutes

later, you will be able to log into the BMC console through its http interface. You will need to use various methods available to you.

For your convenience, the manufacturer includes a helpful information tag (4) that includes the BMC's MAC address you can use to create a static IP address for the BMC or use the address to locate the BMC on your network.

You **must** ensure that the platform has, at a minimum, the versions in Table 3.

Table 3. Minimum Software Revisions

Item	Version
BMC	v3.19.00
EROT	01.03.0114.000-n01-rel-prod
Unlock key	CG1_DOT_package_2023_0829
CPLD	v07
FPGA	v8A
SBIOS	v3A11
VBIOS	96.00.84.00.02

Locating the BMC on the Automatic DHCP Network

To determine your BMC's IP address, use one of the following methods:

- > **Common way:** Open your router's web interface and find the list of issued IP addresses.

The list depends on the router model.

- > **Linux only:** Run the `arp-scan -localnet` command.
- > **Linux, MacOS, and Windows:** Download and run [Angry IP Scanner](#).
- > **Windows PowerShell:** Run the `arp -a` command.

Accessing BMC Information through the System BIOS

With a keyboard and monitor attached to the server, you can determine the BMC settings and information through the System BIOS interface.

1. During initial boot, monitor for the BIOS options screen and follow instructions to enter the setup.
2. In the System BIOS menu, press **tab** until the **Server Mgmt** option is selected.
3. Press **Enter** to select the **BMC network configuration** option.

Figure 2. BMC network configuration Option



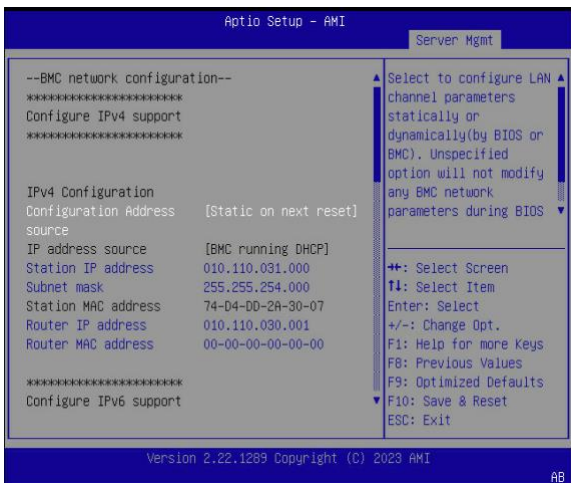
Setting a Static IP for the BMC through the System BIOS

For network configurations that require a static IP, this IP address can also be configured through the System BIOS settings. Refer to "Accessing BMC Information through the System BIOS" on page 7." for the steps. In the BMC Network Configuration menu:

1. Select **Configuration Address source**.
2. Select **Static on next reset**.

You can now change the network and router information to configure your BMC into your infrastructure.

Figure 3. Setting the Static IP Address



Accessing the BMC Network Interface

After you locate the IP address that is assigned to the server BMC, open a web browser on the same network and navigate to the IP address. For instance, if your network assigned 192.168.1.100 to your BMC's MAC address, open a browser for a computer on the same network, and enter

https://192.168.1.100 as the web address.



Note: To correctly load the page, https:// is required.

On the landing page, you will be prompted for a username and password. The default username from QCT is **admin**, and the initial password is located on the BMC information tag (4) on the front of the machine. After you enter the username, and you enter the website, you will be prompted to change the login BMC password.



Caution: Remember to record the new password because going forward, it will be needed to log in to the BMC.

Installing the Operating System

Through the BMC's Remote Control menu, the following interfaces are needed to complete the operating system installation:

- > The HTML5 Viewer, a virtual monitor, which is like a physical monitor that is connected to the box.
- > The Serial over Lan (SOL) interface.

1. Select one of the following options:
 - If you are installing the operating system (OS) image through virtual media, mount the ISO in the virtual media interface.
 - If you are installing using a physical drive, burn the ISO image in a [bootable](#) way on the USB drive and place it into the server.
2. Launch the Viewer window and the SOL Console.
3. Use the resulting interfaces to power cycle the server and begin a fresh boot sequence.
 - If the system drive **is** completely blank, it will eventually boot to the mounted media,
 - If the system drive **is not** blank, watch the BIOS init screen and select the option to choose boot media.
4. Select the available USB boot item.

NVIDIA provides guides for each of the major OS distributions that are available in the following links:



Note: Before you review the other guides, we recommend that you always start with this Quick Start Guide.

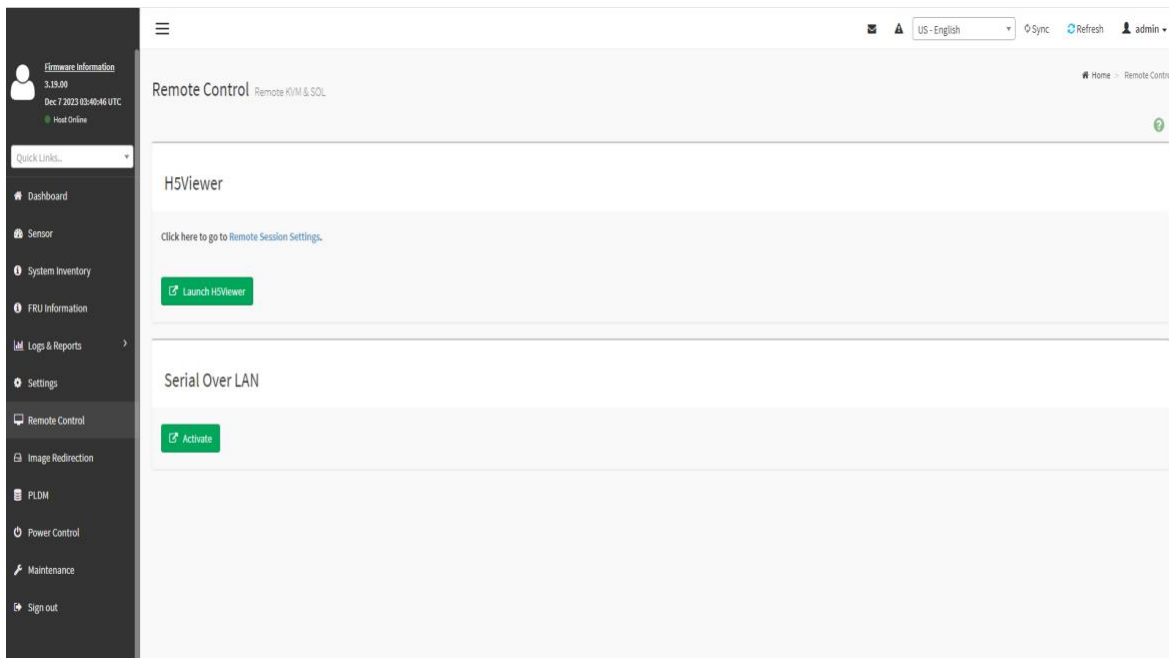
- > [NVIDIA Grace Software with Ubuntu 22.04 Installation Guide](#)
- > [NVIDIA Grace Software with Red Hat Enterprise Linux 9 Installation Guide](#)
- > [NVIDIA Grace Software with SUSE Linux Enterprise Server 15 Installation Guide](#)

Virtually Mounting the ISO File

This section provides information about how to virtually mount the ISO file.


1. On the **Remote Control** tab, click **Launch H5Viewer**.

Figure 4. The Remote Control Tab



2. Next to **CD Image**, click **browser file** and select the ISO file that you want to virtually mount.

Figure 5. Selecting the ISO File



3. To mount the iso image, select **Media Boost** and click start media.

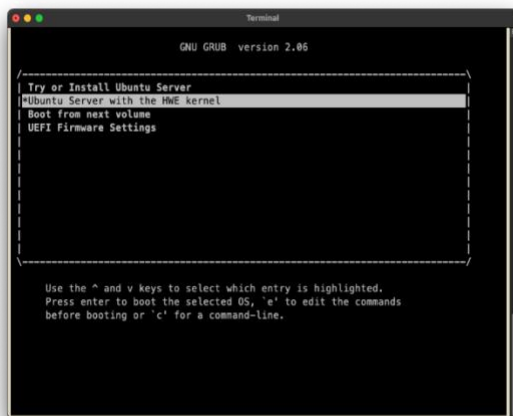


Note: During the flashing process, after the screen goes blank, leave the HTML5 Viewer window open. If the window is closed, flashing will fail.

Installing Ubuntu on the Grace Server

1. After booting the ISO image, through the BMC's virtual media mount or through a physical media drive, the GRUB boot menu will appear.
2. Select **Ubuntu Server with the HWE kernel**.

Figure 6. Ubuntu Server with the HWE Kernel Option



3. The HTML5 Viewer window will disappear and remain black for the rest of this process.

This is a known bug in the temporary kernel.

4. To complete the installation, use the Serial over LAN Console (SOL Console).
5. Complete the installations by selecting the options that are appropriate for you.
6. After the installation has completed, select **Reboot Now**.
7. After the system reboots, log in using the credentials that were specified during the installation.
8. Run the following commands to update the system and install the NVIDIA optimized Ubuntu kernel variant and reboot:

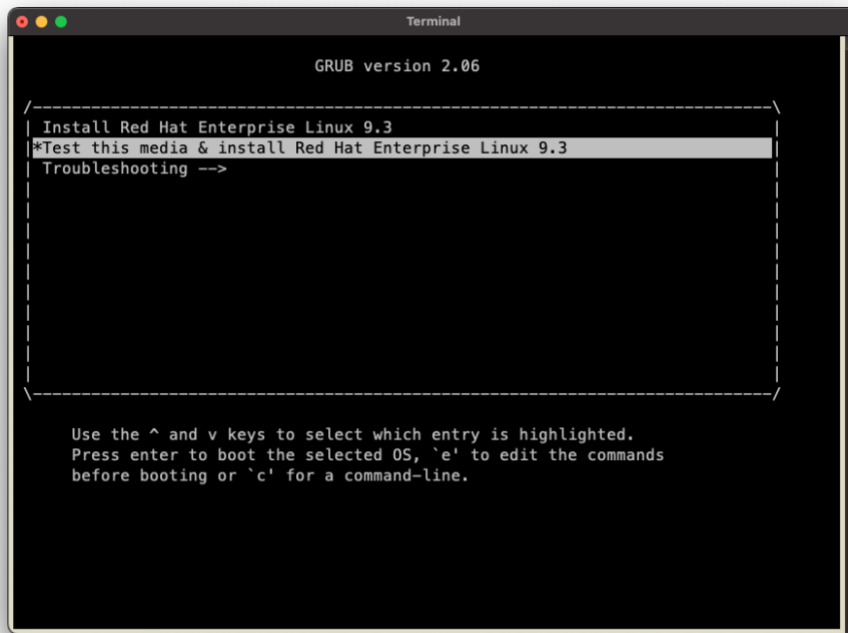
```
sudo DEBIAN_FRONTEND=noninteractive apt purge linux-image-$(uname -r) linux-headers-$(uname -r) linux-modules-$(uname -r) -y
sudo apt update
sudo apt install linux-nvidia-64k-hwe-22.04 -y
sudo reboot now
```

9. The system will reboot and the installation of Ubuntu on the Grace server is complete.

Installing Red Hat Enterprise Linux 9.3 on a Grace Server

1. After booting the ISO image, through the BMC's virtual media mount or through a physical media drive, the GRUB boot menu is displayed.
2. Select **Test this media & install Red Hat Enterprise Linux 9.3**.

Figure 7. Testing the Media



3. The HTML5 Viewer window will disappear and remain black for the rest of this process.
4. To complete the installation, use the SOL Console.
5. Complete the installations by selecting the options that are appropriate for you.
 - a. From the Software Selection menu, select the 64k kernel from the Kernel Options submenu.

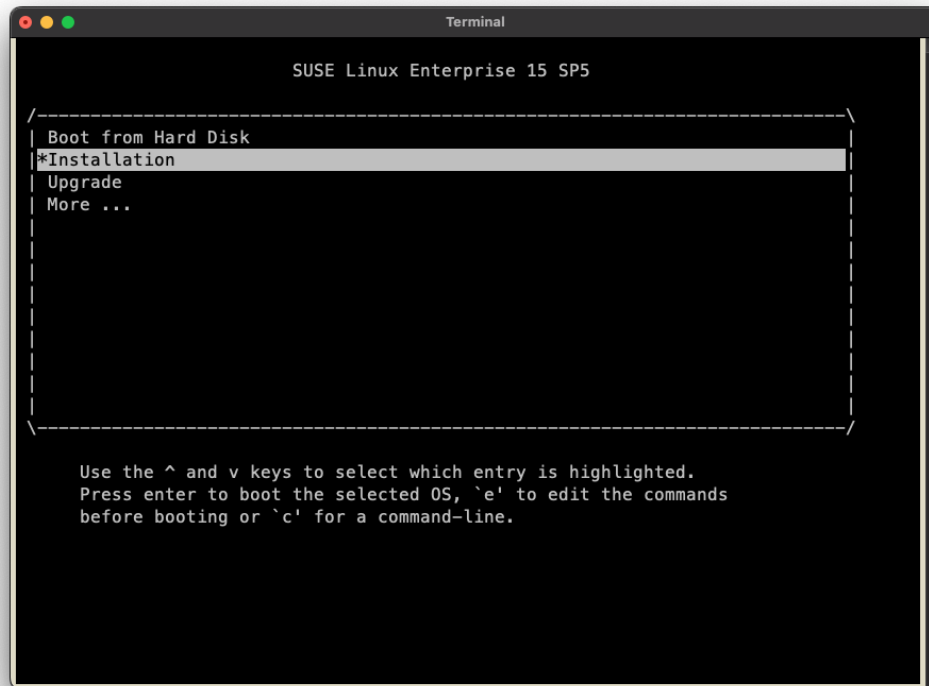
Begin the installation and allow the system to reboot.

6. After the system reboots, the installation of Red Hat Enterprise Linux 9.3 on the Grace server is complete.

Installing SUSE Linux Enterprise Server 15 SP5 on a Grace Server

1. After booting the ISO image, through the BMC's virtual media mount or through a physical media drive, the GRUB splash menu is displayed.
2. To display the GRUB boot menu, press t.

Figure 8. SLES 15 SP5 GRUB Boot Menu



3. Highlight **Installation** and press **e** to edit the boot entry.
 4. Append `modprobe.blacklist=ast` to the end of the list of kernel boot parameters.
 5. Boot the entry by clicking **Ctrl-X** or pressing **F10**.
 5. The HTML5 Viewer window will disappear and remain black for the rest of this process.
 6. To complete the installation, use the SOL Console.
 7. Complete the installations by selecting the options that are appropriate for you.
 - a. From the **Installation Settings** summary, select the **Change menu**.
 - b. Select the **Software** submenu to alter the installation to install the 64k kernel as the sole, default kernel for the system.
 - c. From the software change pane navigate to the **Search Phrase** box, and search for **kernel-64kb**.
 - d. Select the **kernel-64kb** package.
 - e. Return to the **Search Phrase** box and search for **kernel-default**.
 - f. Deselect the **kernel-default** package.
 - g. Navigate to and select **Accept** to update the planned installation.
- Begin the installation and allow the system to reboot.
8. After the system reboots, the installation of SUSE Linux Enterprise Server 15 SP5 on the Grace server is complete.

Installing the NVIDIA GPU Driver and the CUDA Toolkit

There are various supported ways to install and maintain NVIDIA's GPU Driver and CUDA SDK. Refer to the [NVIDIA CUDA Installation Guide for Linux](#) for more information about the supported installation methods. This quick start guide provides the steps for the Network/Public Repository version of the GPU Driver and CUDA installation method.



Note: The R535.129.03 driver is the **minimum** level required for the Hopper GPU.

> [Ubuntu 22.04 Instructions](#)

Run the following commands:

```
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-12-2 -y

sudo apt-get install nvidia-kernel-open-535 cuda-drivers-535 -y

sudo reboot
```

> [RHEL 9.2 Instructions](#)

Run the following commands:

```
sudo dnf install kernel-64k-devel-$(uname -r | sed 's/+64k//g')

sudo dnf install kernel-headers-$(uname -r | sed 's/+64k//g')

sudo dnf install https://dl.fedoraproject.org/pub/epel/epel-release-latest-9.noarch.rpm

sudo dnf config-manager --add-repo
https://developer.download.nvidia.com/compute/cuda/repos/rhel9/sbsa/cuda-rhel9.repo

sudo dnf clean expire-cache

sudo dnf install cuda-toolkit-12-2
```

```
sudo dnf module install nvidia-driver:535-open
sudo systemctl status nvidia-persistenced
sudo reboot
```

> [SLES 15 SP5 Instructions](#)

Run the following commands:

```
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
```

Post-CUDA Installation Checks

For most users, we recommend that you follow post-installation steps in [Post-Installation Actions..](#)

Checking the CUDA Installation

After OS, Driver, and CUDA are installed, to verify that everything is correctly installed and initialized, we recommend that you do a complete system power cycle.

To do a quick check and determine whether Grace Hopper is ready to call the NVIDIA System Management Interface check, run the `nvidia-smi` command.

Figure 9. The `nvidia-smi` Command Output

```

[nvidia@host ~]$ nvidia-smi
Wed Nov  8 20:07:46 2023

+-----+
| NVIDIA-SMI 535.129.03                Driver Version: 535.129.03   CUDA Version: 12.2   |
+-----+-----+-----+-----+-----+-----+
| GPU  Name           Persistence-M   Bus-Id        Disp.A    Volatile Uncorr. ECC  |
| Fan  Temp            Perf             Pwr:Usage/Cap     Memory-Usage  GPU-Util  Compute M.  |
|              Perf                Pwr:Usage/Cap     Memory-Usage  GPU-Util  Compute M.  |
|-----+-----+-----+-----+-----+-----+
|   0   GH200  120GB                Off          00000009:01:00.0 Off          |
| N/A   33C           P0              128W / 900W      4MiB / 97871MiB      0%          Default    |
|              Perf                Pwr:Usage/Cap     Memory-Usage  GPU-Util  Compute M.  |
|-----+-----+-----+-----+-----+
| Processes:                               GPU Memory  |
|  GPU   GI    CI          PID    Type   Process name          Usage        |
|-----+-----+-----+-----+-----+
| No running processes found                |
+-----+

```

To check and determine whether the CPU and GPU memory subsystems are up and functional, run the following commands.

```

sudo nvidia-persistenced
lsmem

```

Figure 10. The `sudo nvidia-persistenced` Command Output

```

nvidia@gctcg1:~$ sudo nvidia-persistenced
nvidia@gctcg1:~$ lsmem
RANGE                                SIZE  STATE  REMOVABLE      BLOCK
0x0000000080000000-0x000000787fffffff 480G  online  yes            4-963
0x0000400000000000-0x00004017bfffffff  95G  online  yes 131072-131261

Memory block size:      512M
Total online memory:    575G
Total offline memory:    0B

```

Installing or Upgrading the ConnectX-7 Driver and Software



Note: In the future, the driver version will be 23.10 driver, but the currently available version is 23.07.

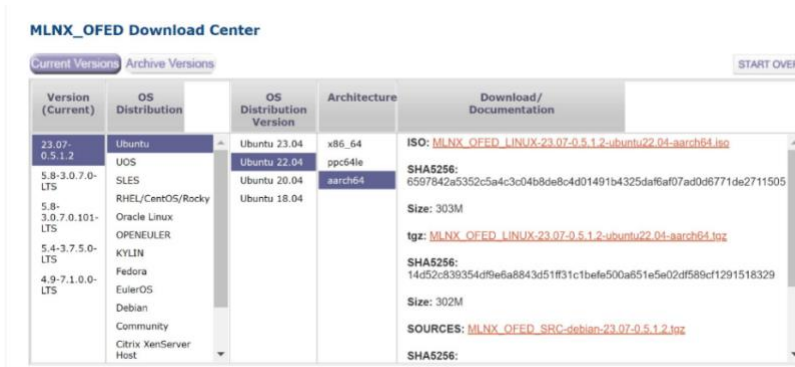
This is a more concise version of some detailed installation instructions that can be found in [NVIDIA Networking Linux Driver Installation](#).

1. Obtain the appropriate driver from [NVIDIA Network Download](#).

Table 4. Driver Information

Component	Description
Version	23.07
OS Distribution	<Ubuntu, RHEL, SLES>
OS Distribution Version	<22.04, 9.2, ?? >
Architecture	Aarch64
Download/Documentation	Select Link for ISO

Figure 11. The NVIDIA Network Download Site



2. Download the ISO file onto the server.

For example, https://www.mellanox.com/downloads/ofed/MLNX_EN-23.07-0.5.1.2/mlnx-en-23.07-0.5.1.2-ubuntu22.04-aarch64.iso.

3. Create a mount point for the ISO image.

For example, run `sudo mkdir -p /mnt/iso`.

4. Mount the ISO image.

```
sudo mount -o loop <ISO IMAGE FILE> /mnt/iso
```

5. Use the `mlnxofedinstall` file in the image mount to install the driver and tools.

```
cd /mnt/iso
sudo ./mlnxofedinstall
```

6. Clean up after the installation.

```
cd ~
sudo umount /mnt/iso
rm <ISO IMAGE FILE>
```

7. Complete one of the following tasks:

- Hard reboot the machine.
- Restart the ConnectX-7 device.

```
sudo /etc/init.d/openibd restart
```

Checking the Driver/Firmware on ConnectX-7



Note: The Firmware will not show an update before the system has been rebooted or card restarted.

1. To check the driver or firmware on ConnectX-7, run the following command.

```
sudo ibdev2netdev -v
```

Here is the sample output:

```
nvidia@qctcg1:/mnt/iso$ sudo ibdev2netdev -v
0000:01:00.1 mlx5_0 (MT4129 - MCX755106AC-HEAT) NVIDIA ConnectX-7 HHL adapter Card, 200GbE / NDR200 IB, Dual-port QSFP112, PCIe 5.0 x16 with x16 PCIe extension option, Crypto and Secure Boot fw 28.36.1010 port 1 (DOWN) =
=> enp1s0f1np1 (Down)
```

2. Use the `net-tools` package `ifconfig` to get the name of the CX-7 device.

```
ethtool -I <nic name>
```

For example, `ethtool -i enp1s0f1np1`

Figure 12. Outputs Showing With and Without Reboot

Without Reboot	With Reboot
<pre>nvidia@qctcg1:/mnt/iso\$ ofed_info -s MLNX_OFED_LINUX-23.10-0.0.1.0: nvidia@qctcg1:/mnt/iso\$ ethtool -i enp1s0f1np1 driver: mlx5_core version: 23.10-0.0.1 firmware-version: 28.36.1010 (MT_0000001045) expansion-rom-version: bus-info: 0000:01:00.1 supports-statistics: yes supports-test: yes supports-eprom-access: no supports-register-dump: no supports-priv-flags: yes nvidia@qctcg1:/mnt/iso\$</pre> <p>Firmware version is still 28.36.</p>	<pre>nvidia@qctcg1:~\$ ofed_info -s MLNX_OFED_LINUX-23.10-0.0.1.0: nvidia@qctcg1:~\$ ethtool -i enp1s0f1np1 driver: mlx5_core version: 23.10-0.0.1 firmware-version: 28.38.1002 (MT_0000001045) expansion-rom-version: bus-info: 0000:01:00.1 supports-statistics: yes supports-test: yes supports-eprom-access: no supports-register-dump: no supports-priv-flags: yes</pre> <p>Firmware version is now 28.38.</p>

Changing the ConnectX-7 Port Mode from IB to Ethernet

To change the ConnectX-7 port mode from IB to Ethernet, run the following command:

```
sudo mst start
sudo mst status -v
# Take the correct pciconf ID related to you NIC from above command
sudo mlxconfig -d /dev/mst/mt4129_pciconf<ID> set LINK_TYPE_P1=ETH LINK_TYPE_P2=ETH
```

where **mt4129** is specific to the ConnectX-7 device. This number might change depending on the version of the installed network card.

Performance Benchmarking Guide

A selection of tests to stress the GraceHopper architecture (refer to NVOnline: **1104321** for more information).

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, 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.

