



DGX Software with CentOS 8

Installation Guide

Table of Contents

Chapter 1. Introduction.....	1
1.1. Related Documentation.....	1
1.2. Prerequisites.....	1
1.2.1. Access to Repositories.....	2
1.2.1.1. NVIDIA Repositories.....	2
1.2.1.2. CentOS Repositories.....	2
1.2.2. Network File System.....	2
1.2.3. BMC Password.....	2
Chapter 2. Installing CentOS 8.....	3
2.1. Obtaining CentOS 8.....	3
2.2. Booting CentOS 8 ISO Locally.....	3
2.3. Booting the CentOS 8 ISO Remotely.....	4
2.3.1. Booting the ISO Image on the DGX-1 Remotely.....	4
2.3.2. Booting the ISO Image on the DGX-2 or DGX A100 Remotely.....	5
2.4. Installing CentOS.....	7
2.4.1. Installing with Kickstart.....	11
2.4.2. Disk Partitioning for DGX-1, DGX Station, and DGX Station A100.....	12
2.4.3. Disk Partitioning with Encryption for DGX-1, DGX Station, and DGX Station A100.....	13
2.4.4. Disk Partitioning for DGX-2 and DGX A100.....	15
2.4.5. Disk Partitioning with Encryption for DGX-2 and DGX A100.....	16
Chapter 3. Installing the DGX Software.....	18
3.1. Configuring a System Proxy.....	18
3.2. Enabling the DGX Software Repository.....	18
3.3. Installing Required Components.....	19
3.4. Installing Optional Components.....	22
3.5. Installing the Optional NVIDIA Desktop Theme.....	23
Chapter 4. Using the NVIDIA Mellanox InfiniBand Drivers.....	24
4.1. Determining the MLNX_OFED Version to Install.....	24
4.2. Installing the NVIDIA Mellanox InfiniBand Drivers.....	24
4.3. Updating the NVIDIA Mellanox InfiniBand Drivers.....	26
Appendix A. Changing the BMC Login.....	28
A.1. Changing the BMC Login on the DGX-1.....	28
A.2. Changing the BMC Login on the DGX-2 or DGX A100.....	29
Appendix B. Installing the Mellanox ConnectX-6 Firmware.....	30

Appendix C. Changing NVIDIA Driver Branches..... 31

Chapter 1. Introduction

The NVIDIA® DGX™ systems are shipped with DGX™ OS which incorporates the NVIDIA DGX software stack built upon the Ubuntu Linux distribution. Instead of running the Ubuntu distribution, you can run CentOS on the DGX system and still take advantage of the advanced DGX features.

This document explains how to install and configure the NVIDIA DGX software stack on DGX systems installed with CentOS.



Important: NVIDIA acknowledges the wide use of CentOS and understands that it is a community-developed derivative of the NVIDIA supported Red Hat Enterprise Linux. Support for CentOS is available directly from the CentOS community. NVIDIA ensures that NVIDIA provided software runs on tested CentOS versions and will try to identify and correct issues related to NVIDIA provided software.



Note: While it may be possible to use other derived Linux distributions besides CentOS, not all have been tested and qualified by NVIDIA. Refer to the [DGX Software for Red Hat Enterprise Linux 8 Release Notes](#) for the list of tested and qualified software and Linux distributions.

1.1. Related Documentation

- ▶ [NVIDIA DGX Software for Red Hat Enterprise Linux 8 - Release Notes](#)
- ▶ [NVIDIA DGX-1 User Guide](#)
- ▶ [NVIDIA DGX-2 User Guide](#)
- ▶ [NVIDIA DGX A100 User Guide](#)
- ▶ [NVIDIA DGX Station User Guide](#)

1.2. Prerequisites

The following are required (or recommended where indicated).

1.2.1. Access to Repositories

The repositories can be accessed from the internet.

If you are using a proxy server, then follow the instructions in the section [Configuring a System Proxy](#) to make sure the system can access the necessary URIs.

1.2.1.1. NVIDIA Repositories

After installing CentOS on the DGX system, you must enable the NVIDIA DGX software repository (<https://repo.download.nvidia.com>). The repository includes the NVIDIA drivers and software for supporting DGX systems.

See the section [Enabling the Repositories](#) for instructions on how to enable the repository.

1.2.1.2. CentOS Repositories

Installation of the DGX Software over CentOS requires access to several additional repositories.

- ▶ **CentOS Software Collections Repository:** `centos-release-scl`
This repository is required by the NVSM tool for Python 3.
- ▶ **CentOS Testing Repository:** `centos-scl0-rh-testing`
This repository is required by the NVSM tool for Python 3.

1.2.2. Network File System

On DGX servers, the data drives are meant to be used as a cache. DGX Station users can follow the same usage, or can alternatively opt to use these drives for storage. When using the data drives as cache, a network file system (NFS) is recommended to take advantage of the cache file system provided by the DGX software stack.

1.2.3. BMC Password

The DGX BMC comes with default login credentials as specified in [Appendix B: Changing the BMC Login](#).



Important:

NVIDIA recommends disabling the default username and creating a unique BMC username and strong password as soon as possible. Refer to [Appendix B: Changing the BMC Login](#) for instructions.

Chapter 2. Installing CentOS 8

There are several methods for installing CentOS as described in the [CentOS Installation Guide \(https://docs.centos.org/en-US/centos/install-guide/\)](https://docs.centos.org/en-US/centos/install-guide/).

See the [DGX Software for Red Hat Enterprise Linux Release Notes](#) for the version of Red Hat Enterprise Linux 8 that is qualified and tested for use with the DGX software.

For convenience, this section describes how to install CentOS using the Quick Install method, and shows when to reclaim disk space in the process. It describes a minimal installation. If you have a preferred method for installing CentOS, then you can skip this section but be sure to reclaim disk space occupied by the existing Ubuntu installation.

The interactive method described here installs CentOS on DGX using a connected monitor and keyboard and USB stick with the ISO image, or remotely through the remote console of the BMC.

2.1. Obtaining CentOS 8

Obtain the CentOS 8 ISO image and store on your local disk or create a boot USB drive formatted for UEFI. See [Downloading CentOS \(https://docs.centos.org/en-US/centos/install-guide/downloading/#chap-download\)](https://docs.centos.org/en-US/centos/install-guide/downloading/#chap-download) for instructions.



Important: Refer to the release notes for any critical information regarding supported releases or dependencies.

2.2. Booting CentOS 8 ISO Locally

1. Plug the USB flash drive containing the CentOS 8 ISO image into the DGX.
2. Connect a monitor and keyboard directly to the DGX.
3. Boot the system and press **F11** when the NVIDIA logo appears to get to the boot menu.
4. Select the USB volume name that corresponds to the inserted USB flash drive, and boot the system from it.
5. Follow the instructions at [Installing CentOS](#)

2.3. Booting the CentOS 8 ISO Remotely

Skip this chapter if you are using a monitor and keyboard for installing locally, or if you are installing on a DGX Station. The DGX Station **cannot** be booted remotely.

2.3.1. Booting the ISO Image on the DGX-1 Remotely

Skip this chapter if you are using a monitor and keyboard for installing locally.

- ▶ For instructions applicable to the NVIDIA DGX-2 or DGX A100, see [Booting the ISO Image on the DGX-2 or DGX A100 Remotely](#)

1. Connect to the BMC and change user privileges.
 - a). Open a Java-enabled web browser within your LAN and go to `http://<BMC-ip-address>/`, then log in.
Use Firefox or Internet Explorer. Google Chrome is not officially supported by the BMC.
 - b). From the top menu, click **Configuration** and then select **Users**.
 - c). Select the user name that you created for the BMC, then click **Modify User**.
 - d). In the **Modify User** dialog, select the **VMedia** checkbox to add it to the extended privileges for the user, then click **Modify**.

2. Set up the ISO image as virtual media and reboot the system.
 - a). From the top menu, click **Remote Control** and select **Remote KVM**.

b). Click **Launch KVM** to open the remote console.
 - c). From the top menu bar in the KVM window, click **Browse File** and select the ISO image, then click **Start Media**.

The CD image should now be connected.

- d). From the top menu bar in the KVM window, click **Power** and then select **Reset Server**.
3. Boot the CD ROM image.

Typically, the default boot order does not boot the CDROM image. This can be changed in the BIOS or as a one-time option in the boot menu. To bring up the boot menu, press F11 at the beginning of the boot process. Pressing **F11** will display Show Boot Options at the top of the virtual display before entering the boot menu. Use the 'soft' keyboard (Menu → Keyboard Layout → SoftKeyboard → <Language>) to bring up a virtual keyboard if pressing the physical key has no effect.

- a). In the boot menu, select **UEFI: AMI Virtual CDROM 1.00** as the boot device and then press **ENTER**.
- b). Follow the instructions at [Installing CentOS](#).

2.3.2. Booting the ISO Image on the DGX-2 or DGX A100 Remotely


Skip this chapter if you are using a monitor and keyboard for installing locally.

1. Connect to the BMC and ensure the required user privileges are set.
 - a). Open a browser within your LAN and go to `https://<BMC-ip-address>/`, then log in.
 - b). From the left-side menu, click **Settings** and then select **User Management**.
 - c). Click the card with the user name that you created for the BMC.
 - d). In the **User Management Configuration** dialog, make sure the **VMedia Access** checkbox is selected, then click **Save**.
2. Set up the ISO image as virtual media.
 - a). From the left-side menu, click **Remote Control**.

The screenshot displays the NVIDIA DGX A100 BMC web interface. The top left shows system information: Apr 24 2020 20:19:22 CST, FW: 9.11.12, IP: 172.18.205.154, MAC: 9CFF35C2587E, Chassis Part: 929-22987-2510-000, Chassis SN: 1234567890123, Host Online, and Chassis Identify LED. The dashboard includes four main cards: Power-On Hours (89 d 7 hrs), Pending Deassertions (181), Access Logs (40), and GPU Info (8). Below these is a Firmware table with columns for BMC, BIOS, MB FPGA, MID FPGA, GB FPGA, and PSU 0-5. The Sensor Monitoring section shows a green status 'All sensors are good now!' and a 'Currently recovered' indicator.

Firmware	BMC	BIOS	MB FPGA	MID FPGA	GB FPGA	PSU 0	PSU 1	PSU 2	PSU 3	PSU 4	PSU 5
Primary	9.11.12*	0.00	9.00.02	9.00.02	3.01	1.04	1.04	1.04	1.04	1.04	1.04
Secondary	9.08.11	0.10*	N/A	N/A	N/A	1.04	1.04	1.04	1.04	1.04	1.04
Communication	N/A	N/A	N/A	N/A	N/A	1.04	1.04	1.04	1.04	1.04	1.04

- b). Select **Launch KVM**.



- c). From the top menu bar in the KVM window, click **Browse File** and select the ISO image, then click **Start Media**.

The CD image should now be connected.

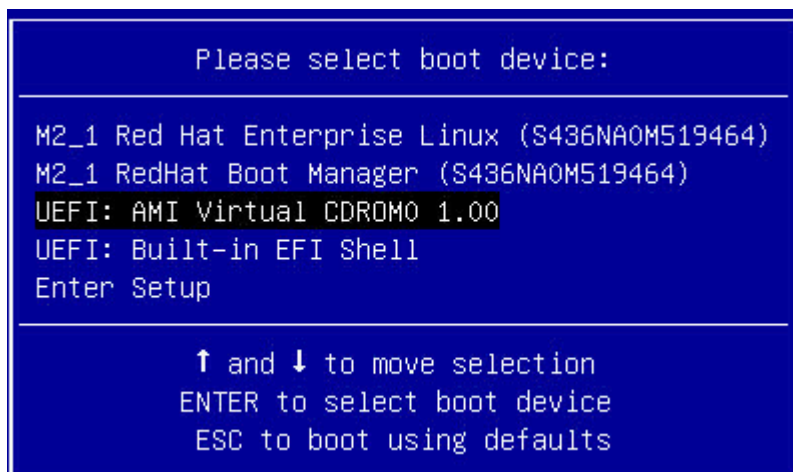
- d). From the top menu bar in the KVM window, click **Power** and then select **Reset Server**.
3. Boot from the virtual media.

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

- a). To bring up the boot menu, press **F11** at the beginning of the boot process.

Pressing **F11** will display **Entering Boot Menu** in the virtual display before entering the boot menu.

- b). In the boot menu, select **UEFI: AMI Virtual CDROM0 1.00** as the boot device and then press **ENTER**.



If you are booting from the DGX-2, the boot device is labeled **UEFI: Virtual CDROM 1.00**.

c). Follow the instructions at [Installing CentOS](#)

2.4. Installing CentOS

This section assumes you have booted the CentOS ISO image. For DGX Station, this can only be performed [locally](#); for all other platforms, this can be performed [locally](#) or [remotely](#).


If you wish to automate the installation with a kickstart file, see the section [Installing with Kickstart](#)

1. After booting the ISO image, the GRUB menu for the installer appears.



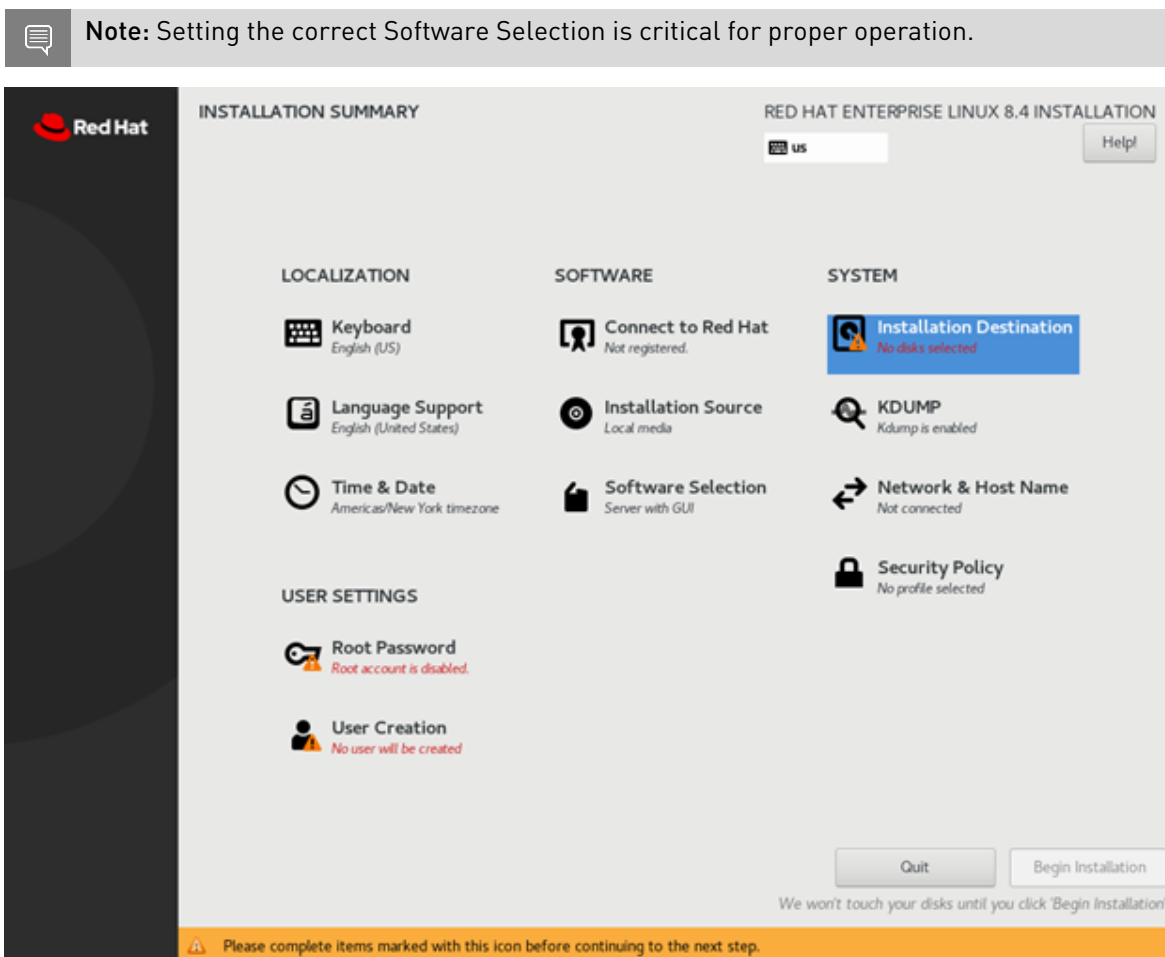
2. Select `Install Red Hat Enterprise Linux` and press `'e'` to edit the selected option.
 - a). Add the `"nomodeset"` parameter as shown below.

`"nomodeset"` prevents the in-box Nouveau driver from loading as it may not properly support the GPUs used in the DGX systems, resulting in display issues.



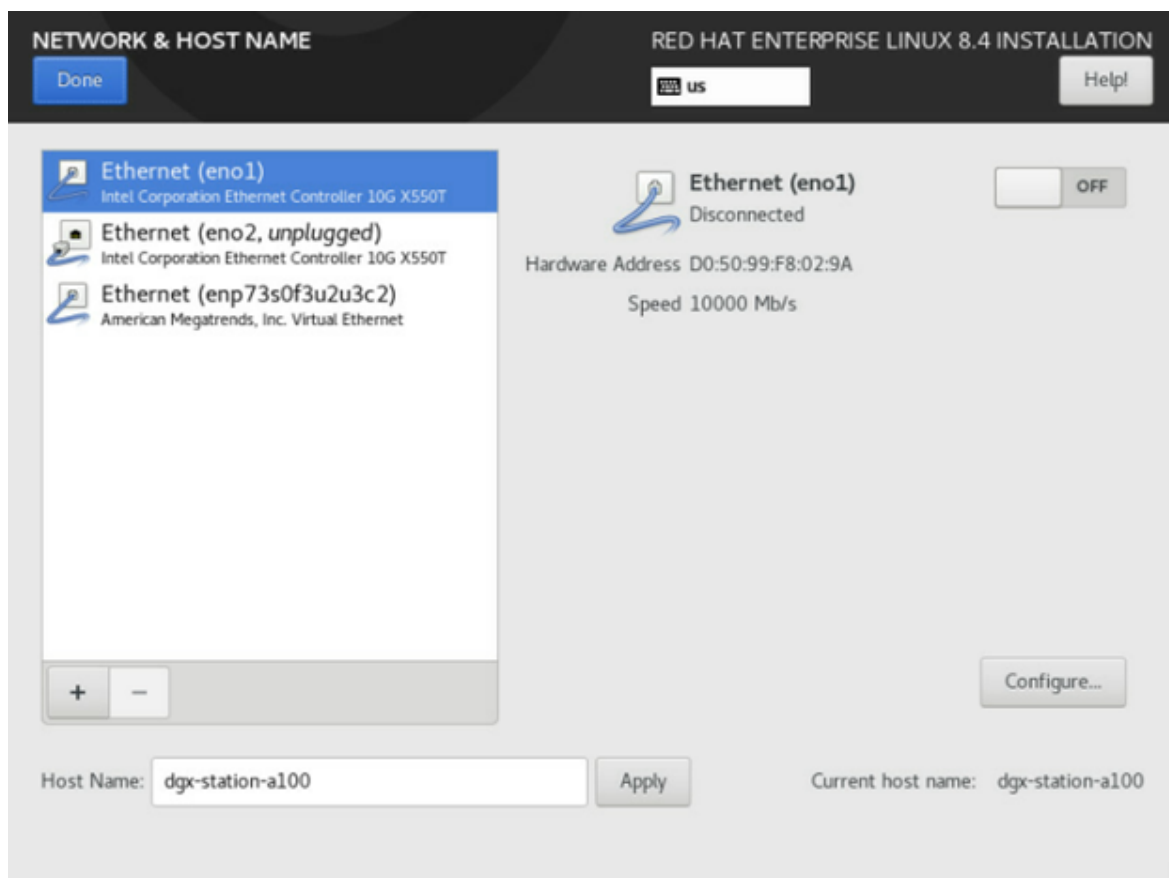
- b). Press **Ctrl-x** to save your changes and start the installation.
3. Refer to the [CentOS Quick Installation Guide \(https://docs.centos.org/en-US/centos/install-guide/Simple_Installation/\)](https://docs.centos.org/en-US/centos/install-guide/Simple_Installation/) for guidance on using the installer.
4. Configure the language, region, date, time, keyboard, and other configuration options you may need from the **Installation Summary** screen.
5. Set the **Software Selection** to the correct value.
 - ▶ For DGX Station and DGX Station A100: Set to **Server with GUI**.

- ▶ For DGX Servers (DGX 1-1, DGX-2, DGX A100): Set to **Minimal Install**.

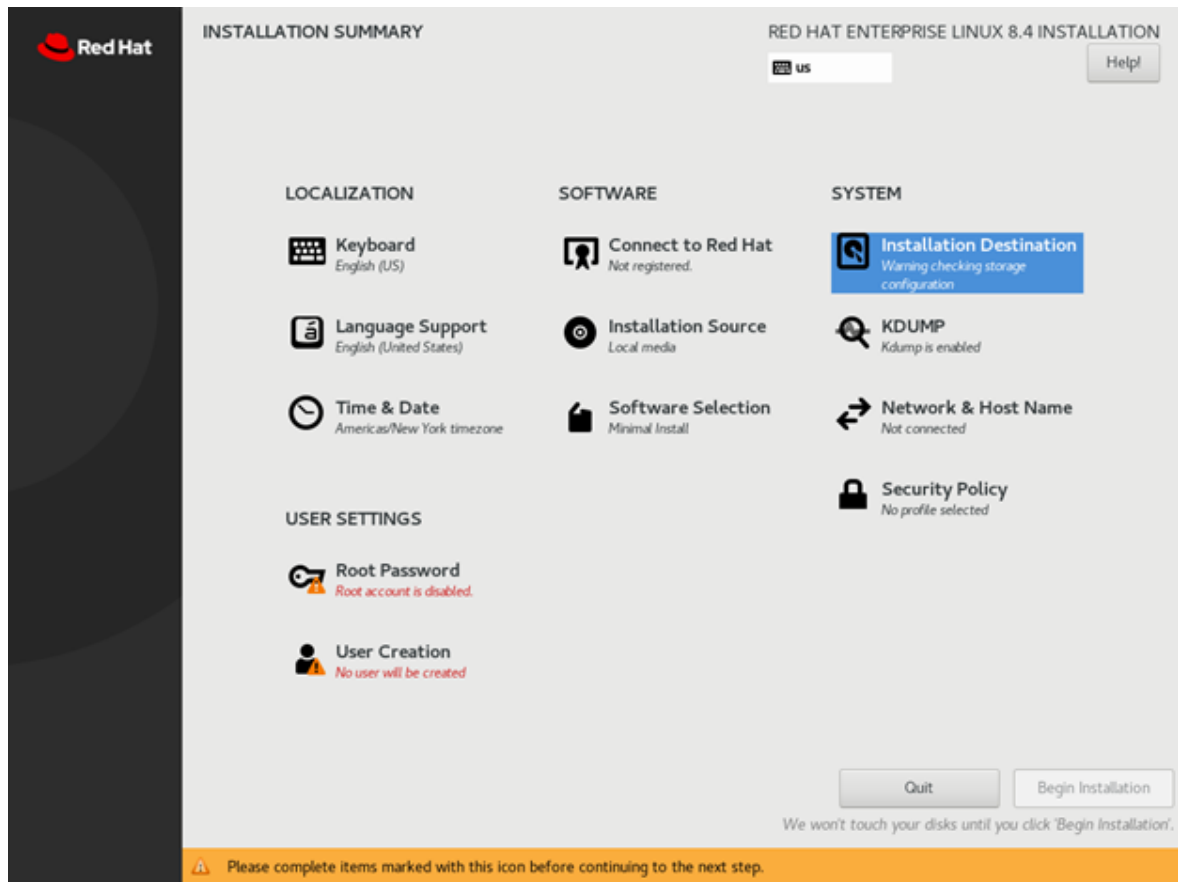


- Partition your disks based on the platform and whether or not you would like to use encryption:
 - ▶ [Disk Partitioning for DGX-1, DGX Station, and DGX Station A100](#)
 - ▶ [Disk Partitioning with Encryption for DGX-1, DGX Station, and DGX Station A100](#)
 - ▶ [Disk Partitioning for DGX-2 and DGX A100](#)
 - ▶ [Disk Partitioning with Encryption for DGX-2 and DGX A100](#)
- Configure Ethernet.

From the **Network & Host Name** section, select and enable the Ethernet device. This defaults to DHCP and can be changed for static IP configurations under **Configure**.



8. From the **INSTALLATION SUMMARY** screen, set your password (at **User Settings > Root Password**) and create a new user (at **User Settings > User Creation**), then click **Begin installation** to start the installation.



When the system reboots, be sure and eject any still-mounted installation discs.

9. Complete the initial CentOS setup.

- ▶ If you installed using the *Server with GUI* base environment, the Initial Setup starts automatically where you can accept the license agreement and register the system. See the [Red Hat instructions](#) for details.
- ▶ If you are using CentOS you will need to accept the license agreement.
- ▶ If you installed with any other base environment, log on to the system as root user and then register the system.

```
subscription-manager register --auto-attach --username=user_name --
password=password
```

See [How to register and subscribe a system to the Red Hat Customer Portal using Red Hat Subscription-Manager](#) for further information.

2.4.1. Installing with Kickstart

Kickstart provides a method of automating the installation process by providing a file which contains the answers to the questions that would be asked during installation. NVIDIA provides kickstart templates for all supported platforms in the following locations:

- ▶ <https://repo.download.nvidia.com/baseos/el/el-files/8/el8-ks/dgxstation-ks.cfg>
- ▶ <https://repo.download.nvidia.com/baseos/el/el-files/8/el8-ks/dgx2-ks.cfg>
- ▶ <https://repo.download.nvidia.com/baseos/el/el-files/8/el8-ks/dgx1-ks.cfg>
- ▶ <https://repo.download.nvidia.com/baseos/el/el-files/8/el8-ks/dgx-a100-ks.cfg>
- ▶ <https://repo.download.nvidia.com/baseos/el/el-files/8/el8-ks/dgxstation-a100-ks.cfg>

For more information on using kickstart files with Red Hat 8, refer to the [Starting Kickstart Installations](#) chapter in Red Hat's documentation.

1. Customize the kickstart template files.

The template files contain tags of the form `<CHANGE_YOUR_XXXX>`. Replace the tags with your site-specific information, such as language, keyboard, time zone, hostname, etc.



Important: Do not specify `'cdrom'` or `'ignoredisk --only-use='` in the template files. This causes the USB CDRROM drive to be ignored, which results in an installation failure since Anaconda is unable to find the install media. Instead of specifying `cdrom` in the template file, use `harddrive --partition=sdX --dir=/` as a workaround. Refer to the following for more information:

- ▶ https://bugzilla.redhat.com/show_bug.cgi?id=1914955
- ▶ <https://access.redhat.com/solutions/5783271>

2. Once you have a kickstart file customized for your installation, place that file in a location that can be accessed by NFS, FTP, HTTP, or HTTPS.
3. Edit the grub menu.

Boot from the installation medium and then press `'e'` at the grub menu to edit the grub entry.

- ▶ Add the `nomodeset` option as you normally would.
- ▶ Add the `inst.ks=<kickstart-file-URL>` option.

Example

```
inst.ks=https://192.168.1.2/kickstart/dgx-a100-cfg.ks
```

4. Reboot the system to apply the settings.

2.4.2. Disk Partitioning for DGX-1, DGX Station, and DGX Station A100



Note: The screenshots in the following section are taken from a DGX-1. The screens for the DGX Station and DGX Station A100 installations can present slightly different information for such things as disk size, disk space available, interface names, etc.

1. At the **Installation Destination** screen, select the first drive (`sda`) to install on.

Under **Storage Configuration**, click the **Custom** radio button and then click **Done**.

This brings up the **Manual Partitioning** window.

- Expand the drop-down menu for the device you've selected and delete all existing partitions until there are none remaining.
- Click the **+** button on the bottom left to create a new partition.

Create the `/boot/efi` partition with a size of 512 MiB.

- Select **Standard Partition** for the Device Type.

Under **Devices**, ensure only the drive you chose in step 1 is selected. Click **Update Settings** to confirm your changes.

- Click the **+** button on the bottom left again to create another new partition.

Create the `/` partition, this time leaving the **Desired Capacity** field empty. This lets the installer know to use the remaining capacity of the disks.

- For the `/` partition, again select **Standard Partition** for the Device Type.

Under **Devices**, ensure that only the drive you chose in step 1 is selected. For the File System, select **XFS**. Click **Update Settings** to confirm your changes.

- Click **Done**.

This causes a yellow warning bar to appear on the bottom because no swap partition has been created.

- Click **Done** again and then click **Accept Changes** to write all of our customizations to disk.

- Return to step 5 in the [Installing Red Hat Enterprise Linux](#) section.

2.4.3. Disk Partitioning with Encryption for DGX-1, DGX Station, and DGX Station A100



Note: The screenshots in the following section are taken from a DGX-1. The screens for the DGX Station and DGX Station A100 installations can present slightly different information for such things as disk size, disk space available, interface names, etc.

- At the **Installation Destination** screen, select the first drive (`sda`) to install on. Under **Storage Configuration**, click the **Custom** radio button and then click **Done**.

- This brings up the **Manual Partitioning** window.
- Expand the drop-down menu for the device you've selected and delete all existing partitions until there are none remaining.
 - Click the **+** button on the bottom left to create a new partition.
Create the `/boot/efi` partition with a size of 512 MiB.
 - Select **Standard Partition** for the Device Type.

Under **Devices**, ensure only the drive you chose in step 1 is selected.

Click **Update Settings** to confirm your changes.

- Click the **+** button on the bottom left again to create another new partition.
Create the `/boot` partition with a size of 2 GiB.
- For the `/boot` partition, again select **Standard Partition** for the Device Type.

Under **Devices**, ensure that only the drive you chose in step 1 is selected. For the File System, select **XFS**.

Click **Update Settings** to confirm your changes.

- Click the **+** button on the bottom left again to create another new partition.
Create the `/` partition, this time leaving the **Desired Capacity** field empty. This lets the installer know to use the remaining capacity of the disks.
- For the `/` partition, again select **Standard Partition** for the Device Type.

Be sure to check the **Encrypt** checkbox, and choose **luks2** for the **LUKS Version**.

For the **File System**, select **XFS**. Again confirm that this partition uses both of the devices you selected in step 2.

Click **Update Settings** to confirm your changes.

A prompt will now appear, asking for a LUKS password.

- Be sure to choose a strong and secure password.

After entering your password twice, click **Save Passphrase**.

- Click **Done**.

This causes a yellow warning bar to appear on the bottom because no swap partition has been created.

11. Click **Done** again and then click **Accept Changes** to write all of our customizations to disk.

12. Return to step 5 in the [Installing Red Hat Enterprise Linux](#) section.

2.4.4. Disk Partitioning for DGX-2 and DGX A100



Note: The screenshots in the following section are taken from a DGX A100. The screens for the DGX-2 installation can present slightly different information for such things as disk size, disk space available, interface names, etc.

1. At the **Installation Destination** screen, select both M.2 disks (that is, the 894,25GB drives for DGX-2, or the 1.8 TB drives for DGX A100).

Under **Storage Configuration**, click the **Custom** radio button and then click **Done**.

This brings up the **Manual Partitioning** window.

2. Expand the drop-down menu for the device you've selected and delete all existing partitions until there are none remaining.
3. Click the **+** button on the bottom left to create a new partition.
Create the `/boot/efi` partition with a size of 512 MiB.
4. Change the **Device Type** to **RAID**, change the **RAID Level** to **RAID1**, then click **Update Settings** and confirm that this partition covers both of the devices you selected in step 2.

5. Click the **+** button on the bottom left again to create another new partition.

Create the `/` partition, this time leaving the **Desired Capacity** field empty. This lets the installer know to use the remaining capacity of the disks.

6. For the `/` partition, again select **RAID** for the **Device Type**, **RAID1**, and the **RAID Level**.
For the **File System**, select **XFS**.

Click **Update Settings** to confirm this partition uses both of the devices you selected in step 2.

7. Click **Done**.

This causes a yellow warning bar to appear on the bottom because no swap partition has been created.

8. Click **Done** again and then click **Accept Changes** to write all of our customizations to disk.
9. Return to the instructions in the [Installing Red Hat Enterprise Linux](#) section.
This should be at the step to configure Ethernet.

2.4.5. Disk Partitioning with Encryption for DGX-2 and DGX A100



Note: The screenshots in the following section are taken from a DGX A100. The screens for the DGX-2 installation can present slightly different information for such things as disk size, disk space available, interface names, etc.

1. At the **Installation Destination** screen, select both M.2 disks (that is, the 894,25GB drives for DGX-2, or the 1.8 TB drives for DGX A100).

Under **Storage Configuration**, click the **Custom** radio button and then click **Done**.

This brings up the **Manual Partitioning** window.

2. Expand the drop-down menu for the device you've selected and delete all existing partitions until there are none remaining.
3. Click the **+** button on the bottom left to create a new partition.
Create the `/boot/efi` partition with a size of 512 MiB.
4. Change the **Device Type** to **RAID**, change the **RAID Level** to **RAID1**, and confirm that this partition covers both of the devices you selected in step 2 by clicking **Modify** in the **Device(s)** section and making sure both disks are still selected.

Note that the partition information on the left may not reflect both devices yet.

Click **Update Settings** to confirm your changes, and verify that the `/boot/efi` partition shown lists the partition name and not the individual device.

5. Click the **+** button on the bottom left again to create another new partition.
Create the `/boot` partition with a size of **2 GiB**.
6. For the `/boot` partition, again select **RAID** for the **Device Type**, **RAID1**, and the **RAID Level**.
For the **File System**, select **XFS**.

Again confirm that this partition uses both of the devices you selected in step 2 by clicking **Modify** in the **Device(s)** section and making sure both disks are still selected.

Click **Update Settings** to confirm your changes.

7. Click the **+** button on the bottom left again to create another new partition.

Create the `/` partition, this time leaving the **Desired Capacity** field empty. This lets the installer know to use the remaining capacity of the disks.

8. For the `/` partition, again select RAID for the Device Type, RAID1 and the RAID Level. Be sure to check the **Encrypt** checkbox, and choose **luks2** for the **LUKS Version**.

For the **File System**, select **XFS**. Again confirm that this partition uses both of the devices you selected in step 2.

Click **Update Settings** to confirm your changes, and **Done** to begin exiting the partition menu.

A prompt will now appear asking for a LUKS password.

9. Be sure to choose a strong and secure password.

After entering your password twice, click **Save Passphrase**.

10. Click **Done**.

This causes a yellow warning bar to appear on the bottom because no swap partition has been created.

11. Click **Done** again and then click **Accept Changes** to write all of our customizations to disk.

12. Return to step 5 in the [Installing Red Hat Enterprise Linux](#) section.

Chapter 3. Installing the DGX Software

This section requires that you have already installed CentOS 8 or derived operating system on the DGX system. You can skip this section if you already installed the DGX software stack during a kickstart install.

3.1. Configuring a System Proxy

If your network requires use of a proxy, then

- ▶ Edit the file `/etc/dnf/dnf.conf` and make sure the following lines are present in the `[main]` section, using the parameters that apply to your network:

```
proxy=http://<Proxy-Server-IP-Address>:<Proxy-Port>
proxy_username=<Proxy-User-Name>
proxy_password=<Proxy-Password>
```

3.2. Enabling the DGX Software Repository



ATTENTION: By running these commands you are confirming that you have read and agree to be bound by the [DGX Software License Agreement](#). You are also confirming that you understand that any pre-release software and materials available that you elect to install in a DGX may not be fully functional, may contain errors or design flaws, and may have reduced or different security, privacy, availability, and reliability standards relative to commercial versions of NVIDIA software and materials, and that you use pre-release versions at your risk.

Install the NVIDIA DGX Package for Red Hat Enterprise Linux.

```
$ sudo dnf install -y https://repo.download.nvidia.com/baseos/el/el-files/8/nvidia-repo-setup-21.06-1.el8.x86_64.rpm
```

3.3. Installing Required Components

1. Upgrade to the latest software.



Important: Before performing the upgrade, consult the release notes for additional instructions depending on the specific EL8 release.

```
sudo dnf update -y --nobest
```

2. Install DGX tools and configuration files.

- ▶ For DGX-1, install **DGX-1 Configurations**.

```
sudo dnf group install -y 'DGX-1 Configurations'
```

- ▶ For the DGX-2, install **DGX-2 Configurations**.

```
sudo dnf group install -y 'DGX-2 Configurations'
```

- ▶ For the DGX A100, install **DGX A100 Configurations**.

```
sudo dnf group install -y 'DGX A100 Configurations'
```

- ▶ For the DGX Station, install **DGX Station Configurations**.

```
sudo dnf group install -y 'DGX Station Configurations'
```

- ▶ For the DGX Station A100, install **DGX Station A100 Configurations**.

```
sudo dnf group install -y 'DGX Station A100 Configurations'
```

The configuration changes take effect only after rebooting the system, which will be performed after installing the CUDA driver.

3. Configure the `/raid` partition.

All DGX systems support RAID 0 or RAID 5 arrays.

The following commands create a RAID array, mount it to `/raid` and create an appropriate entry in `/etc/fstab`.

- ▶ To create a RAID 0 array:

```
sudo /usr/bin/configure_raid_array.py -c -f
```

- ▶ To create a RAID 5 array:

```
sudo /usr/bin/configure_raid_array.py -c -f -5
```



Note: The RAID array must be configured before installing `nvidia-conf-cachefilesd`, which places the proper SELinux label on the `/raid` directory. If you ever need to recreate the RAID array — which will wipe out any labeling on `/raid` — after `nvidia-conf-cachefilesd` has already been installed, be sure to restore the label manually before restarting `cachefilesd`.

```
sudo restorecon /raid
sudo systemctl restart cachefilesd
```

Optional: If you wish to use your RAID array for caching, install `nvidia-conf-cachefilesd`. This will update the `cachefilesd` configuration to use the `/raid` partition.

```
sudo dnf install -y nvidia-conf-cachefilesd
```

4. Install the NVIDIA CUDA driver.

- a). For non-NVSwitch systems such as DGX-1, DGX Station, and DGX Station A100, install either the R450 or R470 driver using the default profile:

Installing the R450 drivers

```
sudo dnf module install -y nvidia-driver:450/{default,src}
sudo dnf install -y nv-persistence-mode libnvidia-nscq-450
```

Installing the R470 drivers

```
sudo dnf module install -y nvidia-driver:470/{default,src}
sudo dnf install -y nv-persistence-mode libnvidia-nscq-470
```

- b). For NVSwitch systems such as DGX-2 and DGX A100, install either the R450 or R470 driver using the fabric manager (fm) profile:

Installing the R450 drivers

```
sudo dnf module install -y nvidia-driver:450/{fm,src}
sudo dnf install -y nv-persistence-mode nvidia-fm-enable
```

Installing the R470 drivers

```
sudo dnf module install -y nvidia-driver:470/{fm,src}
sudo dnf install -y nv-persistence-mode nvidia-fm-enable
```

5. (DGX Station A100 Only) Install additional packages required for DGX Station A100.

These packages must be installed after installation of the `nvidia-driver` module.

```
sudo dnf install -y nvidia-conf-xconfig nv-docker-gpus
```

6. Reboot the system to load the drivers and to update system configurations.

- a). Issue the reboot.

```
sudo reboot
```

- b). After the system has rebooted, verify that the drivers have been loaded and are handling the NVIDIA devices.

```
nvidia-smi
```

The output should show all available GPUs.

Example:

```
+-----+
+
| NVIDIA-SMI 450.80.02      Driver Version: 450.80.02      CUDA Version: 11.0
|
+-----+-----+-----+-----+
+
| GPU   Name               Persistence-M| Bus-Id        Disp.A | Volatile Uncorr. ECC
| Fan   Temp   Perf    Pwr:Usage/Cap|      Memory-Usage | GPU-Util  Compute M.
|
|                               |              |           |
|                               |              |           |
+=====+=====+=====+=====+
|    0   Tesla V100-SXM2...  On      | 00000000:06:00.0 Off  |
| N/A    35C    P0     42W / 300W |      0MiB / 16160MiB |      0%      Default
|
|                               |              |           |
|                               |              |           |
+-----+-----+-----+-----+
+
|    1   Tesla V100-SXM2...  On      | 00000000:07:00.0 Off  |
| N/A    35C    P0     44W / 300W |      0MiB / 16160MiB |      0%      Default
|
```


- ▶ To administer self-encrypting drives, install the `nv-disk-encrypt` package, issue the following.

```
$ sudo dnf install -y nv-disk-encrypt
$ sudo reboot
```

Refer to the “Managing Self-Encrypting Drives” section in the [DGX A100 User Guide](#) for usage information.

- ▶ To install the NVIDIA Collectives Communication Library (NCCL) Runtime, refer to the [NCCL:Getting Started](#) documentation.
- ▶ To install the CUDA Deep Neural Networks (cuDNN) Library Runtime, refer to the [NVIDIA cuDNN](#) page.
- ▶ To install NVIDIA TensorRT, refer to the [NVIDIA TensorRT](#) page.
- ▶ To install NVIDIA GPUDirect Storage (GDS), issue the following to install the GDS packages.

```
$ sudo dnf install nvidia-gds
```

Be sure to enable GDS within the MLNX_OFED driver if you install the driver. Refer to [Using the NVIDIA Mellanox InfiniBand Drivers](#).

3.5. Installing the Optional NVIDIA Desktop Theme

The DGX Software Repository also provides optional theme packages and desktop wallpapers to give the user-interface an NVIDIA look-and-feel for the DGX Station desktop. These packages would have been installed as part of the **DGX Station Configurations** group.

1. To apply the theme and background images, first open `gnome-tweaks`.
2. Under **Applications**, select one of the **NV-Yaru** themes.
This comes in default, light, and dark variations.
3. Under **Shell**, select the **NV-Yaru-dark** theme.

If this field is grayed out, you may need to reboot the system or restart GDM in order to enable the user-themes extension.

To restart GDM, issue the following.

```
sudo systemctl restart gdm
```

4. Select one of the NVIDIA wallpapers for the background image and lock screen.

Chapter 4. Using the NVIDIA Mellanox InfiniBand Drivers

The DGX software stack for Red Hat-derived operating systems does not include the Mellanox OpenFabrics Enterprise Distribution (MLNX_OFED) for Linux. This is to ensure that the MLNX_OFED driver is in sync with the Red Hat distribution kernel. This section describes how to download, install, and upgrade MLNX_OFED on systems that are running Red Hat Enterprise Linux.

4.1. Determining the MLNX_OFED Version to Install

NVIDIA validates each release of NVIDIA EL8 software with a specific MLNX_OFED version. Consult the [NVIDIA EL8 release notes](#) for the recommended MLNX_OFED version to install for a particular version of NVIDIA EL8 software.

The following table provides a quick reference for tested versions.

Red Hat Enterprise Linux Version	MLNX_OFED Version
8.4	5.4-1.0.3.0
8.5	5.4-3.1.0.0


4.2. Installing the NVIDIA Mellanox InfiniBand Drivers

This section describes how to install MLNX_OFED on systems that do not yet have it installed. It is imperative that a validated MLNX_OFED version is used for the RHEL version that the DGX system is running. Note that the “dnf update” command that is run before installing the NVIDIA driver will update the system to the latest Red Hat Enterprise Linux version.

1. Determine which version of Red Hat Enterprise Linux is installed on the DGX system.

```
cat /etc/redhat-release
```

2. Determine the appropriate MLNX_OFED software bundle to install..
Refer to [Determining the MLNX_OFED Version to Install](#).
3. Download the MLNX_OFED software bundle.
 - a). Visit the [Linux InfiniBand Drivers](#) page, scroll down to the Download wizard, and then click the **Download** tab.
 - b). At the **MLNX_OFED Download Center** matrix, choose
 - ▶ The version to install (you may need to select **Archive Versions**),
 - ▶ **RHEL/CentOS** (under OS Distribution), and
 - ▶ The relevant OS Distribution Version and Architecture.
 - c). Click the desired ISO/tgz package.
To obtain the download link, accept the End User License Agreement.
4. After downloading the correct MLNX_OFED software bundle, proceed with the installation steps.
 - a). Re-visit the [MLNX_OFED Software Releases](#) site and select the MLNX_OFED software version you intend to use.
 - b). Use the side menu to navigate to **Installation->Installing MLNX_OFED**, and follow the instructions.


 **Note:** The system may report that additional software needs to be installed before performing the installation. If such a message appears, install the software and then retry installing the MLNX_OFED driver.

5. If you intend to use NVIDIA GPUDirect Storage (GDS), enable the driver's GDS support according to the instructions at <https://docs.nvidia.com/gpudirect-storage/troubleshooting-guide/index.html#mofed-req-install>.
6. Install `nvidia-mlnx-config`.

```
sudo dnf install -y nvidia-mlnx-config
```
7. Install kernel headers and development packages for your kernel.
These are needed for the ensuing DKMS compilation.

```
sudo dnf install -y kernel-headers-$(uname -r) kernel-devel-$(uname -r)
```
8. After installing the MLNX_OFED drivers, install the NVIDIA peer memory module.

```
sudo dnf install -y nvidia-peer-memory-dkms
```

 **Note:** While in-box drivers may be available, using the in-box drivers is not recommended as they provide lower performance than the official MLNX OFED drivers and they do not support

the GPUDirect™ RDMA feature. For more information on configuring the in-box drivers, see the following Red Hat Enterprise Linux documentation:

- ▶ [Configuring InfiniBand and RDMA Networks](#)

4.3. Updating the NVIDIA Mellanox InfiniBand Drivers

This section describes how to update MLNX_OFED on systems that **already have it installed**. If MLNX_OFED drivers are not yet installed, follow the instructions at [Installing the NVIDIA Mellanox InfiniBand Drivers](#).

The Mellanox InfiniBand Drivers in RPM packages are precompiled for a specific kernel version. It is imperative that a validated MLNX_OFED version is used for the Red Hat Enterprise Linux version that the DGX system has been updated to. There is no need to uninstall the current MLNX_OFED first, because the "mlnxofedinstall" script will automatically uninstall any previously installed versions.



Note: The MLNX_OFED drivers support Red Hat Enterprise Linux weak-modules script. This means that any updates to the kernel within the same Red Hat Enterprise Linux version (for example, 8.4) will not require an update to the MLNX_OFED driver.

1. Upgrade the Red Hat Enterprise Linux release and kernel version.


```
sudo dnf update --nobest
```
2. Determine which version of Red Hat Enterprise Linux is installed on the DGX system.


```
cat /etc/redhat-release
```
3. Determine the appropriate MLNX_OFED software bundle to install..

Refer to [Determining the MLNX_OFED Version to Install](#).
4. Download the MLNX_OFED software bundle.
 - a). Visit the [Linux InfiniBand Drivers](#) page, scroll down to the Download wizard, and then click the **Download** tab.

 - b). At the **MLNX_OFED Download Center** matrix, choose
 - ▶ The version to install (you may need to select **Archive Versions**),
 - ▶ **RHEL/CentOS** (under OS Distribution), and
 - ▶ The relevant OS Distribution Version and Architecture.

c). Click the desired ISO/tgz package.

To obtain the download link, accept the End User License Agreement.

5. Mount the downloaded ISO somewhere on the system.

The following example shows the ISO being mounted on the `/mnt` directory.

```
sudo mount MLNX_OFED_LINUX-<version>.iso /mnt
```

6. Prepare to install the driver.

a). Remove `nvidia-mlnx-config` and `nvidia-peer-memory-dkms`.

```
sudo dnf remove -y nvidia-mlnx-config nvidia-peer-memory-dkms
```

The `mlnxofedinstall` step will remove packages prior to installing new ones, causing `nvidia-mlnx-config` and `nvidia-peer-memory-dkms` to fall out because they depend on some of these removed packages. Removing those components ahead of time avoids issues. These will be reinstalled as a final step.

b). Specify the new kernel version to use when installing the driver.

```
NEXTKERNEL=$(sudo grubby --default-kernel | sed 's/.*vmlinuz\-\//g')
```

7. Install the driver with the `-k` and `-s` flags to specify the new kernel version and kernel source path.

```
sudo /mnt/mlnxofedinstall -k ${NEXTKERNEL} -s /lib/modules/${NEXTKERNEL}/build --force
```



Note: The system may report that additional software needs to be installed before performing the installation. If such a message appears, install the software and then repeat this step.

8. Reboot.

```
sudo reboot
```

9. Reinstall `nvidia-mlnx-config` and `nvidia-peer-memory-dkms`.

```
sudo dnf install -y nvidia-mlnx-config nvidia-peer-memory-dkms
```

10. If you intend to use NVIDIA GPUDirect Storage (GDS), enable the driver's GDS support according to the instructions at <https://docs.nvidia.com/gpudirect-storage/troubleshooting-guide/index.html#mofed-req-install>.

Appendix A. Changing the BMC Login

The NVIDIA DGX server includes a base management controller (BMC) for out-of-band management of the DGX system. NVIDIA recommends disabling the default username and creating a unique username and password as soon as possible.

A.1. Changing the BMC Login on the DGX-1

1. Log into the BMC.
 - a). Open a browser within your LAN and go to `http://<BMC-ip-address>/`.
Use Firefox or Internet Explorer. Google Chrome is not officially supported by the DGX-1 BMC.
 - b). Log in, using `qct.admin/qct.admin` for the User ID/Password.
2. Select Configuration → Users.
3. Add a new user.
 - a). Select an empty field and click Add User.
 - b). Enter new user information and click Add.
4. Log out and then log back in as the new user.
5. Select Configuration → Users.
6. Disable User Access for the user `qct.admin`.
 - a). Select the user `qct.admin` user and select Modify User
 - b). Deselect Enable in User Access and click Modify.
 - c). Ensure User Access is disabled for the user `qct.admin`.

7. Log out.

A.2. Changing the BMC Login on the DGX-2 or DGX A100

1. Log into the BMC.
 - a). Open a browser within your LAN and go to `https://<BMC-ip-address>/`.
 - b). Log in, using admin/admin for the User ID/Password.
2. Select **Settings** from the left-side navigation menu.
3. Select the **User Management** card.
4. Click the green Help icon (?) for information about configuring users, then add a new user with unique username and strong password.
5. Log out and then log back in as the new user.
6. Select **Settings** → **User Management**.
7. Disable the **admin** and **anonymous** users.

Appendix B. Installing the Mellanox ConnectX-6 Firmware

MOFED version 5.1-2.5.8.0 contains adapter card firmware version xx.28.2006, while the latest GA version of the ConnectX-6 firmware is 20.28.4000.

NVIDIA has tested Red Hat Enterprise Linux 8 with the latest ConnectX-6 firmware, 20.28.4000. This can be downloaded directly from the [Mellanox site](#). The ConnectX-6 adapter card versions used in DGX A100 are:

- ▶ OPN MCX653105A-HDA_Ax, PSID MT_0000000223
- ▶ OPN MCX653106A-HDA_Ax, PSID MT_0000000225

Download the latest version from the [Mellanox site](#), and follow the [instructions](#) to burn the new firmware.

Appendix C. Changing NVIDIA Driver Branches

To switch driver branches, you must first remove the existing branch before installing the new branch. Currently, removing the `nvidia-driver` module clears the `GRUB_CMDLINE_LINUX` setting and this can result in the server failing to boot. Be sure to note the current `GRUB_CMDLINE_LINUX` setting and restore it after switching driver branches as explained in this section.

1. Note the existing `GRUB_CMDLINE_LINUX` setting in the file `etc/default/grub`.

Example:

```
GRUB_CMDLINE_LINUX="crashkernel=auto
rd.md.uuid=09a9380c:87edd4b6:8f5d9bbc:45e834c7 rhgb quiet
rd.driver.blacklist=nouveau"
```



Note: The "`rd.driver.blacklist=nouveau`" parameter was added when installing the driver and should not be included in the restoration.

2. Remove and clear the existing stream.

```
$ sudo dnf module remove --all nvidia-driver
$ sudo dnf module reset nvidia-driver
```
3. Follow the instructions in step 4 of the section [Installing Required Components](#) to install the new driver branch.
4. If the `nvidia-peer-memory-dkms` driver is installed, it must be reinstalled to match the new driver branch.

```
$ sudo dnf reinstall -y nvidia-peer-memory-dkms
```
5. In the file `/etc/default/grub`, remove any extra instances of `GRUB_CMDLINE_LINUX` and manually edit the file to restore the original setting (except for the blacklist parameter).

Example:

```
GRUB_CMDLINE_LINUX="crashkernel=auto
rd.md.uuid=09a9380c:87edd4b6:8f5d9bbc:45e834c7 rhgb quiet"
```

6. Reboot the system.

```
sudo reboot
```

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 NON-INFRINGEMENT, 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, DGX, DGX-1, DGX-2, DGX A100, DGX Station, and DGX Station A100 are trademarks and/or registered trademarks of NVIDIA Corporation in the United States and other countries. Other company and product names may be trademarks of the respective companies with which they are associated.

Copyright

© 2022 NVIDIA Corporation. All rights reserved.