



DGX Software with CentOS

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..... | 4 |
| 2.1. Obtaining CentOS..... | 4 |
| 2.2. Booting CentOS ISO Locally..... | 4 |
| 2.3. Booting the CentOS ISO Remotely on the DGX-1, DGX-2, or DGX A100..... | 5 |
| 2.3.1. Booting the ISO Image on the DGX-1 Remotely..... | 5 |
| 2.3.2. Booting the ISO Image on the DGX-2 Remotely..... | 8 |
| 2.3.3. Booting the ISO Image on the DGX A100 Remotely..... | 12 |
| 2.4. Installing CentOS..... | 15 |
| 2.4.1. Installing on the DGX-1, DGX Station, or DGX Station A100..... | 16 |
| 2.4.2. Installing on the DGX-2..... | 23 |
| 2.4.3. Installing on the DGX A100..... | 33 |
| Chapter 3. Installing the DGX Software..... | 44 |
| 3.1. Configuring a System Proxy..... | 44 |
| 3.2. Enabling the Repositories..... | 44 |
| 3.3. Installing Required Components..... | 45 |
| 3.3.1. Installing DGX Tools and Updating Configuration Files..... | 45 |
| 3.3.2. Configuring the /raid Partition..... | 46 |
| 3.3.2.1. Configuring the /raid Partition as an NFS Cache..... | 46 |
| 3.3.2.2. Configuring the /raid Partition for Local Persistent Storage..... | 47 |
| 3.3.3. Installing and Loading the NVIDIA CUDA Drivers..... | 47 |
| 3.3.4. Installing the NVIDIA Container Runtime..... | 48 |
| 3.4. Installing Diagnostic Components..... | 49 |
| 3.5. Replicating the EFI System Partition on DGX-2 or DGX A100..... | 50 |
| 3.6. Installing Optional Components..... | 51 |
| 3.7. Applying an NVIDIA Look and Feel to the Desktop User Interface..... | 51 |
| 3.8. Managing CPU Mitigations..... | 54 |
| 3.8.1. Determining the CPU Mitigation State of the DGX System..... | 54 |

| | |
|--|-----------|
| 3.8.2. Disabling CPU Mitigations..... | 55 |
| 3.8.3. Re-enabling CPU Mitigations..... | 55 |
| Chapter 4. Using the NVIDIA Mellanox InfiniBand Drivers..... | 57 |
| 4.1. Determining the MLNX_OFED Version to Install..... | 57 |
| 4.2. Installing the NVIDIA Mellanox InfiniBand Drivers..... | 57 |
| 4.3. Updating the NVIDIA Mellanox InfiniBand Drivers..... | 59 |
| Chapter 5. Running Containers..... | 61 |
| Chapter 6. Configuring Storage - NFS Mount and Cache..... | 62 |
| Appendix A. Changing the BMC Login..... | 64 |
| A.1. Changing the BMC Login on the DGX-1..... | 64 |
| A.2. Changing the BMC Login on the DGX-2 or DGX A100..... | 69 |
| Appendix B. Using Custom DGX Software Utilities for the DGX Station..... | 71 |
| B.1. Rebuilding or Re-Creating the DGX Station RAID Array..... | 71 |
| B.2. Changing the RAID Level of the RAID Array..... | 72 |
| B.3. EL7-20.01 Only: Checking the Health of the DGX Station..... | 73 |
| B.4. EL7-20.01 Only: Collecting Information for Troubleshooting the DGX Station..... | 74 |
| Appendix C. Expanding the DGX Station RAID Array..... | 76 |

Chapter 1. Introduction

The NVIDIA® DGX™ systems (DGX-1, DGX-2, and DGX A100 servers, and NVIDIA DGX Station™ and DGX Station A100 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 7 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 - 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.



Note:

You can use `yum-config-manager` to conveniently enable certain repositories. To use `yum-config-manager`, first install the yum utilities.

```
sudo yum -y install yum-utils
```

1.2.1.1. NVIDIA Repositories

► NVIDIA DGX Software Repository

After installing CentOS on the DGX system, you must enable the NVIDIA DGX software repository. 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 repositories.

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-sclo-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

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 Linux distributions that are 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

Obtain the CentOS 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.

2.2. Booting CentOS ISO Locally

1. Plug the USB flash drive containing the CentOS 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 UEFI 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 ISO Remotely on the DGX-1, DGX-2, or DGX A100

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, see [Booting the ISO Image on the DGX-2 Remotely](#)
 - ▶ For instructions applicable to the NVIDIA DGX A100, see [Booting the ISO Image on the 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 **User Management**.
 - 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 **Console Redirection**.



- b). Click **Java Console** to open the remote JViewer window. Make sure pop-up blockers are disabled for this site.
 - c). From the JViewer top menu bar, click **Media** and then select **Virtual Media Wizard**.

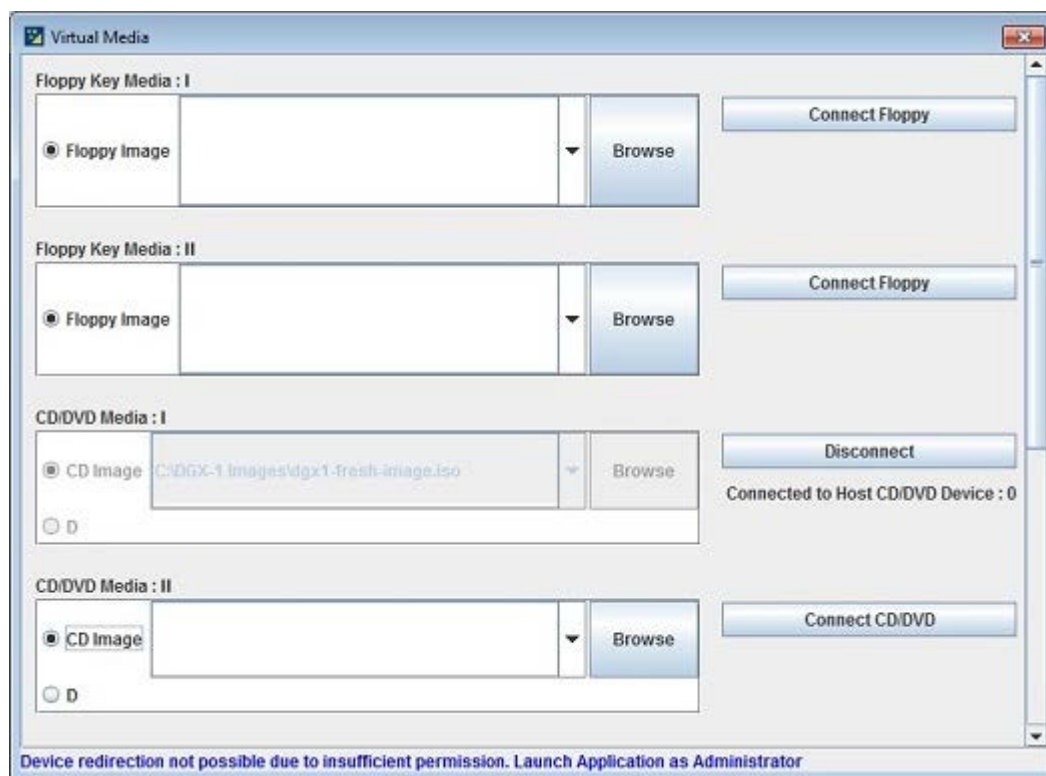


- d). From the CD/DVD Media: I section of the Virtual Media dialog, click **Browse** and then locate the CentOS ISO file on your system and click **Open**.

You can ignore the device redirection warning at the bottom of the Virtual Media wizard as it does not affect the ability to re-image the system.

- e). Click **Connect CD/DVD**, then click **OK** at the Information dialog.

The Virtual Media window shows that the ISO image is connected.



f). Close the window.

The CD ROM icon in the menu bar turns green to indicate that the ISO image is attached.



g). From the top menu, click **Power** and then select **Reset Server**.



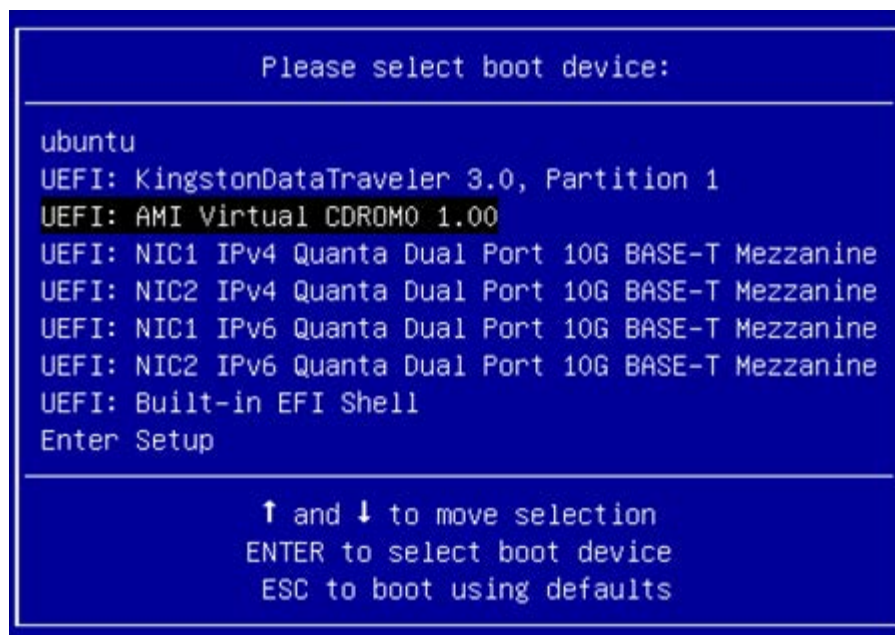
h). Click **Yes** and then **OK** at the Power Control dialogs, then wait for the system to power down and then come back online.

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**.

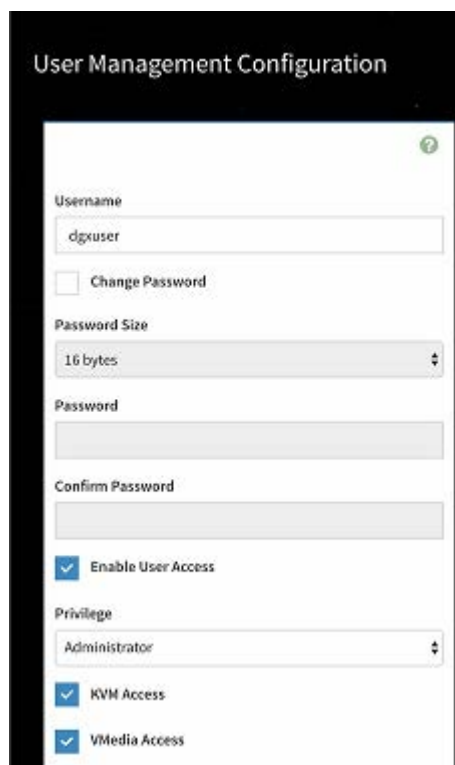


- b). Follow the instructions at [Installing CentOS](#).

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

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

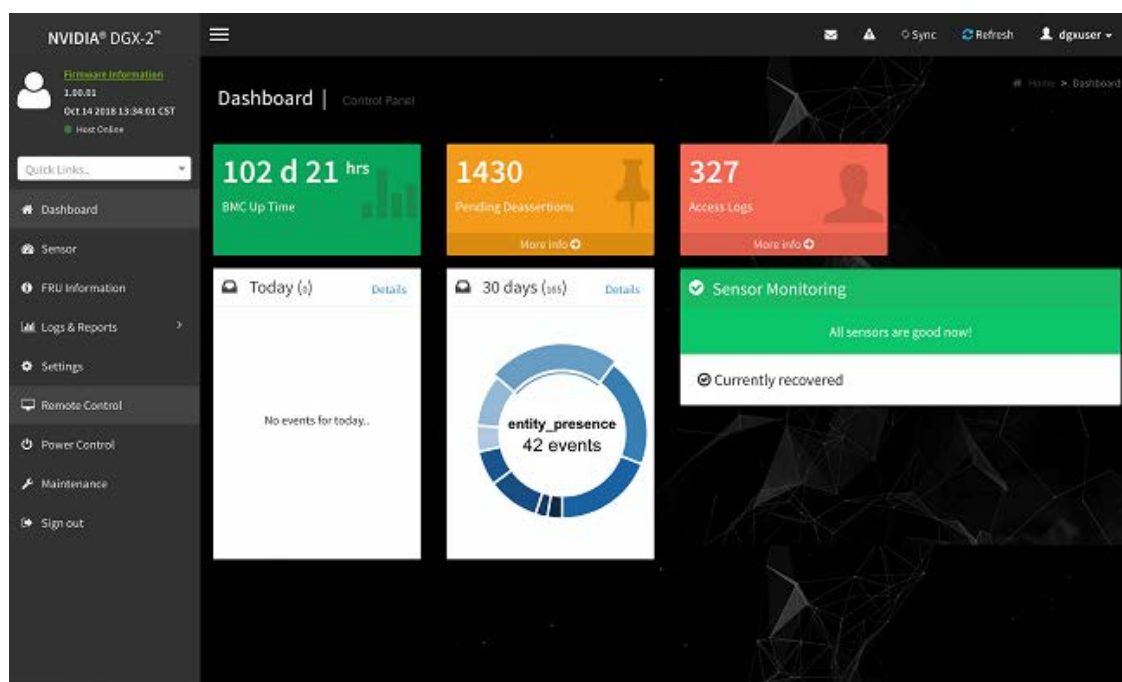
1. Connect to the BMC and ensure the required user privileges are set.
 - a). Open a browser within your LAN and go to `https://<BMC-ip-address>/`, then log in.
 - b). From the left-side menu, click **Settings** and then select **User Management**.
 - c). Click the card with the user name that you created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




The image shows a 'User Management Configuration' dialog box. It contains the following fields and options:

- Username:** A text field containing 'digxuser'.
- Change Password:** An unchecked checkbox.
- Password Size:** A dropdown menu set to '16 bytes'.
- Password:** A text field.
- Confirm Password:** A text field.
- Enable User Access:** A checked checkbox.
- Privilege:** A dropdown menu set to 'Administrator'.
- KVM Access:** A checked checkbox.
- VMedia Access:** A checked checkbox.

2. Set up the ISO image as virtual media.
 - a). From the left-side menu, click **Remote Control**.



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.

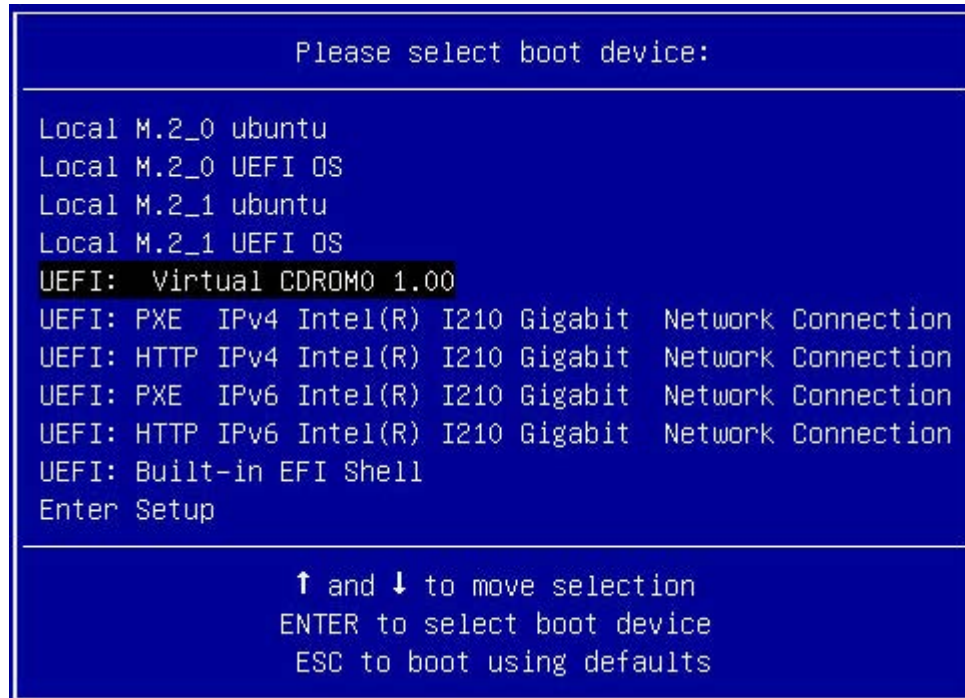
Typically, the default boot order does not boot the CDROM image. This can be changed in the BIOS or as a one-time option in the boot menu.

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: Virtual CDROM 1.00** as the boot device and then press **ENTER**.



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

2.3.3. Booting the ISO Image on the 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**.

User Management Configuration

Username
dgxuser

☐ Change Password

Password Size
16 bytes

Password

Confirm Password

☒ Enable User Access

Privilege
Administrator

☒ KVM Access

☒ VMedia Access

2. Set up the ISO image as virtual media.
 - a). From the left-side menu, click **Remote Control**.

NVIDIA DGX™ A100

Apr 24 2020 20:18:22 CST
FW: 9.11.12
IP: 172.18.205.154
MAC: SCFF35C2507E
Chassis Part: 100-22967-2510-000
Chassis SN: 1234567890123
Host Online
Chassis Identify LED

Quick Links...

- Dashboard
- Sensor
- System Inventory
- FRU Information
- GPU Information
- Logs & Reports
- Settings
- Remote Control
- Power Control
- Chassis ID LED Control
- Maintenance
- Sign out

Dashboard Control Panel

89 d 7 hrs Power-On Hours

181 Pending Disconnections

40 Access Logs

8 GPU Info


| 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 |

Front View

Rear View

Sensor Monitoring
All sensors are good now!
Currently recovered

- 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 CDROM 1.00** as the boot device and then press **ENTER**.



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

2.4. Installing CentOS

2.4.1. Installing on the DGX-1, DGX Station, or DGX Station A100

This section assumes you have booted the CentOS ISO image, either [locally](#) or [remotely](#) for DGX-1, or [locally](#) for DGX Station. or DGX Station A100

- ▶ For instructions applicable to the DGX-2, see [Installing on the DGX-2](#).
- ▶ For instructions applicable to the DGX A100, see [Installing on the DGX A100](#).

The screenshots in these instructions are taken from a DGX-1 system. The DGX Station installation can have slight differences with respect to disk size, disk space available, interface names, etc.


1. After booting the ISO image through either the BMC or from the USB drive, start the installation.



- ▶ **On the DGX-1:** Select **Install CentOS** and then press **Enter** to start the installation.
- ▶ **On the DGX Station:** Update the `linuxefi` parameters to prevent the Nouveau driver in CentOS from being loaded and then start the installation.

The Nouveau driver does not properly support the GPUs in the DGX Station, which may cause display issues. The GPUs in the DGX Station require that drivers that are installed as explained in [Installing and Loading the NVIDIA CUDA Drivers](#).

- a). Select **Install CentOS** and then type **e**.
- b). Append the `nomodeset` option to the `linuxefi` parameters.



- c). Press **Ctrl+X** to save your changes and start the installation.

The CentOS installer starts.

For guidance on using the installer, refer to [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/).


2. Configure the language, region, date, time, keyboard, and other configuration options you may need from the **Installation Summary** screen.
3. Select the software packages that you want to install.

For detailed instructions, refer to [Software Selection in CentOS Installation Guide \(https://docs.centos.org/en-US/centos/install-guide/SoftwareSpoke-x86/\)](https://docs.centos.org/en-US/centos/install-guide/SoftwareSpoke-x86/).

4. 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**.



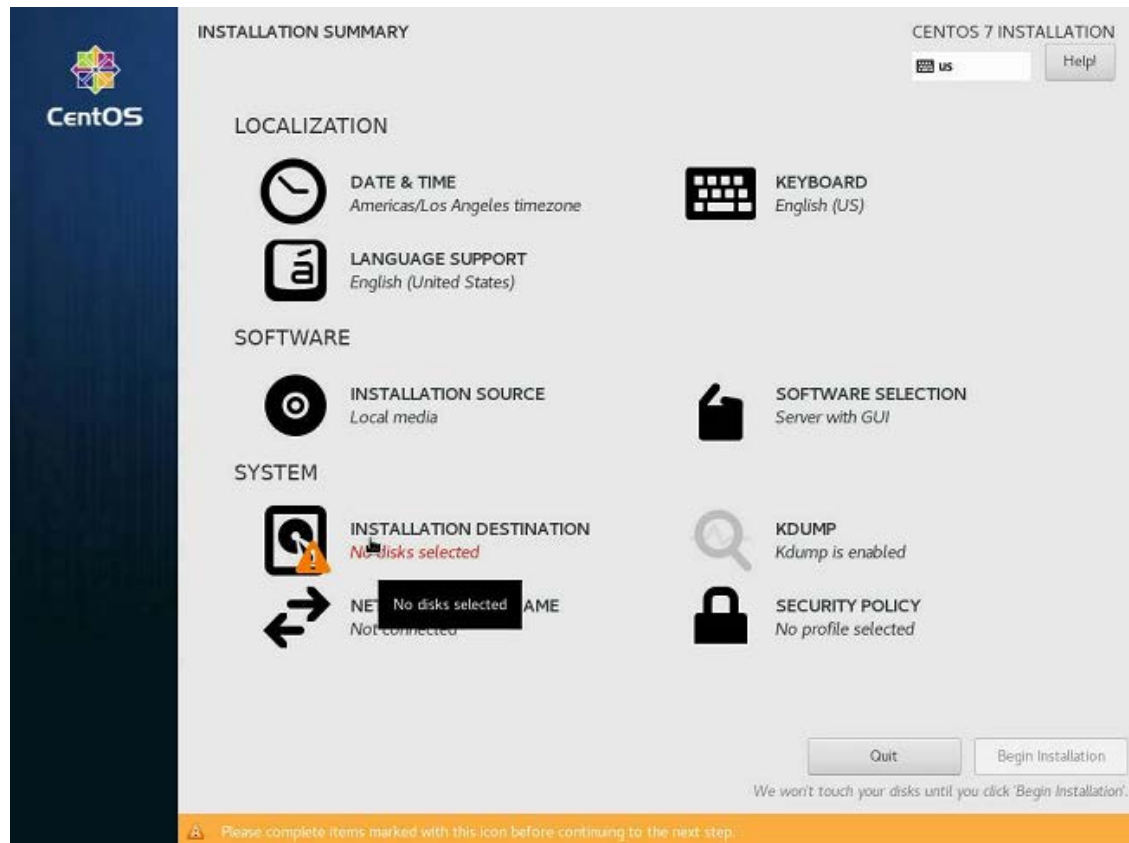
Note: Setting the correct Software Selection is critical for proper operation.



- 5.
6. Set up the system drive.

This step removes the Ubuntu installation in order to reclaim space for the CentOS installation.

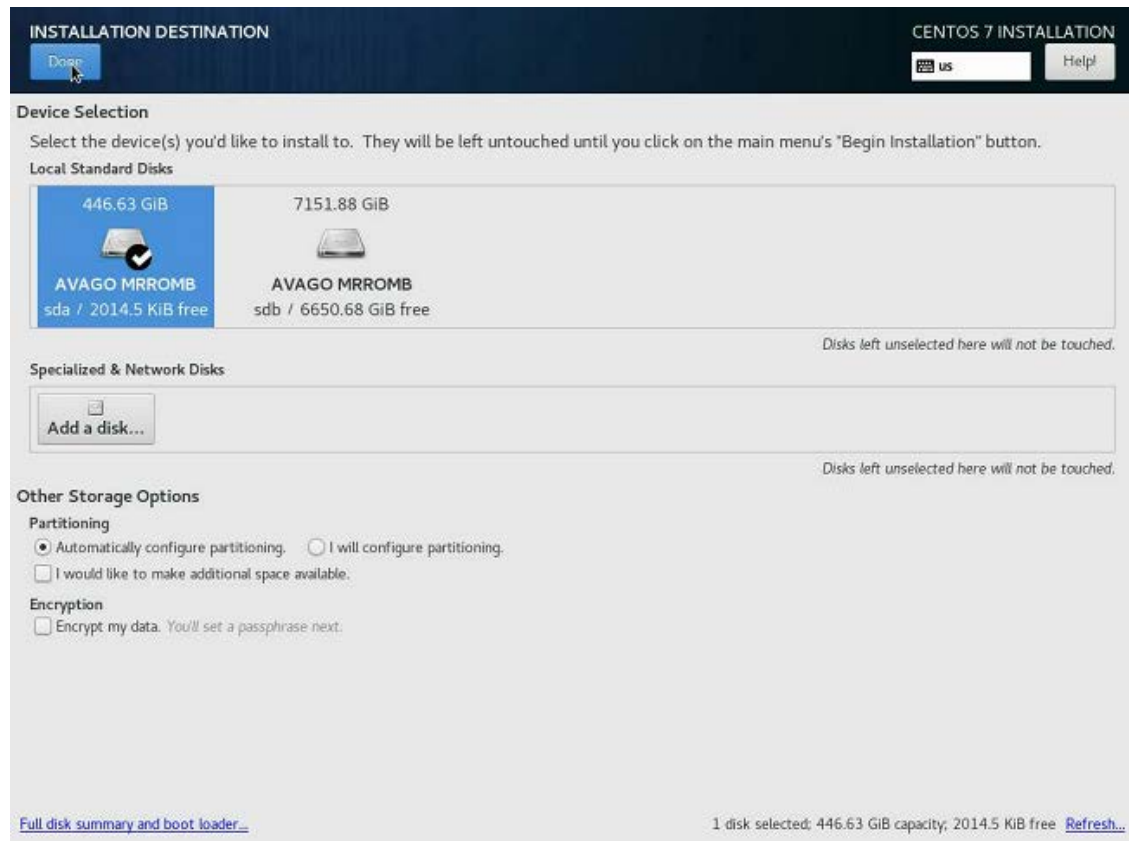
- a). From the Installation Summary screen, click **INSTALLATION DESTINATION**.



- b). At the **Installation Destination** screen, select the first drive (sda) as the installation drive.
- c). Under **Other Storage Options**, click the **I will configure partitioning** radio button then click **Done**.



Note: If the software is being installed on the DGX Station, the number of local standard disks, their labels, capacities, and the amount of free space shown may be different than the examples shown in the screen capture.



The **Manual Partitioning** dialog box appears.

- d). Expand the **Unknown** drop-down menu for the device you've selected, select a partition and then click the "-" button to delete the partition.

Repeat for other existing partitions listed in the **Unknown** drop down until none remain.

- e). Create a new partition.

Click the "+" button on the bottom left and set the following fields:

- ▶ Mount point: **/boot/efi**
- ▶ Desired Capacity: **512 MiB**
- ▶ Device Type: **Standard Partition**

Under **Devices**, ensure only the sda drive you chose is selected, then click **Update Settings** to confirm your changes.

f). Create another partition.

Click the "+" button on the bottom left and set the following fields:

- ▶ Mount point: **/**
- ▶ Desired Capacity: [leave empty to let the installer know to use remaining capacity of the disks]
- ▶ Device Type: **Standard Partition**
- ▶ File System: **XFS**

Under **Devices**, ensure only the sda drive you chose is selected, then click **Update Settings** to confirm your changes.

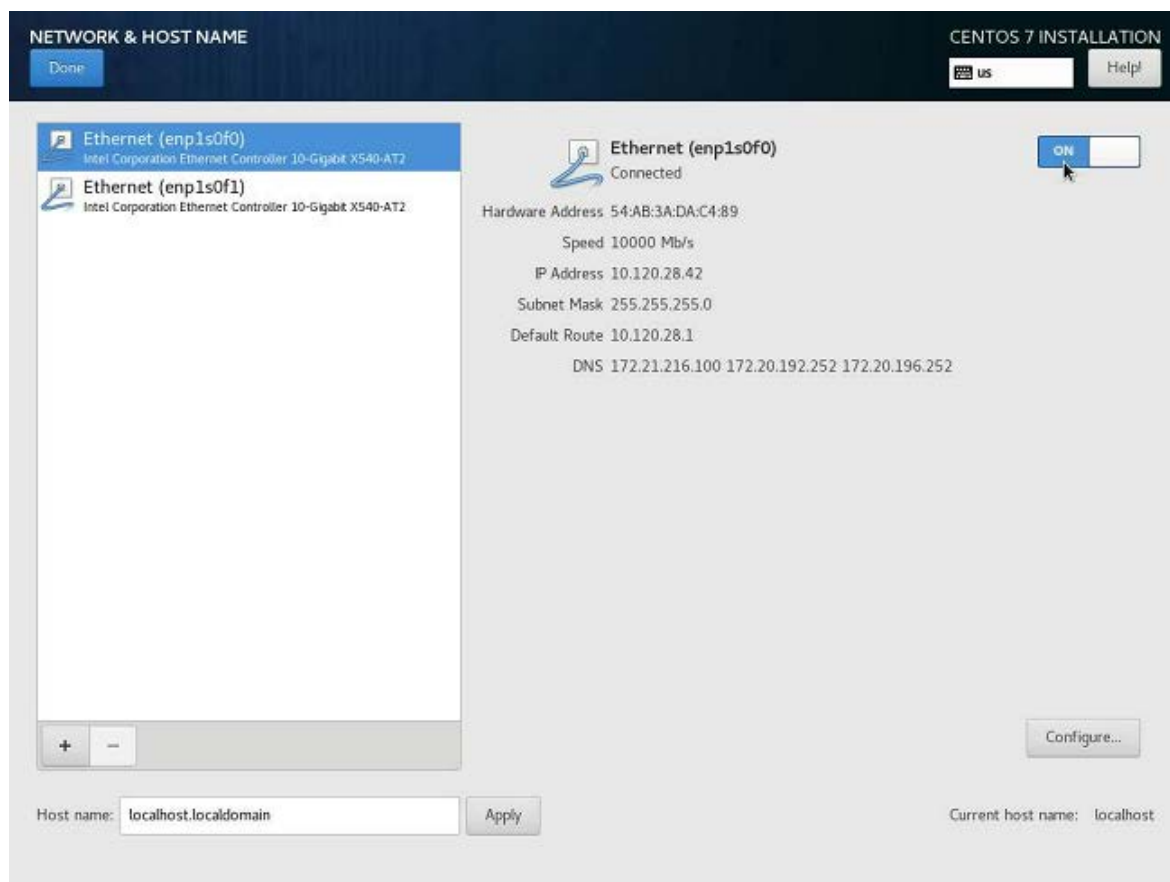
g). Click **Done** when finished.

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

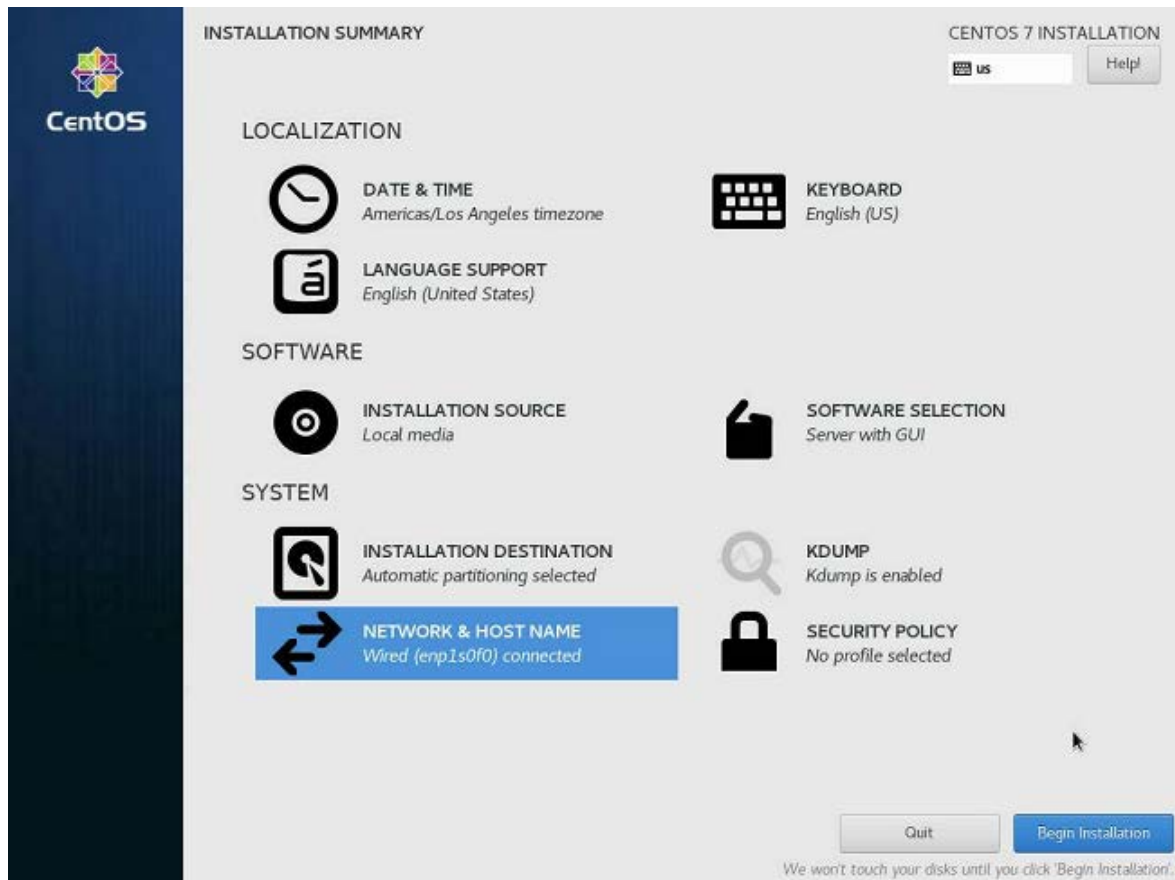
Click **Done** again and then click **Accept Changes** to write all the customizations to disk.

7. Configure Ethernet.

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, click **Begin installation** to start the installation.



- While the installation process is running, set your password (ROOT PASSWORD) and create a new user (USER CREATION) from the Configuration screen.
- When the installation completes, click **Reboot** to reboot the system.

2.4.2. Installing on the DGX-2

This section assumes you have booted the CentOS ISO image, either [locally](#) or [remotely](#).

- After booting the ISO image through either the BMC or from the USB drive, select **Install CentOS** and then press **Enter** to start the installation.



The CentOS installer starts.

For guidance on using the installer, refer to [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/).


2. Configure the language, region, date, time, keyboard, and other configuration options you may need from the Installation Summary screen.
3. Select the software packages that you want to install.

For detailed instructions, refer to [Software Selection in CentOS Installation Guide \(https://docs.centos.org/en-US/centos/install-guide/SoftwareSpoke-x86/\)](https://docs.centos.org/en-US/centos/install-guide/SoftwareSpoke-x86/).

4. 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**.



Note: Setting the correct Software Selection is critical for proper operation.



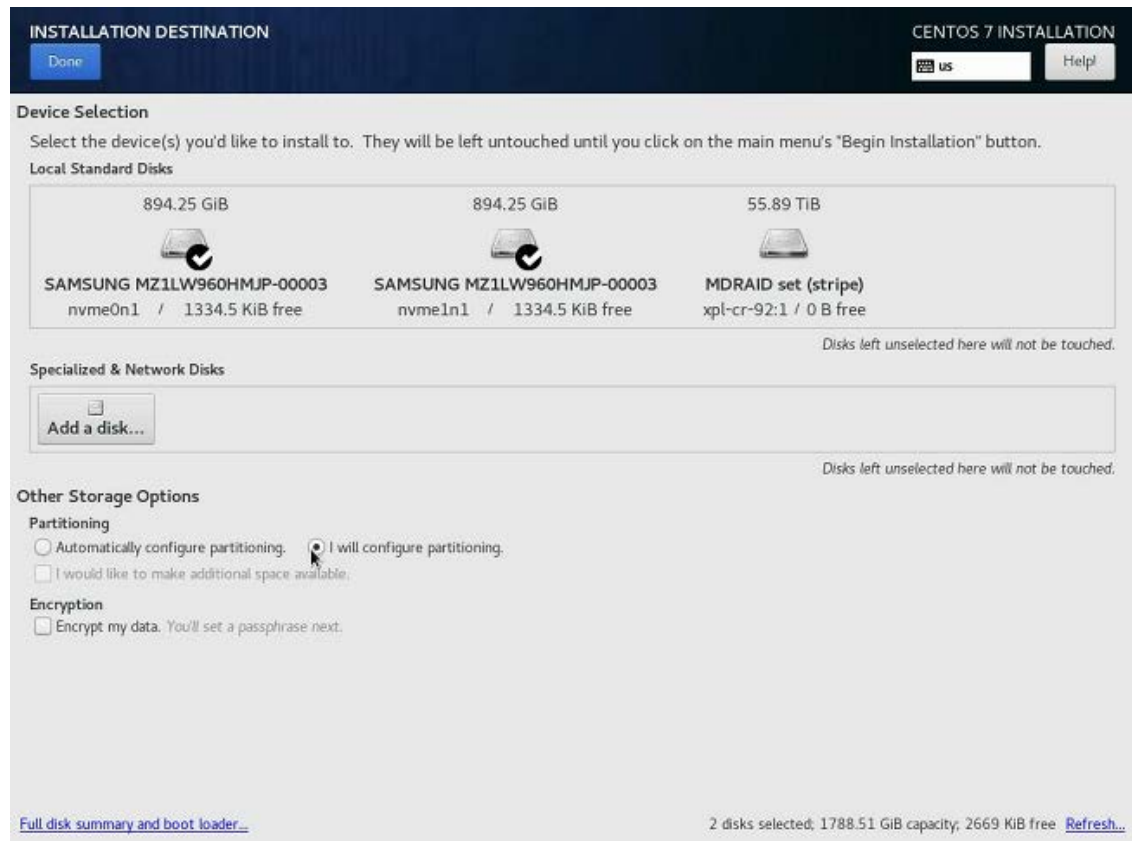
5. Set up the boot disks.

The DGX OS, which is installed in the factory before shipping out the DGX-2, creates a RAID-1 array on the two M.2 NVMe boot disks. Therefore, it is necessary to select both boot disks when installing RHEL 7 on a DGX-2.

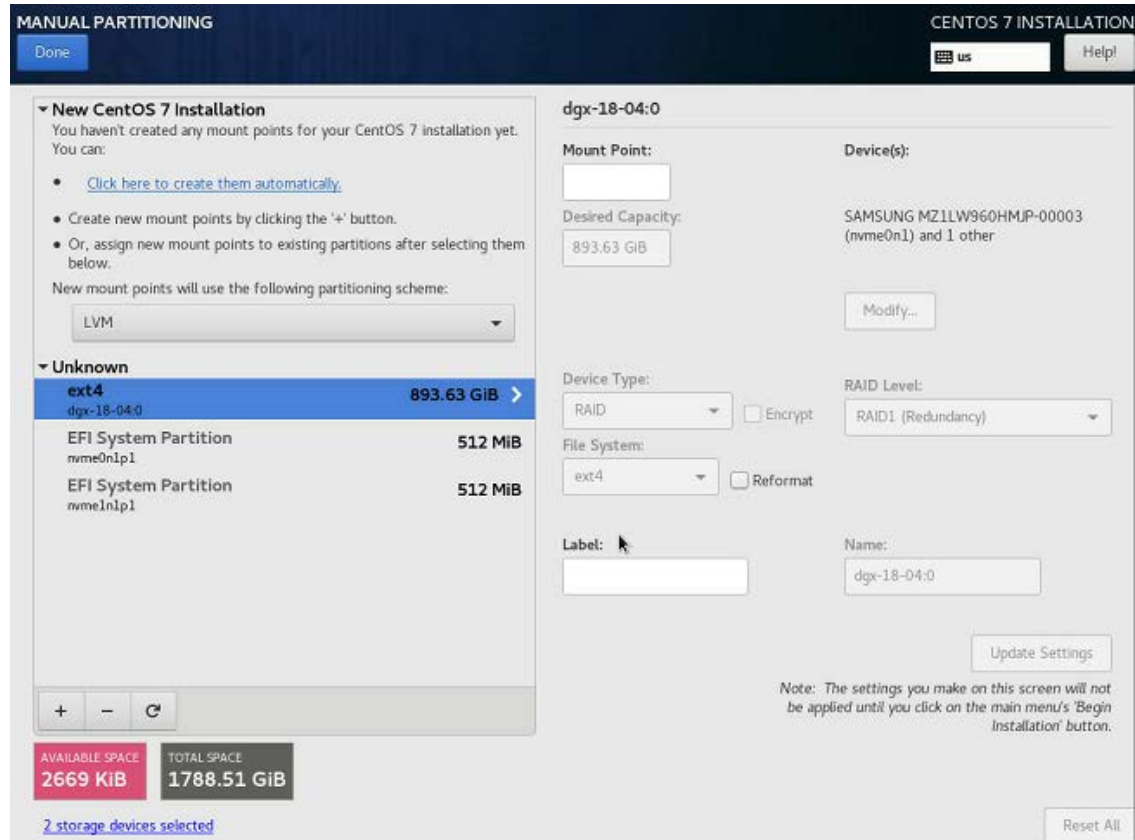
a). From the Installation Summary screen, click **INSTALLATION DESTINATION**.



- b). At the Installation Destination screen, select both M.2 disks (i.e. the 894.25 GB ones) and the **"I will configure partitioning"** radio button, then click **Done**.




- The Manual Partitioning window appears.
- c). Expand the **Unknown** drop-down menu.



This shows the disks that CentOS has identified. They will be

- ▶ the RAID-1 root partition (ext4, 893 GB),
 - ▶ the first EFI system partition (ESP) (nvme0n1p1, 512 MB), and
 - ▶ the second ESP (nvme1n1p1, 512 MB).
- d). Select the **ext4** partition and reformat it as **xfs**, then set the mount point to **/** and label it as **"md0"**

MANUAL PARTITIONING CENTOS 7 INSTALLATION

[Done](#)  [us](#) [Help!](#)

▼ New CentOS 7 Installation

You haven't created any mount points for your CentOS 7 installation yet. You can:

- [Click here to create them automatically.](#)
- Create new mount points by clicking the '+' button.
- Or, assign new mount points to existing partitions after selecting them below.

New mount points will use the following partitioning scheme:

LVM

▼ Unknown

| Partition | Size |
|-----------------------------------|-------------------|
| ext4 dgx-18-04:0 | 893.63 GiB |
| EFI System Partition nvme0n1p1 | 512 MiB |
| EFI System Partition nvme1n1p1 | 512 MiB |

+ - ↺

AVAILABLE SPACE: **2669 KiB** TOTAL SPACE: **1788.51 GiB**

[2 storage devices selected](#)

dgx-18-04:0

Mount Point:

Desired Capacity:

Device(s): SAMSUNG MZ1LW960HMP-00003 (nvme0n1) and 1 other

[Modify...](#)

Device Type: RAID ☐ Encrypt

RAID Level: RAID1 (Redundancy)

File System: xfs ☒ Reformat

Label:

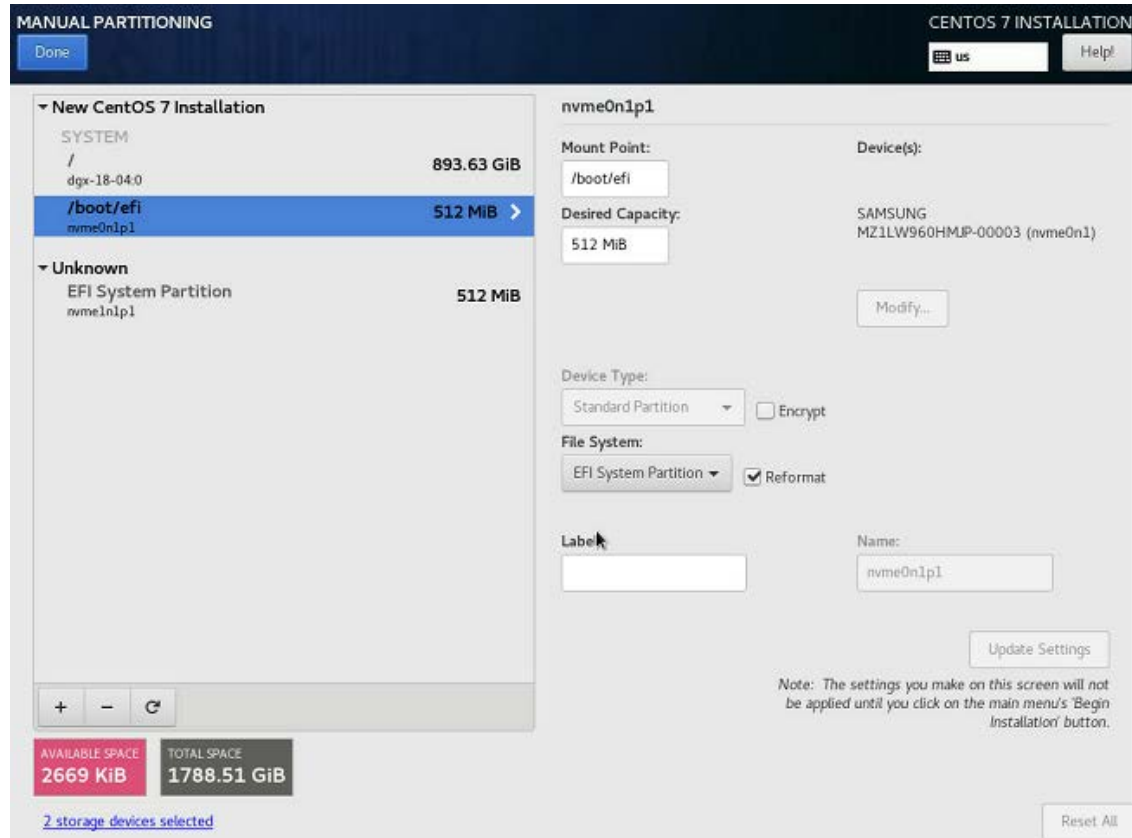
Name:

[Update Settings](#)

Note: The settings you make on this screen will not be applied until you click on the main menu's 'Begin Installation' button.

[Reset All](#)

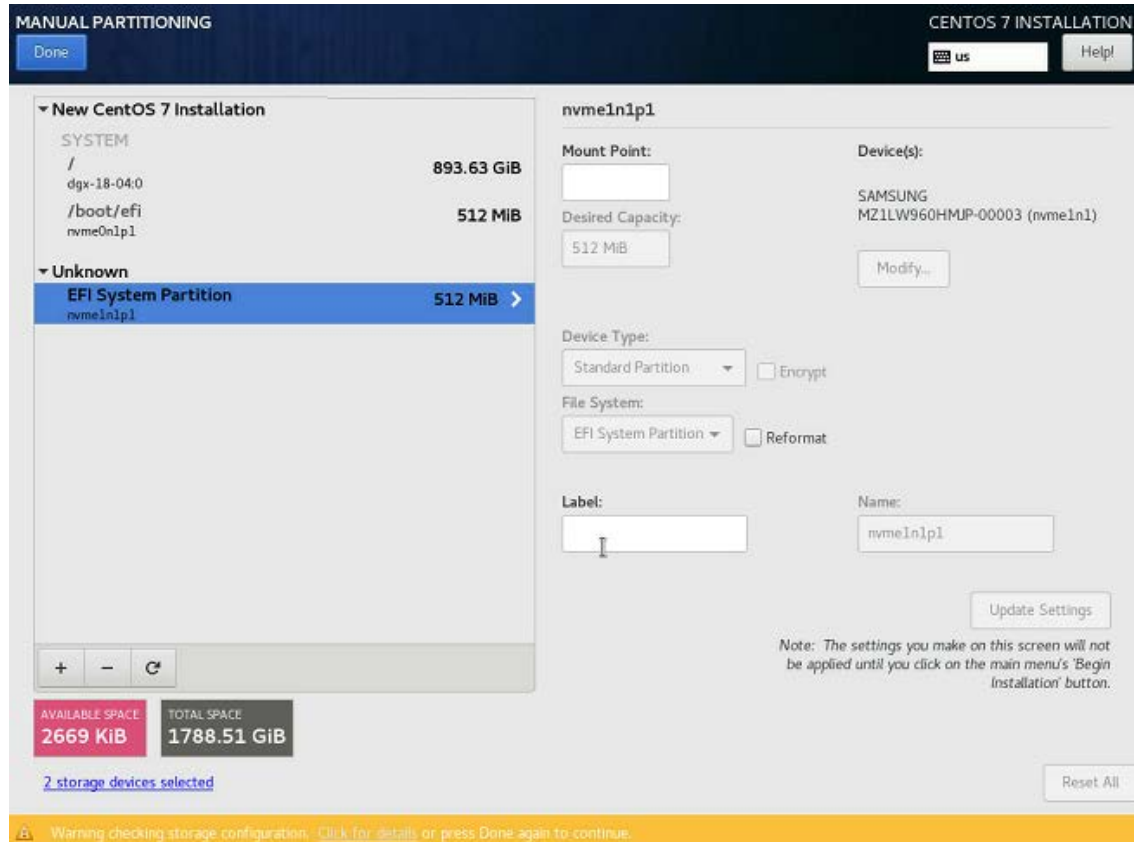
- e). Click **Update Settings**.
6. Create an EFI partition.
 - a). Select the **nvme0n1p1** partition and reformat it as an **EFI System Partition**, then set the mount point to `/boot/efi` and click **Update Settings**.



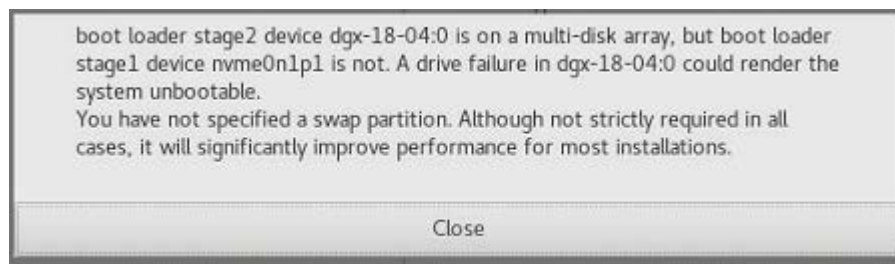
At this point, the two partitions you configured will be under the **New CentOS Installation** drop down menu, and the **nvme1n1p1** EFI partition will still be in **Unknown**. We can't create a second ESP because CentOS won't allow you to create two partitions mounted on the same mount point, i.e. `/boot/efi`, even if they are on different disks. You can replicate the ESP after installing the "DGX System Management" yum group, which provides a tool to accomplish this.

b). Click **Done**.

A yellow-highlighted warning appears on the bottom of the screen.

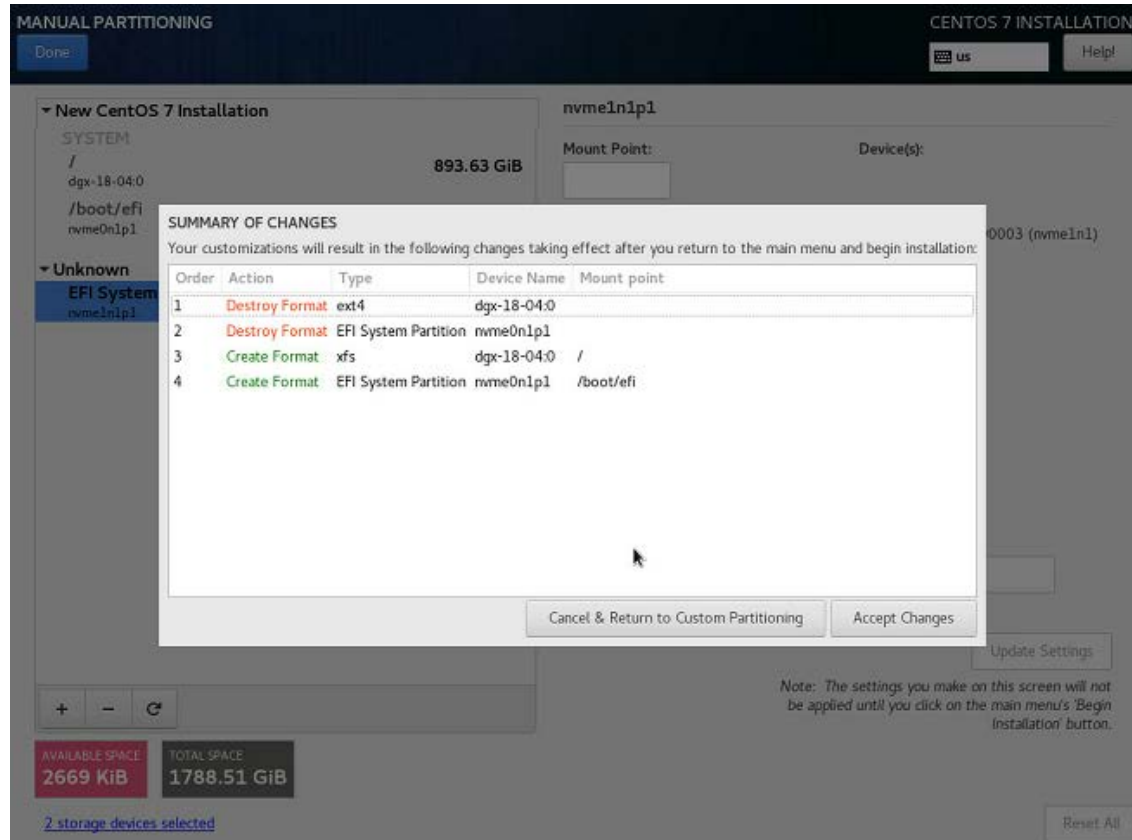


- c). Click the yellow-highlighted warning message.
The following message appears.



The message appears because the DGX OS from the factory has created a RAID 1 array for the root partition, and has replicated the EFI system partitions.

- d). Close the warning and then press **Done**.
A summary of the partitioning steps appears.



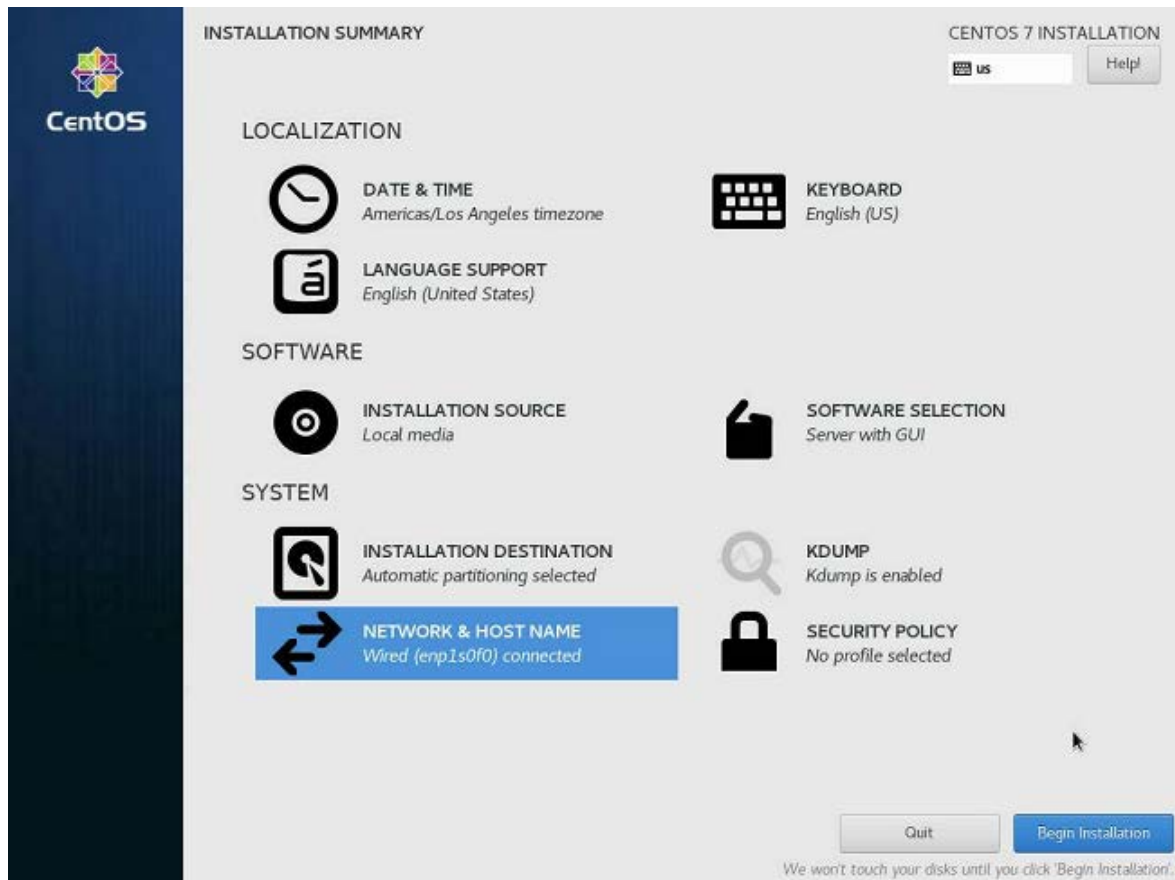
e). Click **Accept Changes** to get back to the main Installation Summary Screen.

There should no longer be a yellow warning icon on the **System: Installation Destination** option.

7. Configure Ethernet.

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, click **Begin installation** to start the installation.



- a). While the installation process is running, set your password (ROOT PASSWORD) and create a new user (USER CREATION) from the Configuration screen.
- b). When the installation completes, click **Reboot** to reboot the system.

2.4.3. Installing on the DGX A100

This section assumes you have booted the CentOS ISO image, either [locally](#) or [remotely](#).

For instructions applicable to the DGX-1 and the DGX Station, see [Installing on the DGX-1, DGX Station, or DGX Station A100](#).

For instructions applicable to the DGX-2, see [Installing on the DGX-2](#).

1. After booting the ISO image through either the BMC or from the USB drive, select **Install CentOS** and then press **Enter** to start the installation.



The CentOS installer starts.

For guidance on using the installer, refer to [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/).


2. Configure the language, region, date, time, keyboard, and other configuration options you may need from the **Installation Summary** screen.
3. Select the software packages that you want to install.

For detailed instructions, refer to [Software Selection in CentOS Installation Guide \(https://docs.centos.org/en-US/centos/install-guide/SoftwareSpoke-x86/\)](https://docs.centos.org/en-US/centos/install-guide/SoftwareSpoke-x86/).

4. 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**.



Note: Setting the correct Software Selection is critical for proper operation.



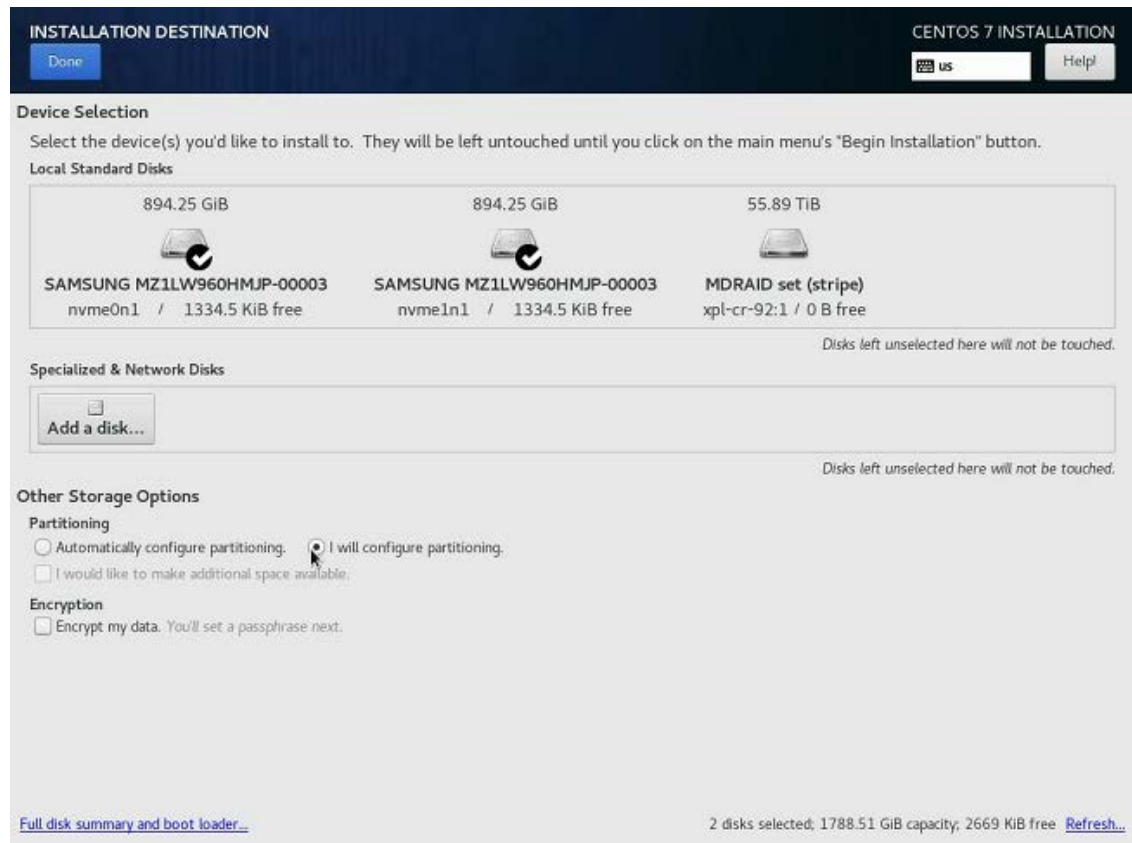
5. Set up the boot disks.

The DGX OS, which is installed in the factory before shipping out the DGX A100, creates a RAID-1 array on the two M.2 NVMe boot disks. Therefore, it is necessary to select both boot disks when installing RHEL 7 on a DGX A100.

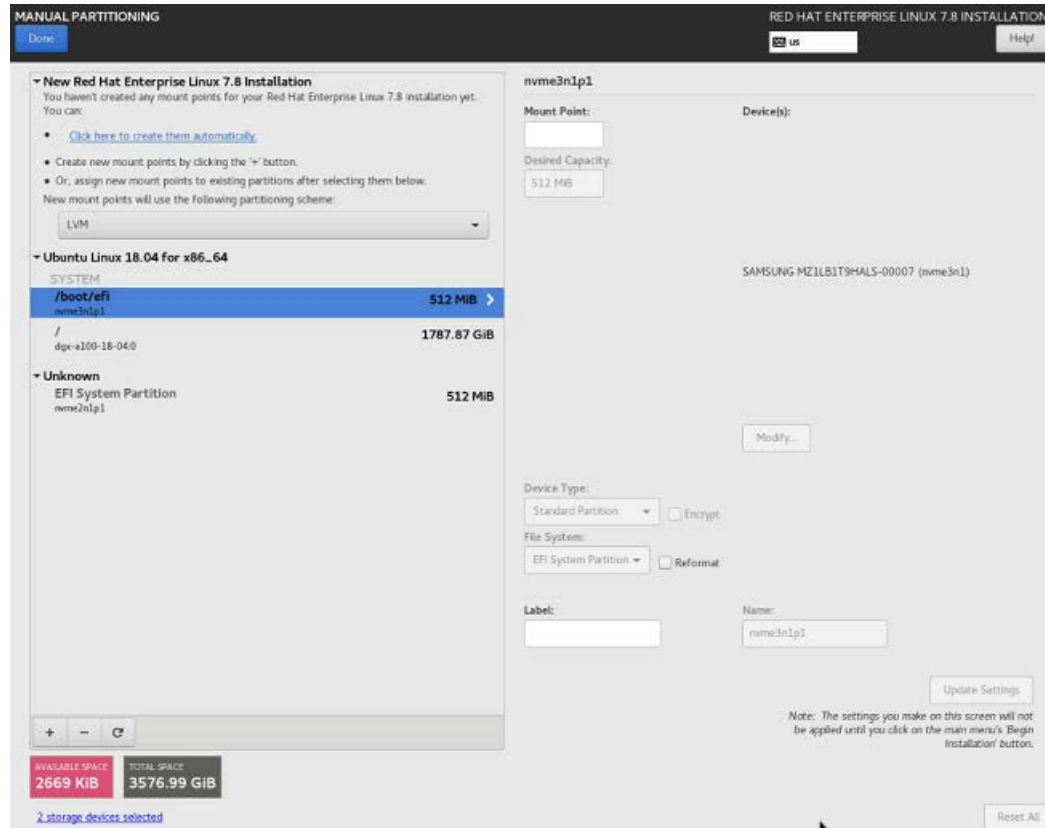
a). From the Installation Summary screen, click **INSTALLATION DESTINATION**.



- b). At the Installation Destination screen, select both M.2 disks (i.e. 1788 GiB on DGX A100) and the "I will configure partitioning" radio button, then click **Done**.

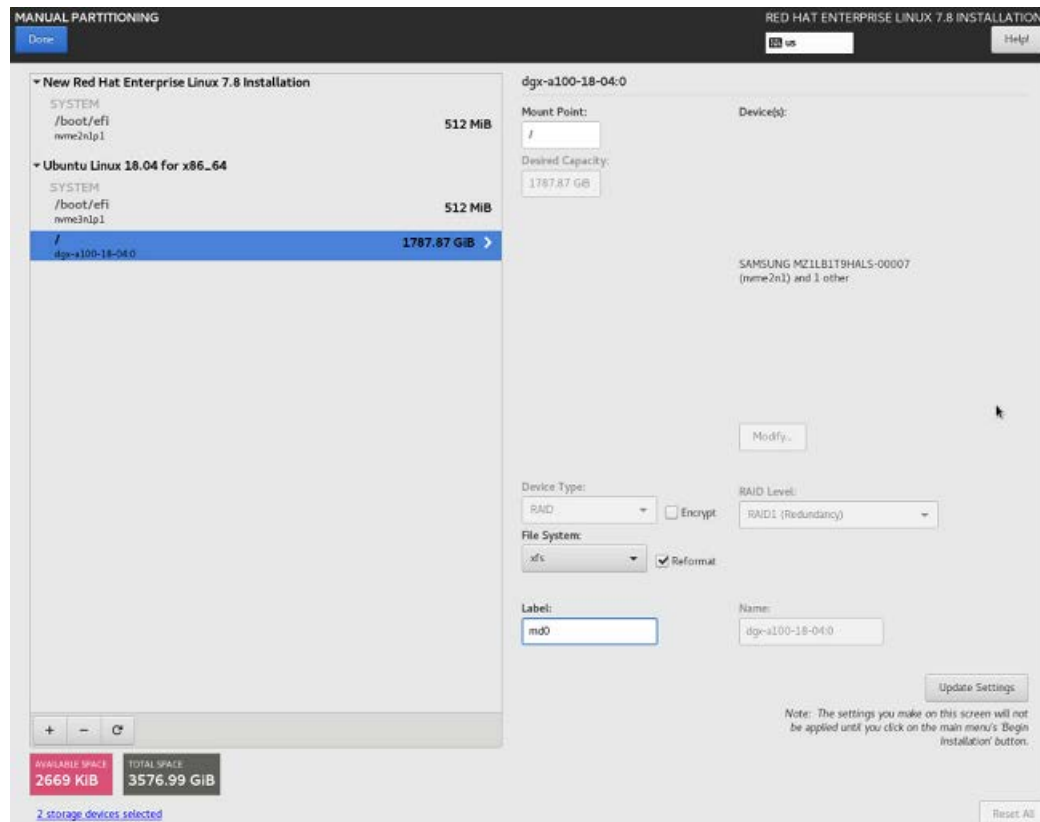


- The Manual Partitioning window appears.
- c). Expand the **Unknown** drop-down menu.

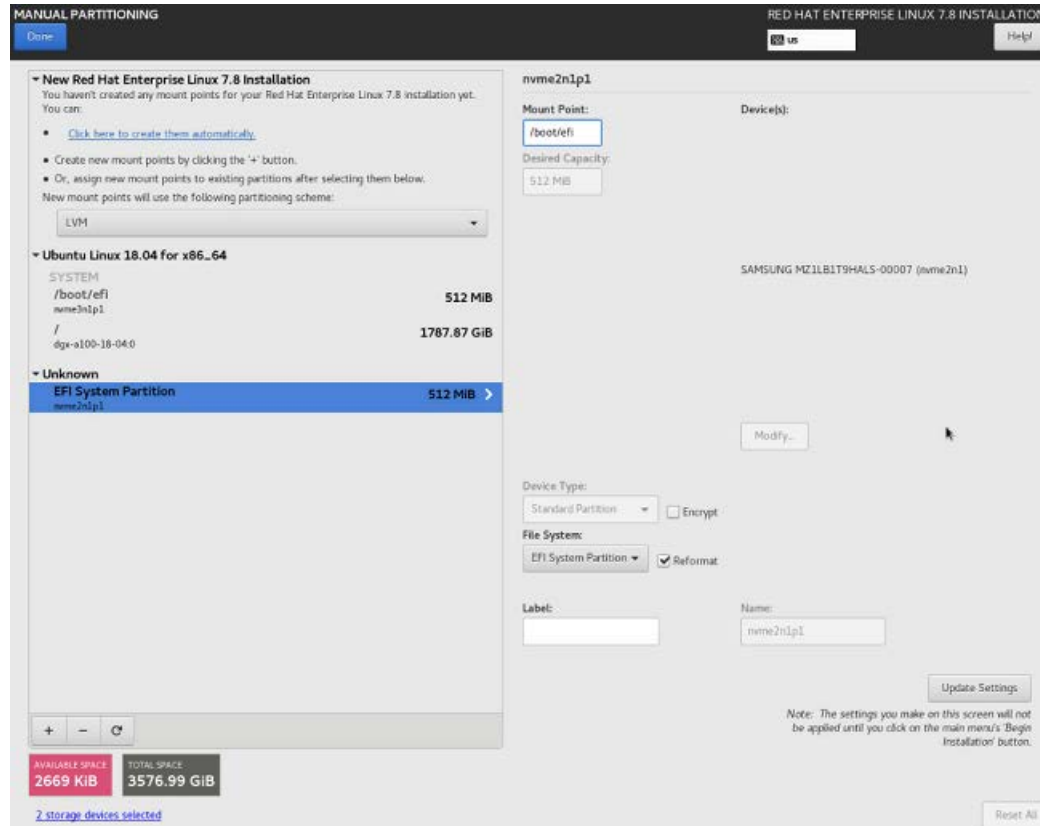


This shows the disks that CentOS has identified. They will be

- ▶ the RAID-1 root partition (ext4, 1787.87 GB),
 - ▶ the first EFI system partition (ESP) (nvme3n1p1, 512 MB), and
 - ▶ the second ESP (nvme2n1p1, 512 MB).
- d). Select the **ext4** partition and reformat it as **xfs**, then set the mount point to **/** and label it as **"md0"**



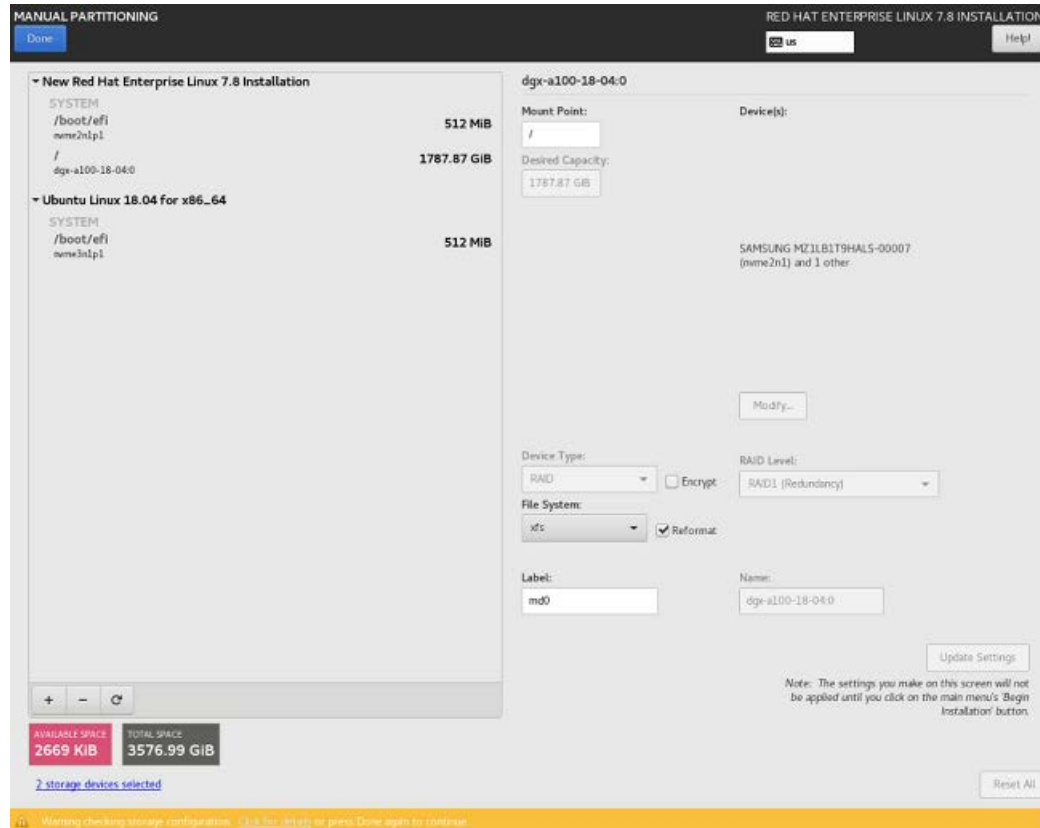
- e). Click **Update Settings**.
6. Create an EFI partition.
 - a). Select the second EFI partition, in the Unknown section (**nvme2n1p1** on DGX A100), reformat it as an **EFI System Partition**, then set the mount point to `/boot/efi` and click **Update Settings**.




At this point, the two partitions you configured will be under the **New CentOS Installation** drop down menu, and the Ubuntu EFI partition (**nvme3n1p1** on DGX A100) will be under **Ubuntu**. We can't create a second ESP because CentOS won't allow you to create two partitions mounted on the same mount point, i.e. `/boot/efi`, even if they are on different disks. You can replicate the ESP after installing the "DGX System Management" yum group, which provides a tool to accomplish this.

b). Click **Done**.

A yellow-highlighted warning appears on the bottom of the screen.

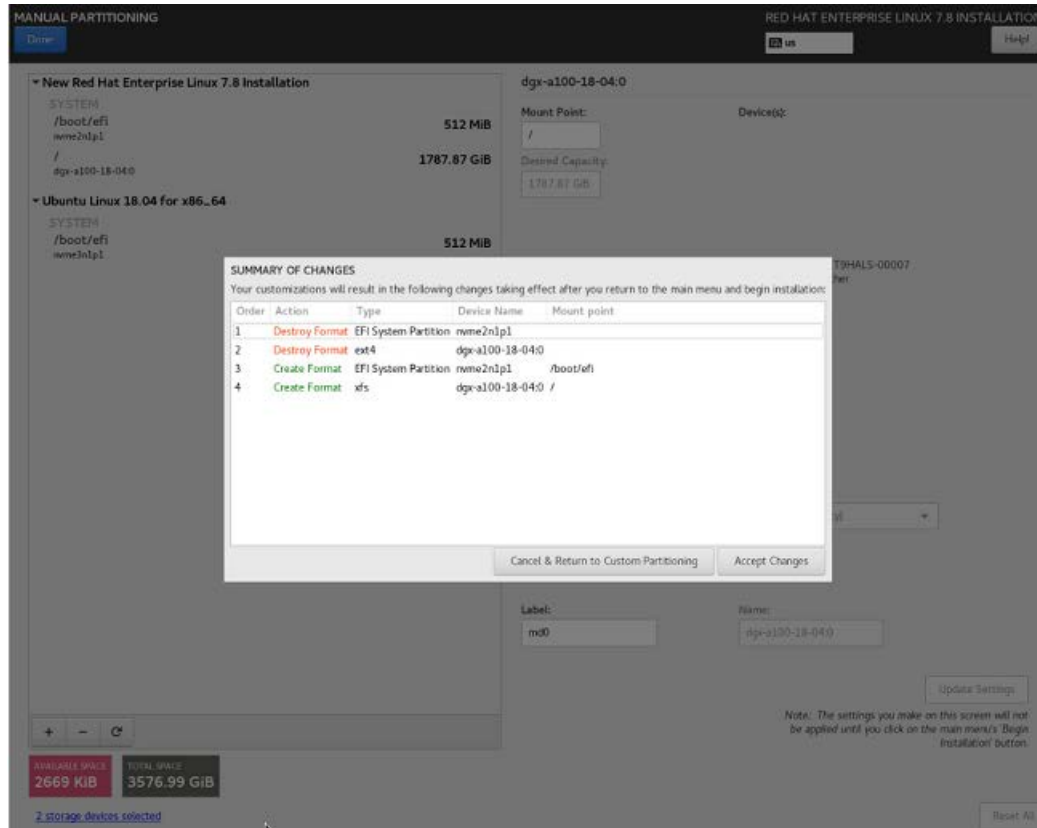


- c). Click the yellow-highlighted warning message.
The following message appears.



The message appears because the DGX OS from the factory has created a RAID 1 array for the root partition, and has replicated the EFI system partitions.

- d). Close the warning and then click **Done**.
A summary of the partitioning steps appears.



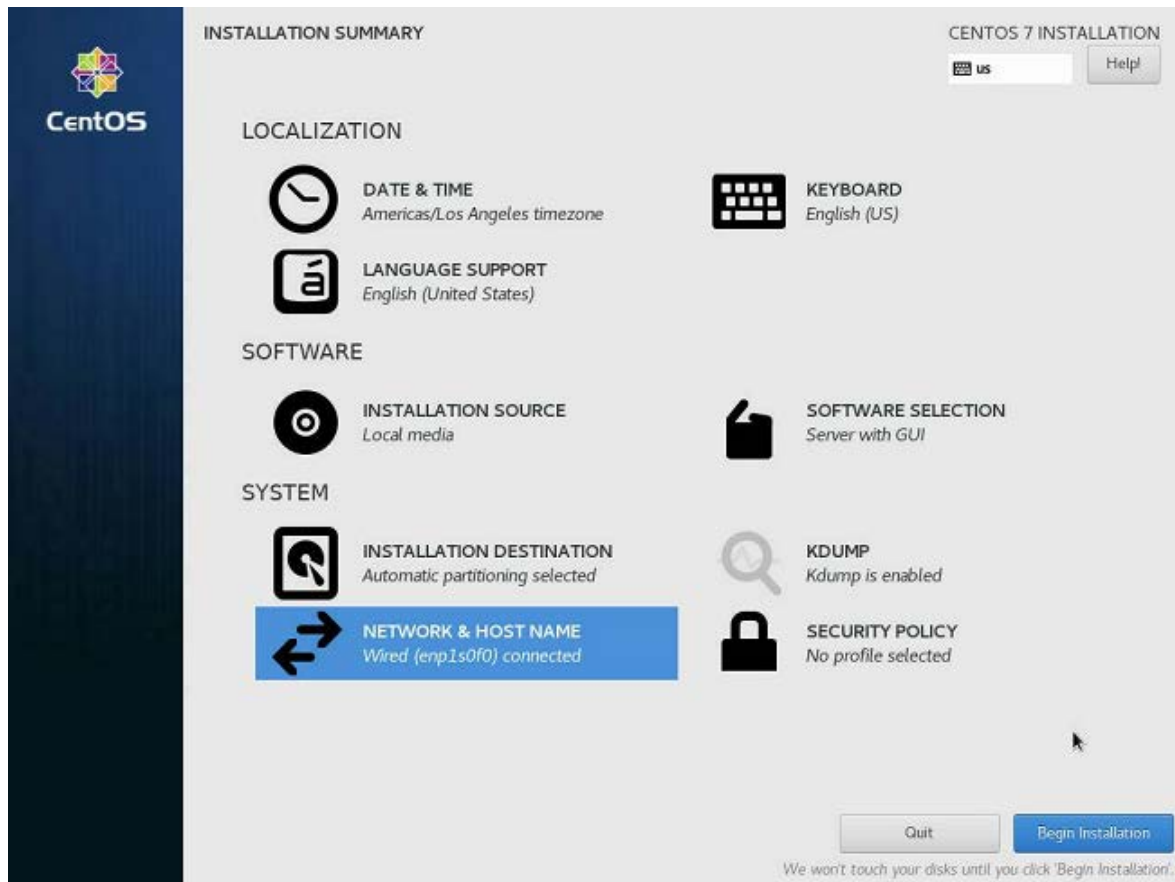
e). Click **Accept Changes** to get back to the main Installation Summary Screen.

There should no longer be a yellow warning icon on the **System: Installation Destination** option.

7. Configure Ethernet.

Select and enable the Ethernet device. This defaults to DHCP and can be changed for static IP configurations under **Configure**. This screen also lets you set the hostname.

8. From the INSTALLATION SUMMARY screen, click **Begin installation** to start the installation.



- While the installation process is running, set your password (ROOT PASSWORD) and create a new user (USER CREATION) from the Configuration screen.
- When the installation completes, click **Reboot** to reboot the system.

Chapter 3. Installing the DGX Software

This section requires that you have already installed CentOS" or derived operating system on the DGX server.

3.1. Configuring a System Proxy

If your network requires use of a proxy, then

- ▶ Edit the file `/etc/yum.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>
```

- ▶ Make sure that the following domains are not blocked and that the system can access them.
 - ▶ `cdn.redhat.com`
 - ▶ `international.download.nvidia.com`

3.2. Enabling the Repositories

1. Run the following commands to install the DGX software installation package and enable the NVIDIA 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 yum install -y https://international.download.nvidia.com/dgx/repos/rhel-
files/dgx-repo-setup-21.11-1.el7.x86_64.rpm
```

The next steps describe how to enable new driver branch repositories. Newer versions of other various packages can be made available through newer driver branch repositories.

For example, enabling rhel7-updates would make 2.8.3 the highest NCCL version available, and enabling rhel7-r450-cuda11-0 afterwards would bump the highest NCCL version to 2.11.4.

2. Enable the R450/CUDA 11.0 repository to move to the R450 package.

The step is required for DGX A100 and DGX Station A100, but optional for other DGX platforms.

- ▶ Either edit `/etc/yum.repos.d/nvidia-dgx-7.repo` and set `enabled=1`,

```
[nvidia-dgx-7-r450-cuda11-0]
name=NVIDIA DGX EL7 R450-CUDA11-0
baseurl=https://international.download.nvidia.com/dgx/repos/rhel7-r450-
cuda11-0/
enabled=1
gpgcheck=1
gpgkey=file:///etc/pki/rpm-gpg/RPM-GPG-KEY-dgx-cosmos-support
```

- ▶ Or (if you have the `yum-utils` package installed), issue the following.

```
sudo yum-config-manager --enable nvidia-dgx-7-r450-cuda11-0
```

3. (Optional) Enable the R470/CUDA 11.4 repository to move to the R470 package.

If you are moving from the R418 repo, then enable the R450 repo first as described in the previous step.

- ▶ Either edit `/etc/yum.repos.d/nvidia-dgx-7.repo` and set `enabled=1`,

```
[nvidia-dgx-7-r470-cuda11-4]
name=NVIDIA DGX EL7 R470-CUDA11-4
baseurl=https://international.download.nvidia.com/dgx/repos/rhel7-r470-
cuda11-4/
enabled=1
gpgcheck=1
gpgkey=file:///etc/pki/rpm-gpg/RPM-GPG-KEY-dgx-cosmos-support
```

- ▶ Or (if you have the `yum-utils` package installed), issue the following.

```
sudo yum-config-manager --enable nvidia-dgx-7-r470-cuda11-4
```

If building CUDA 11.4 Samples, install a version of GCC greater than or equal to 5.1.0. For instance, install GCC 9.1.1 by installing Red Hat Developer Toolset 9:

```
sudo yum install -y devtoolset-9
```

In this example, build CUDA samples using the following command in place of the “make” command:

```
scl enable devtoolset-9 make
```

3.3. Installing Required Components

3.3.1. Installing DGX Tools and Updating Configuration Files

1. Install the kernel development tools.
 - a). Install the kernel-devel package.

The kernel-devel package provides kernel headers required for the NVIDIA CUDA driver. Use the following command to install the kernel headers for the kernel version that is currently running on the system.

```
sudo yum install -y "kernel-devel-uname-r == $(uname -r)"
```

- b). Ensure that you have installed the latest version of gcc installed, as older versions may not support all of the features required to build the driver.

```
sudo yum install -y gcc
```

2. Install DGX tools and configuration files.

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

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

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

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

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

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

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

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

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

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

The configuration changes take effect only after rebooting the system, which is covered in the next step.

3. Update the kernel.

- a). Issue the following.

```
$ sudo yum update
```

Performing this update also updates the installed Red Hat Enterprise Linux 7 distribution to the latest version. To check the latest Red Hat Enterprise Linux 7 version, visit <https://access.redhat.com/articles/3078>.

- b). Reboot the server into the updated kernel.

```
$ sudo reboot
```

3.3.2. Configuring the /raid Partition

The DGX servers and the DGX Station include multiple SSDs for data caching or data storage. Configure these SSDs as a RAID array in a partition mounted at /raid. For the DGX servers, these SSDs are intended to be used as a data cache for NFS mounted directories. For the DGX Station, these SSDs are intended to be used either for local persistent storage or as a data cache for NFS mounted directories.

3.3.2.1. Configuring the /raid Partition as an NFS Cache

If you are using the data SSDs for caching NFS reads, configure these SSDs as a RAID 0 or RAID 5 array, mounted at /raid and update the `cachefilesd` configuration to use the /raid partition.

1. Configure the RAID array.

This will create the RAID group, mount it to `/raid`, and create an appropriate entry in `/etc/fstab`.

- ▶ To configure a RAID 0 array, issue the following:

```
sudo configure_raid_array.py -c -f
```

- ▶ To configure a RAID 5 array, issue the following:

```
sudo configure_raid_array.py -c -5 -f
```



Note:

The RAID array must be configured before installing `dgx-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 `dgx-conf-cachefilesd` has already been installed, be sure to restore the label manually before restarting `cachefilesd`.

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

2. Install `dgx-conf-cachefilesd` to update the `cachefilesd` configuration to use the `/raid` partition.

```
sudo yum install -y dgx-conf-cachefilesd
```

3.3.2.2. Configuring the `/raid` Partition for Local Persistent Storage

If you are using the data SSDs in the DGX Station for local persistent storage, configure these SSDs as a RAID 0 or RAID 5 array, mounted at `/raid`.

RAID 0 provides the maximum storage capacity, but does not provide any redundancy. If a single SSD in the array fails, all data stored on the array is lost. RAID 5 provides some level of protection against failure of a single SSD but with lower storage capacity than RAID 0.

- ▶ To configure a RAID 0 array, run the following command.

```
sudo configure_raid_array.py -c -f
```

- ▶ To configure a RAID 5 array, run the following command.

```
sudo configure_raid_array.py -c -f -5
```

These commands will create the RAID group, mount it to `/raid`, and create an appropriate entry in `/etc/fstab`.

3.3.3. Installing and Loading the NVIDIA CUDA Drivers

1. Install the driver package.

This will build and install the driver kernel modules. The installation of the `dkms-nvidia` package can take approximately five minutes.

```
sudo yum install -y cuda-drivers dgx-persistence-mode
```

2. **(DGX Station 100 Only)**: Install additional required Station A100 packages.

These packages must be installed after installing the cuda-drivers package.

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

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

```
sudo reboot
```

4. 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: Output from a DGX-1 system

| | | | | | | | | | |
|---|--|--------------------|--|-----------------|--|--------------|--|----------------------|--|
| NVIDIA-SMI 450.51.05 Driver Version: 450.51.05 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 | | ...00:06:00.0 | | Off | | 0 | |
| N/A 33C | | P0 45W / 300W | | 0MiB / 32480MiB | | 0% | | Default | |
| ----- | | | | | | | | | |
| 1 Tesla V100-SXM2... | | On | | ...00:07:00.0 | | Off | | 0 | |
| N/A 35C | | P0 44W / 300W | | 0MiB / 32480MiB | | 0% | | Default | |
| ----- | | | | | | | | | |
| : | | | | : | | : | | | |
| ----- | | | | | | | | | |
| 7 Tesla V100-SXM2... | | On | | ...00:8A:00.0 | | Off | | 0 | |
| N/A 34C | | P0 44W / 300W | | 0MiB / 32480MiB | | 0% | | Default | |
| ----- | | | | | | | | | |
| Processes: | | | | | | | | GPU Memory | |
| GPU | | PID | | Type | | Process name | | Usage | |
| ===== | | | | | | | | | |
| No running processes found | | | | | | | | | |
| ----- | | | | | | | | | |

3.3.4. Installing the NVIDIA Container Runtime

1. Install Docker 1.13 from the centos-extras repository.

```
sudo yum install -y docker
```

2. Install the **NVIDIA Container Runtime** group.

```
sudo yum groupinstall -y 'NVIDIA Container Runtime'
```

3. **(DGX Station A100 only):** Restart the `nv-docker-gpus` service.

```
$ sudo systemctl restart nv-docker-gpus
```

4. Run the following command to verify the installation.

- ▶ If the R418 driver package is installed:

```
sudo docker run --security-opt label=type:nvidia_container_t --rm nvcr.io/  
nvidia/cuda:10.0-base nvidia-smi
```

- ▶ If the R450 driver package is installed:

```
sudo docker run --security-opt label=type:nvidia_container_t --rm nvcr.io/  
nvidia/cuda:11.0-base nvidia-smi
```

The output should show all available GPUs.

Example: Output on a DGX -1 system

| | | | |
|------------|-----------|---------------------------|--------------------|
| NVIDIA-SMI | 450.51.05 | Driver Version: 450.51.05 | 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 | ...00:06:00.0 | Off | | 0 |
| N/A 33C P0 45W / 300W | 0MiB / 32480MiB | 0% | Default | | |
| 1 Tesla V100-SXM2... | On | ...00:07:00.0 | Off | | 0 |
| N/A 35C P0 44W / 300W | 0MiB / 32480MiB | 0% | Default | | |
| : | : | : | : | | |
| 7 Tesla V100-SXM2... | On | ...00:8A:00.0 | Off | | 0 |
| N/A 34C P0 44W / 300W | 0MiB / 32480MiB | 0% | Default | | |
| Processes: | | | | | |
| GPU | PID | Type | Process name | GPU Memory Usage | |
| No running processes found | | | | | |

See the section [Running Containers](#) for more information about this command. For a description of `nvcr.io`, see the [NGC Registry Spaces](#) documentation.

3.4. Installing Diagnostic Components

NVIDIA System Management (NVSM) provides the diagnostic components for NVIDIA DGX systems. NVSM is a software framework for monitoring NVIDIA DGX **server** nodes in a data center. It includes active health monitoring, system alerts, and log generation. The NVSM CLI can also be used for checking the health of and obtaining diagnostic information for DGX Station workstations.

Note: The diagnostic components for NVIDIA DGX systems require Python 3. It is available from the [CentOS Software Collections Repository \(https://wiki.centos.org/AdditionalResources/Repositories/SCL\)](https://wiki.centos.org/AdditionalResources/Repositories/SCL). The [Fedora EPEL](#) repository also contains a version of Python 3; however, this combination has not been tested.

1. Install the [CentOS Software Collections Repository \(https://wiki.centos.org/AdditionalResources/Repositories/SCL\)](https://wiki.centos.org/AdditionalResources/Repositories/SCL).

```
sudo yum install centos-release-scl
```

2. Enable the testing repository.

a). Install the yum utilities if you have not already done so.

```
sudo yum -y install yum-utils
```

b). Use yum-config-manager to enable the testing repository.

```
sudo yum-config-manager --enable centos-scl-rh-testing
```

3. Install Python 3.6.

```
sudo yum install -y rh-python36
```

Important: The diagnostic components for NVIDIA DGX systems are not supported with the `python3` package. Be sure to only install the `rh-python36` package.

4. Install the DGX System Management package group.

```
sudo yum groupinstall -y 'DGX System Management'
```

For information about using NVSM, see the [NVIDIA System Management](#) documentation.

3.5. Replicating the EFI System Partition on DGX-2 or DGX A100

This section applies only to the NVIDIA DGX-2 and DGX A100.

Once the 'DGX System Management' group is installed, the 'nvsm' tool can be used to replicate the EFI system partition (ESP) onto the second M.2 drive.



Important: Run these steps **ONLY IF**

- ▶ You are installing CentOS on the NVIDIA DGX-2 or DGX A100, and
- ▶ You installed CentOS on the RAID 1 array per instructions in the section [Installing on DGX-2](#) or [Installing on DGX A100](#) as appropriate.

1. Start the NVSM tool.

```
sudo nvsm
```

2. Navigate to /systems/localhost/storage/volumes/md0.

```
nvsm-> cd /systems/localhost/storage/volumes/md0
```

3. Start the rebuild process.

```
nvsm(/systems/localhost/storage/volumes/md0)-> start rebuild
```

- a). At the first prompt, specify the second M.2 disk.

```
PROMPT: In order to rebuild this volume, a spare drive
        is required. Please specify the spare drive to
        use to rebuild md0.
```

```
Name of spare drive for md0 rebuild (CTRL-C to cancel): nvme1n1
```

This should be the M.2 disk on which you did NOT install the ESP.

- ▶ If you followed the instructions in the section [Installing on DGX-2](#), this should be 'nvme1n1'
- ▶ If you followed the instructions in the section [Installing on DGX A100](#), this should be 'nvme2n1'

- b). At the second prompt, confirm that you want to proceed.

```
WARNING: Once the volume rebuild process is started, the
        process cannot be stopped.
```

```
Start RAID-1 rebuild on md0? [y/n] y
```

Upon successful completion, the following message should appear indicating that the ESP has been replicated:

```
/systems/localhost/storage/volumes/md0/rebuild started at 2019-03-07
14:40:55.844542
```

```
RAID-1 rebuild exit status: ESP_REBUILT
```

If necessary, the RAID 1 array is rebuilt after the ESP has been replicated.

```
Finished rebuilding RAID-1 on volume md0
```

```
100.0% [=====]
```

```
Status: Done
```


3.6. Installing Optional Components

The DGX is fully functional after installing the components as described in [Installing Required Components](#). If you intend to launch NGC containers (which incorporate the CUDA toolkit, NCCL, cuDNN, and TensorRT) on the DGX system, which is the expected use case, then you can skip this section.

If you intend to use your DGX system as a development system for running deep learning applications on bare metal, then install the optional components as described in this section.

1. Install the CUDA toolkit.

```
sudo yum install cuda
```

2. Install the NVIDIA Collectives Communication Library (NCCL) Runtime.

```
sudo yum groupinstall 'NVIDIA Collectives Communication Library Runtime'
```

3. Install the CUDA Deep Neural Networks (cuDNN) Library Runtime.

```
sudo yum groupinstall 'CUDA Deep Neural Networks Library Runtime'
```

4. Install NVIDIA TensorRT.

```
sudo yum install tensorrt
```

3.7. Applying an NVIDIA Look and Feel to the Desktop User Interface

If the GNOME Desktop is installed, you can optionally apply an NVIDIA look and feel to the desktop user interface by applying the NVIDIA theme to applications and the shell, and using NVIDIA images for the desktop background and lock screen.

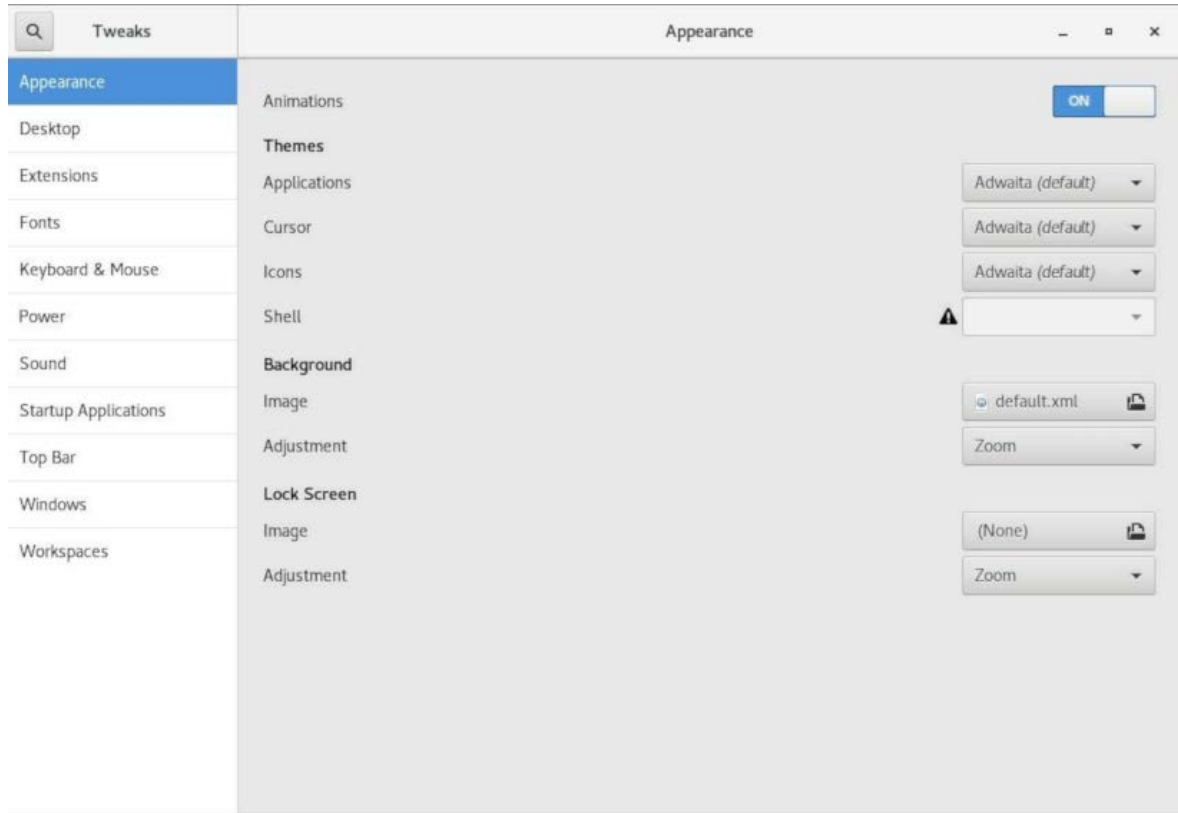
The GNOME Desktop must already be installed and running on your system. If **SOFTWARE SELECTION** was set to **Server with GUI** when you installed CentOS, the GNOME Desktop is already installed. If the GNOME Desktop is not installed, you must install the `x` Window System and GNOME package groups.

1. Install the DGX Desktop Theme package group.

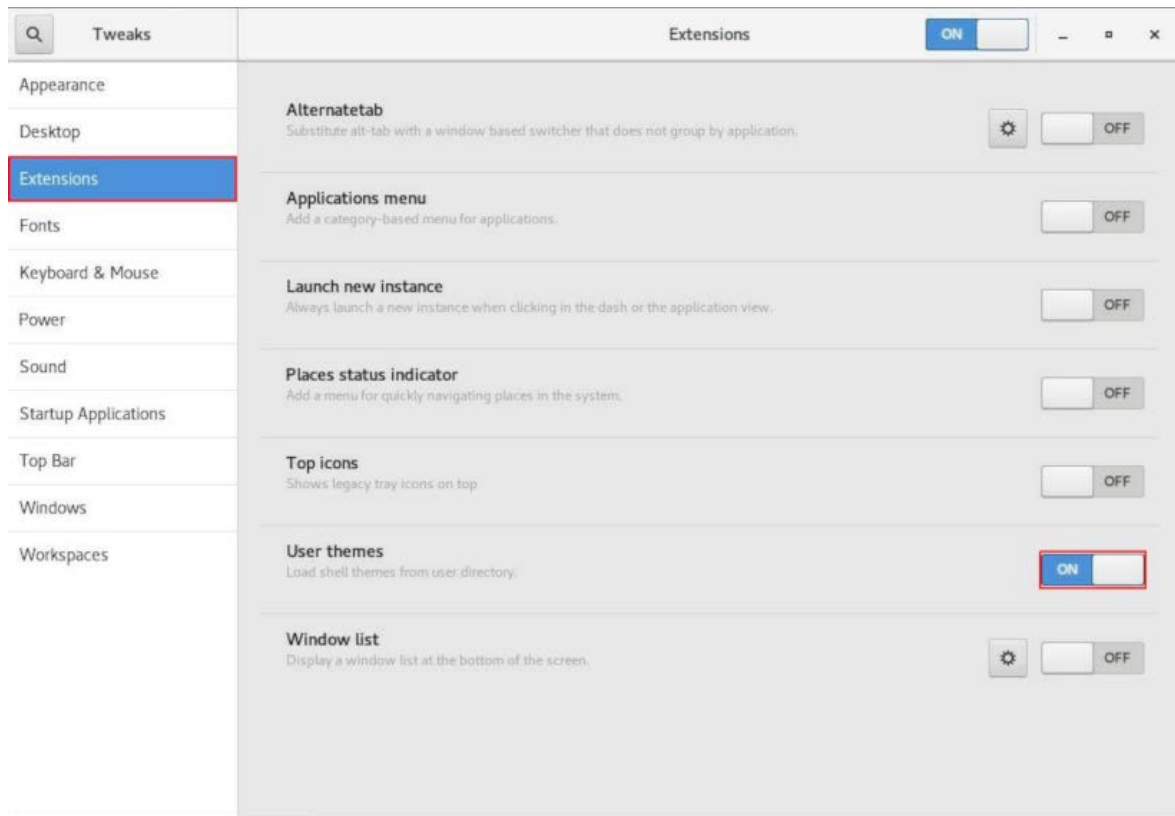
```
sudo yum groupinstall -y 'DGX Desktop Theme'
```

2. Start `gnome-tweaks`.

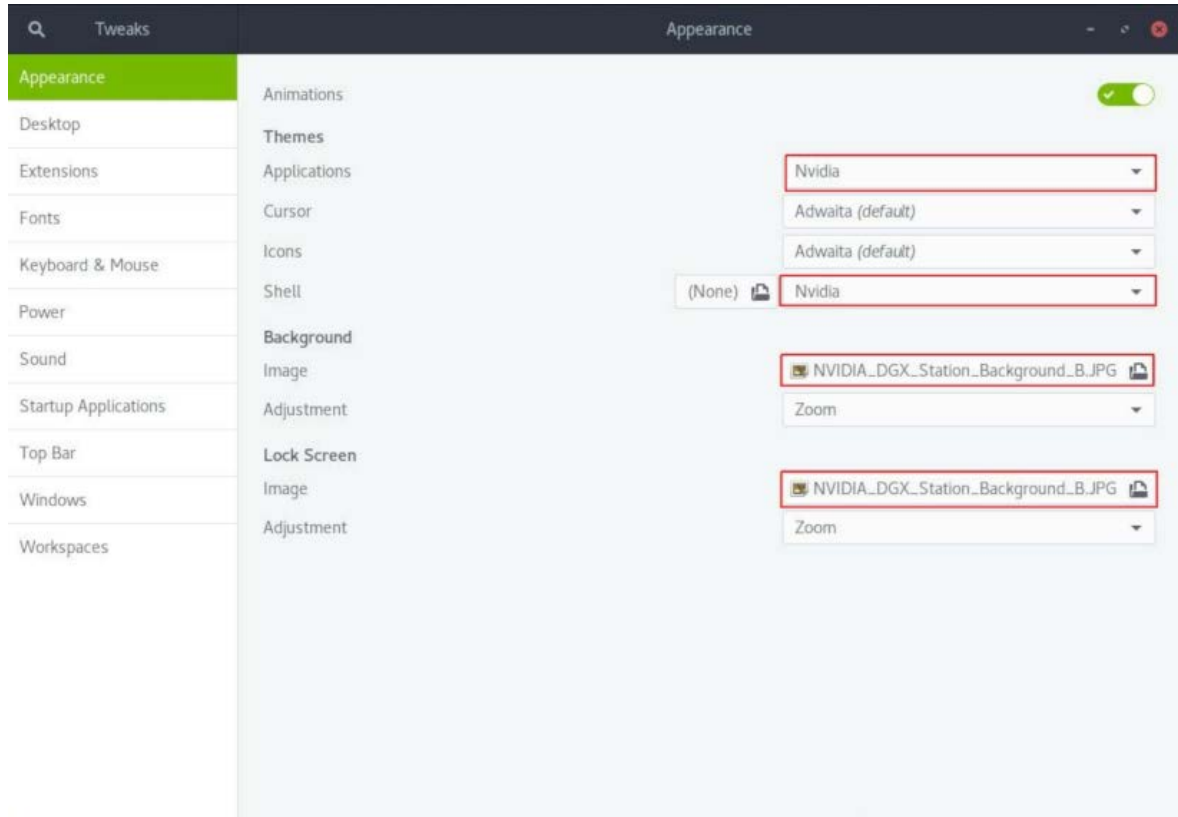
3. In the **Appearance** window that opens, under **Tweaks**, click **Extensions**.



4. In the **Extensions** window that opens, set **Extensions** in the title bar and **User themes** to **ON**.



5. Stop and restart `gnome-tweaks`.
6. In the **Appearance** window that opens, apply the NVIDIA theme to applications and the shell, and use NVIDIA images for the desktop background and lock screen.
 - a). Under **Themes**, in the drop-down lists for **Applications** and **Shell**, click **Nvidia**.
 - b). Under **Background** and **Lock Screen**, click the **Image** file selector.
 - c). In the **Image** window that opens, select an NVIDIA DGX Station background image, for example, `NVIDIA_DGX_Station_Background_B.JPG`, and click **Open**.



3.8. Managing CPU Mitigations

DGX Software for CentOS includes security updates to mitigate CPU speculative side-channel vulnerabilities. These mitigations can decrease the performance of deep learning and machine learning workloads.

If your installation of DGX systems incorporates other measures to mitigate these vulnerabilities, such as measures at the cluster level, you can disable the CPU mitigations for individual DGX nodes and thereby increase performance. This capability is available starting with DGX Software for CentOS version EL7-20.02.

3.8.1. Determining the CPU Mitigation State of the DGX System

If you do not know whether CPU mitigations are enabled or disabled, issue the following.

```
$ cat /sys/devices/system/cpu/vulnerabilities/*
```

- CPU mitigations are enabled if the output consists of multiple lines prefixed with Mitigation:.

Example


```
KVM: Mitigation: Split huge pages
Mitigation: PTE Inversion; VMX: conditional cache flushes, SMT vulnerable
Mitigation: Clear CPU buffers; SMT vulnerable
Mitigation: PTI
Mitigation: Speculative Store Bypass disabled via prctl and seccomp
Mitigation: usercopy/swapgs barriers and __user pointer sanitization
Mitigation: Full generic retpoline, IBPB: conditional, IBRS_FW, STIBP:
conditional, RSB filling
Mitigation: Clear CPU buffers; SMT vulnerable
```

- CPU mitigations are disabled if the output consists of multiple lines prefixed with Vulnerable.

Example

```
KVM: Vulnerable
Mitigation: PTE Inversion; VMX: vulnerable
Vulnerable; SMT vulnerable
Vulnerable
Vulnerable
Vulnerable: __user pointer sanitization and usercopy barriers only; no swapgs
barriers
Vulnerable, IBPB: disabled, STIBP: disabled
Vulnerable
```

3.8.2. Disabling CPU Mitigations



CAUTION: Performing the following instructions will disable the CPU mitigations provided by the DGX Software for CentOS.

1. Apply the `dgx*-no-mitigations` profile.

- On a DGX-2 system, issue

```
$ sudo tuned-adm profile dgx2-no-mitigations
```
- On a DGX-1 system, issue

```
$ sudo tuned-adm profile dgx-no-mitigations
```
- On a DGX Station workstation, issue

```
$ sudo tuned-adm profile dgxstation-no-mitigations
```

2. Reboot the system.

3. Verify CPU mitigations are disabled.

```
$ cat /sys/devices/system/cpu/vulnerabilities/*
```

The output should include several `vulnerable` lines. See [Determining the CPU Mitigation State of the DGX System](#) for example output.

3.8.3. Re-enabling CPU Mitigations

1. Apply the `dgx*-performance` package.

- On a DGX-2 system, issue

```
$ sudo tuned-adm profile dgx2-performance
```
- On a DGX-1 system, issue

```
$ sudo tuned-adm profile dgx-performance
```
- On a DGX Station workstation, issue


```
$ sudo tuned-adm profile dgxstation-performance
```

2. Reboot the system.
3. Verify CPU mitigations are enabled.

```
$ cat /sys/devices/system/cpu/vulnerabilities/*
```

The output should include several Mitigations lines. See [Determining the CPU Mitigation State of the DGX System](#) for example output.

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 EL7 software with a specific MLNX_OFED version. Consult the [NVIDIA EL7 release notes](#) for the recommended MLNX_OFED version to install for a particular version of NVIDIA EL7 software.

The following table provides a quick reference for tested versions.

| Red Hat Enterprise Linux Version | MLNX_OFED LTS Version |
|----------------------------------|-----------------------|
| 7.9 | 4.9-2.2.6.0 |
| 7.8 | 4.9-0.1.7.0 |
| 7.7 | 4.7-3.2.9.0 |

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 “yum 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 **LTS Download** tab.

NVIDIA EL7 software is tested only with LTS versions of MLNX_OFED.

- 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.

5. Install nvidia-mlnx-config.

```
sudo yum install -y nvidia-mlnx-config
```

6. Install kernel headers and development packages for your kernel.

These are needed for the ensuing DKMS compilation.

```
sudo yum install -y kernel-headers-$(uname -r) kernel-devel-$(uname -r)
```

7. After installing the MLNX_OFED drivers, install the NVIDIA peer memory module.

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

8. Load the nv_peer_mem module, either

- ▶ Manually, by issuing `sudo systemctl start nv_peer_mem`, or
- ▶ Set up the system to start it automatically on every system boot as follows.
 - a). Create a file `/etc/modules-load.d/nv-peer-mem.conf` with contents `"nv_peer_mem"`.
 - b). Issue `sudo dracut --force /boot/initramfs-$(uname -r).img $(uname -r)`

c). Reboot the system.



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](#), and
- ▶ [InfiniBand and RDMA Related Software Packages](#)

4.3. Updating the NVIDIA Mellanox InfiniBand Drivers

This section describes how to update MLNX_OFED on systems that already have it installed. 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, 7.9) will not require an update to the MLNX_OFED driver.

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

```
sudo yum update
```
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 **LTS Download** tab.

NVIDIA EL7 software is tested only with LTS versions of MLNX_OFED.

- 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 retry installing the MLNX_OFED driver.

8. Reboot.

```
sudo reboot
```

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

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

10. Load the `nv_peer_mem` module, either

- Manually, by issuing `sudo systemctl start nv_peer_mem`, or
- Set up the system to start it automatically on every system boot as follows.
 - a). Create a file `/etc/modules-load.d/nv-peer-mem.conf` with contents `"nv_peer_mem"`.
 - b). Issue `sudo dracut --force /boot/initramfs-$(uname -r).img $(uname -r)`
 - c). Reboot the system.

Chapter 5. Running Containers

The following is an example of running the CUDA container from the NGC registry.

```
sudo docker run --security-opt label=type:nvidia_container_t --rm nvcr.io/nvidia/cuda:10.0-base nvidia-smi
```



Note: To accommodate SELinux, the DGX software stack includes a package (`nvidia-container-selinux`) that defines a policy for allowing containers to access NVIDIA GPUs. The `--security-opt` option in the command sets the corresponding label type permitting the specified container to access NVIDIA GPUs. If SELinux is removed or disabled, then the `--security-opt` option is not needed.

Chapter 6. Configuring Storage - NFS Mount and Cache

The DGX servers and the DGX Station include multiple SSDs for data caching or data storage.

- ▶ The DGX-1 server contains four SSDs to be used in a RAID 0 configuration.
- ▶ The DGX-2 server contains 8 or 16 SSDs to be used in a RAID 0 configuration.
- ▶ The DGX A100 server contains 4 or 8 SSDs to be used in a RAID 0 configuration.
- ▶ The DGX Station contains three SSDs to be used in a RAID 0 or RAID 5 configuration.

For the DGX servers, these SSDs are intended to be used as a data cache for NFS mounted directories. For the DGX Station, these SSDs are intended to be used either for local persistent storage or as a data cache for NFS mounted directories. If you are using these SSDs as a data cache for NFS mounted directories, you must set up your own NFS mounted directories for long-term data storage.

The following instructions describe how to mount the NFS directories onto the DGX system, and how to cache the NFS using the DGX SSDs for improved performance.

Make sure that you have an NFS server with one or more exports with data to be accessed by the DGX System, and that there is network access between the DGX System and the NFS server.

1. Install the nfs packages.

```
$ sudo yum install nfs-utils
```

2. Configure an NFS mount for the DGX.

- a). Edit the filesystem tables configuration.

```
sudo vi /etc/fstab
```

- b). Add a new line for the NFS mount, using the local mount point of /mnt.

```
<nfs_server>:<export_path> /mnt nfs  
rw,noatime,rsize=32768,wsizes=32768,nolock,tcp,intr,fsc,nofail 0 0
```

/mnt is used here as an example mount point.

Consult your Network Administrator for the correct values for <nfs_server> and <export_path>.

The nfs arguments presented here are a list of recommended values based on typical use cases. However, "fsc" must always be included as that argument specifies use of FS-Cache.

c). Save the changes.

3. Verify the NFS server is reachable.

```
ping <nfs_server>
```

Use the server IP address or the server name provided by your network administrator.

4. Mount the NFS export.

```
sudo mount /mnt
```

/mnt is the example mount point used in step 1.

5. Verify caching is enabled.

```
cat /proc/fs/nfsfs/volumes
```

Look for the text `FSC=yes` in the output. The NFS will be mounted and cached on the DGX System automatically upon subsequent reboot cycles.

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.

The screenshot displays the NVIDIA DGX-1 BMC web interface. The top navigation bar includes links for Dashboard, Server Information, Server Health, Configuration, Remote Control, Maintenance, Firmware Update, and HELP. The Configuration menu is expanded, showing options like Active Directory, DNS, LDAP/E-Directory, Mouse Mode, Network, LAN Port, PEF, RADIUS, Remote Session, Services, SMTP, SOL, SSL, Users, Virtual Media, and SNMP Community. The main content area is divided into two sections: Device Information and Sensor Monitoring.

Device Information

Dashboard gives the overall information about the status of server.

Device Information

Firmware Revision: 3.20.30
Firmware Build Time: Sep 30 2017 14:25:46 CST
BIOS Version: SZW_3A05
Expander Firmware Version: 0600
BMC NIC: Dedicated-NIC
System MAC
- 1 : OCP Mezz :A8:1E:84:C3:5F:3F
- 2 : OCP Mezz :A8:1E:84:C3:5F:40
BMC Date & Time: Thu, 13 Dec 2018 22:46:14
BMC Chipset: AST2400
PSU1 MFR Revision: 01
PSU1 Firmware Version:00.03.07
PSU2 MFR Revision: 01
PSU2 Firmware Version:00.03.07
PSU3 MFR Revision: 01
PSU3 Firmware Version:00.03.07
PSU4 MFR Revision: 01
PSU4 Firmware Version:00.03.07
GPU0 Market Name: Tesla V100-SXM2-32GB
GPU0 GPU Part Number:1DB5-896-A1

Sensor Monitoring

| Status | Sensor | Reading | |
|--------|--------------|----------|---|
| ● | Temp_Outlet | 38 ° C | 🔍 |
| ● | Temp_GPGPU0 | 33 ° C | 🔍 |
| ● | Temp_GPGPU1 | 36 ° C | 🔍 |
| ● | Temp_GPGPU2 | 35 ° C | 🔍 |
| ● | Temp_GPGPU3 | 33 ° C | 🔍 |
| ● | Temp_GPGPU4 | 35 ° C | 🔍 |
| ● | Temp_GPGPU5 | 35 ° C | 🔍 |
| ● | Temp_GPGPU6 | 37 ° C | 🔍 |
| ● | Temp_GPGPU7 | 35 ° C | 🔍 |
| ● | Power_GPGPU0 | 44 Watts | 🔍 |
| ● | Power_GPGPU1 | 44 Watts | 🔍 |
| ● | Power_GPGPU2 | 42 Watts | 🔍 |

3. Add a new user.
 - a). Select an empty field and click Add User.



qct.admin (Administrator) Refresh Print Logout

Dashboard Server Information Server Health Configuration Remote Control

Maintenance Firmware Update HELP

User Management

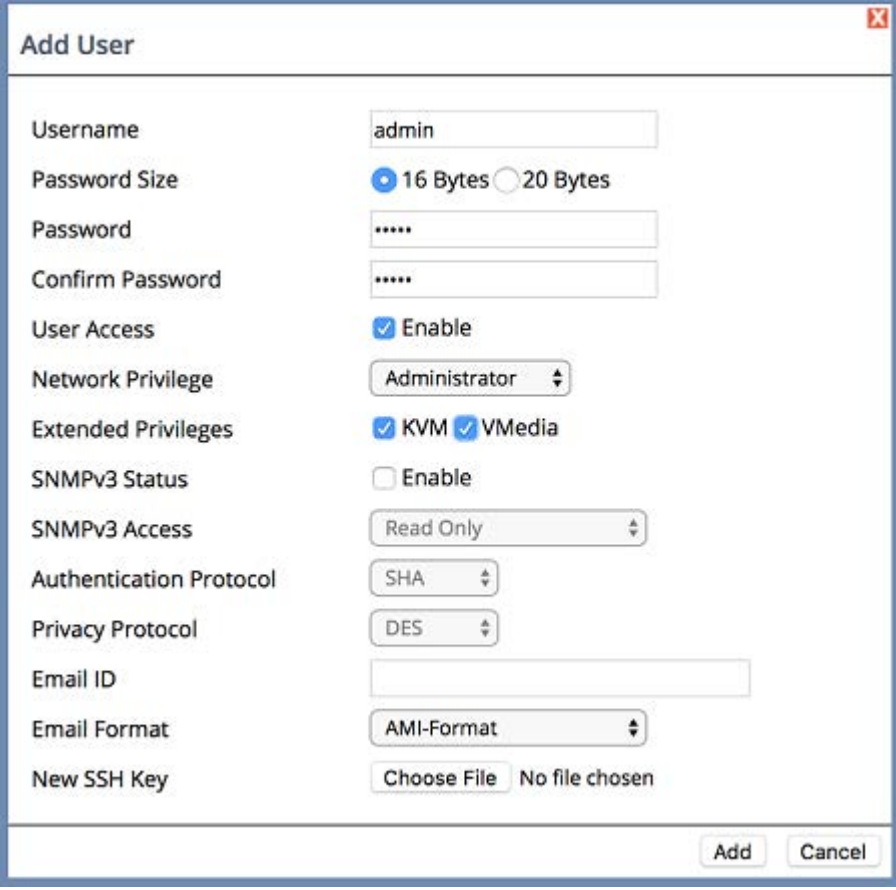
The list below shows the current list of available users. To delete or modify a user, select the user name from the list and click "Delete User" or "Modify User". To add a new user, select an unconfigured slot and click "Add User"

Number of configured users: 2

| UserID | Username | User Access | Network Privilege | SNMPv3 Status | Email ID |
|--------|-----------|-------------|-------------------|---------------|----------|
| 1 | anonymous | Disabled | Administrator | Disabled | ~ |
| 2 | qct.admin | Enabled | Administrator | Disabled | ~ |
| 3 | ~ | ~ | ~ | ~ | ~ |
| 4 | ~ | ~ | ~ | ~ | ~ |
| 5 | ~ | ~ | ~ | ~ | ~ |
| 6 | ~ | ~ | ~ | ~ | ~ |
| 7 | ~ | ~ | ~ | ~ | ~ |
| 8 | ~ | ~ | ~ | ~ | ~ |
| 9 | ~ | ~ | ~ | ~ | ~ |
| 10 | ~ | ~ | ~ | ~ | ~ |

Add User Modify User Delete User

- b). Enter new user information and click Add.

A screenshot of a web-based 'Add User' dialog box. The dialog has a title bar with a close button (X) in the top right corner. The main area contains a list of configuration options for a new user. The 'Username' field is set to 'admin'. 'Password Size' has two radio buttons: '16 Bytes' (selected) and '20 Bytes'. 'Password' and 'Confirm Password' fields are masked with dots. 'User Access' is checked with a blue checkbox and labeled 'Enable'. 'Network Privilege' is a dropdown menu showing 'Administrator'. 'Extended Privileges' has two checked checkboxes: 'KVM' and 'VMedia'. 'SNMPv3 Status' is unchecked. 'SNMPv3 Access' is a dropdown menu showing 'Read Only'. 'Authentication Protocol' is a dropdown menu showing 'SHA'. 'Privacy Protocol' is a dropdown menu showing 'DES'. 'Email ID' is an empty text field. 'Email Format' is a dropdown menu showing 'AMI-Format'. 'New SSH Key' has a 'Choose File' button and the text 'No file chosen'. At the bottom right, there are 'Add' and 'Cancel' buttons.

| | |
|-------------------------|--|
| Username | admin |
| Password Size | <input checked="" type="radio"/> 16 Bytes <input type="radio"/> 20 Bytes |
| Password | ***** |
| Confirm Password | ***** |
| User Access | <input checked="" type="checkbox"/> Enable |
| Network Privilege | Administrator |
| Extended Privileges | <input checked="" type="checkbox"/> KVM <input checked="" type="checkbox"/> VMedia |
| SNMPv3 Status | <input type="checkbox"/> Enable |
| SNMPv3 Access | Read Only |
| Authentication Protocol | SHA |
| Privacy Protocol | DES |
| Email ID | |
| Email Format | AMI-Format |
| New SSH Key | Choose File No file chosen |

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



| | | | | | |
|--|--------------------|---------------|---------------|----------------|------|
| admin (Administrator) Refresh Print Logout | | | | | |
| Dashboard | Server Information | Server Health | Configuration | Remote Control | |
| Maintenance | Firmware Update | | | | HELP |

User Management

The list below shows the current list of available users. To delete or modify a user, select the user name from the list and click "Delete User" or "Modify User". To add a new user, select an unconfigured slot and click "Add User"

Number of configured users: 3

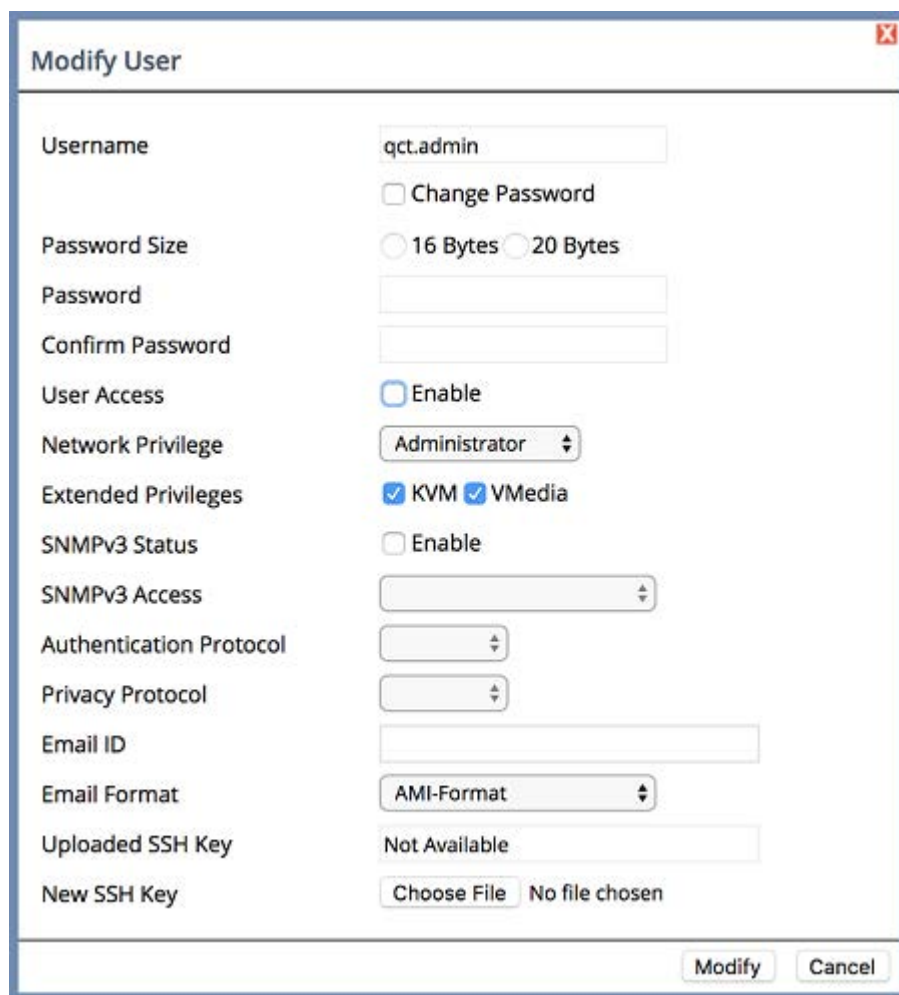
| UserID | Username | User Access | Network Privilege | SNMPv3 Status | Email ID |
|--------|-----------|-------------|-------------------|---------------|----------|
| 1 | anonymous | Disabled | Administrator | Disabled | ~ |
| 2 | qct.admin | Enabled | Administrator | Disabled | ~ |
| 3 | admin | Enabled | Administrator | Disabled | ~ |
| 4 | ~ | ~ | ~ | ~ | ~ |
| 5 | ~ | ~ | ~ | ~ | ~ |
| 6 | ~ | ~ | ~ | ~ | ~ |
| 7 | ~ | ~ | ~ | ~ | ~ |
| 8 | ~ | ~ | ~ | ~ | ~ |
| 9 | ~ | ~ | ~ | ~ | ~ |
| 10 | ~ | ~ | ~ | ~ | ~ |

Add User

Modify User

Delete User

b). Deselect Enable in User Access and click Modify.



Modify User

| | |
|-------------------------|--|
| Username | qct.admin |
| | <input type="checkbox"/> Change Password |
| Password Size | <input type="radio"/> 16 Bytes <input type="radio"/> 20 Bytes |
| Password | <input type="password"/> |
| Confirm Password | <input type="password"/> |
| User Access | <input checked="" type="checkbox"/> Enable |
| Network Privilege | Administrator |
| Extended Privileges | <input checked="" type="checkbox"/> KVM <input checked="" type="checkbox"/> VMedia |
| SNMPv3 Status | <input type="checkbox"/> Enable |
| SNMPv3 Access | |
| Authentication Protocol | |
| Privacy Protocol | |
| Email ID | |
| Email Format | AMI-Format |
| Uploaded SSH Key | Not Available |
| New SSH Key | <input type="button" value="Choose File"/> No file chosen |

c). Ensure User Access is disabled for the user qct.admin.



admin (Administrator) Refresh Print Logout

Dashboard Server Information Server Health Configuration Remote Control

Maintenance Firmware Update HELP

User Management

The list below shows the current list of available users. To delete or modify a user, select the user name from the list and click "Delete User" or "Modify User". To add a new user, select an unconfigured slot and click "Add User"

Number of configured users: 3

| UserID | Username | User Access | Network Privilege | SNMPv3 Status | Email ID |
|--------|-----------|-------------|-------------------|---------------|----------|
| 1 | anonymous | Disabled | Administrator | Disabled | ~ |
| 2 | qct.admin | Disabled | Administrator | Disabled | ~ |
| 3 | admin | Enabled | Administrator | Disabled | ~ |
| 4 | ~ | ~ | ~ | ~ | ~ |
| 5 | ~ | ~ | ~ | ~ | ~ |
| 6 | ~ | ~ | ~ | ~ | ~ |
| 7 | ~ | ~ | ~ | ~ | ~ |
| 8 | ~ | ~ | ~ | ~ | ~ |
| 9 | ~ | ~ | ~ | ~ | ~ |
| 10 | ~ | ~ | ~ | ~ | ~ |

Add User

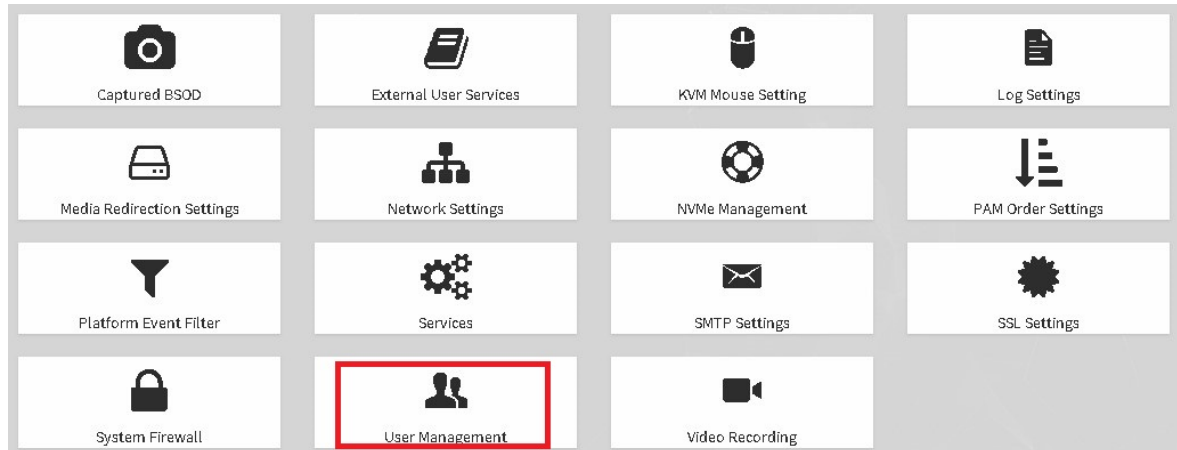
Modify User

Delete User

7. Log out.

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

- Log into the BMC.
 - Open a browser within your LAN and go to `https://<BMC-ip-address>/`.
 - Log in, using admin/admin for the User ID/Password.
- Select **Settings** from the left-side navigation menu.
- 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. Using Custom DGX Software Utilities for the DGX Station

The DGX Software includes custom utilities for maintaining the DGX Station persistent storage. Custom utilities for managing and obtaining diagnostic information for the DGX Station were included **only** in version EL7-20.01 of the DGX Software.

B.1. Rebuilding or Re-Creating the DGX Station RAID Array

Failure of a single drive in a RAID 5 array is a recoverable error but the failure causes data redundancy for the array to be lost. After replacing a single failed SSD in a RAID 5 array, you must rebuild the array to restore data redundancy for the array. Failure of any number of SSDs in a RAID 0 array and failure of more than one SSD in a RAID 5 array are both unrecoverable failures. After replacing the SSDs in response to an unrecoverable failure, you must re-create the array.

If the DGX Station RAID array is degraded because one or more SSDs failed, replace each failed SSD as explained in [DGX Station User Guide](#).

The DGX Station software includes the custom script `configure_raid_array.py` for rebuilding or re-creating the RAID array.

- To rebuild a RAID 5 array after replacing a single failed SSD, run the following command:

```
$ sudo configure_raid_array.py -r
```



Note: The time required to rebuild a RAID 5 array depends on factors such as system load, SSD capacity, and the number of SSDs in the array. Rebuilding the array of three 1.92-terabyte SSDs in the DGX Station may require several hours.

You can monitor the progress of a long-running rebuild by examining the contents of the `/proc/mdstat` file:

```
$ cat /proc/mdstat
Personalities : [raid0] [linear] [multipath] [raid1] [raid6] [raid5] [raid4]
               [raid10]
```



```
md0 : active raid5 sdb[0] sdd[3] sdc[1]
      3750486016 blocks super 1.2 level 5, 512k chunk, algorithm 2 [3/2] [UU_]
      [>.....] recovery = 4.0% (75580956/1875243008)
      finish=438.3min speed=68419K/sec
      bitmap: 2/14 pages [8KB], 65536KB chunk

unused devices: <none>
```

In this example, the rebuild is 4.0% complete and the rebuild is estimated to finish in 438.3 minutes.

- ▶ To re-create a RAID 5 array after replacing more than one failed SSD, run the following command:

```
$ sudo configure_raid_array.py -c -5 -f
```



CAUTION: Specify the `-c` option only if an unrecoverable failure, such as the failure of more than one SSD, has occurred. The `-c` option erases all data in the array.

- ▶ To re-create a RAID 0 array after replacing any number of failed SSDs, run the following command:

```
$ sudo configure_raid_array.py -c -f
```

The RAID array is rebuilt or re-created with the RAID level that you specified.

- ▶ If you re-created a RAID 0 or RAID 5 array, all data that was on the array is erased after array is re-created.
- ▶ If you rebuilt a RAID 5 array, the data on the array is preserved after array is rebuilt.

If you have re-created a RAID 0 or RAID 5 array and have a backup of data on the array that you want to preserve, restore the data from the backup.

B.2. Changing the RAID Level of the RAID Array

During the initial installation of the DGX software on CentOS, the data SSDs in the DGX Station are configured as a RAID 0 or RAID 5 array. If your requirements for redundancy or storage capacity change, you can change the RAID level of the array from the level that was initially configured.

Before changing the RAID level of the DGX Station RAID array, back up all data on the array that you want to preserve. Changing the RAID level of the DGX Station RAID array erases all data stored on the array.

The DGX Station software includes the custom script `configure_raid_array.py`, which you can use to change the level of the RAID array without unmounting the RAID volume.

- ▶ To change the RAID level to RAID 5, run the following command:


```
$ sudo configure_raid_array.py -m raid5
```

Note:

After you change the RAID level to RAID 5, the RAID array is rebuilt. A RAID array that is being rebuilt is online and ready to be used, but a check on the health of the DGX Station reports the status of the RAID volume as unhealthy. Therefore, avoid checking the health of the DGX Station while the RAID array is being rebuilt. For more information, see [EL7-20.01 Only: Checking the Health of the DGX Station](#).

The time required to rebuild the RAID array depends on the workload on the system. On an idle system, the rebuild might be complete within 30 minutes.

- To change the RAID level to RAID 0, run the following command:

```
$ sudo configure_raid_array.py -m raid0
```

To confirm that the RAID level was changed as required, run the `lsblk` command. The entry in the `TYPE` column for each SSD in the RAID array indicates the RAID level of the array.

The following example shows that the RAID level of the array is RAID 0. The name of the RAID volume is `md0` and the mount point of the volume is `/raid`.

```
~$ lsblk
NAME        MAJ:MIN RM  SIZE RO TYPE  MOUNTPOINT
sda          8:0    0  1.8T  0 disk
└─sda1       8:1    0   487M  0 part  /boot/efi
└─sda2       8:2    0  1.8T  0 part  /
sdb          8:16   0  1.8T  0 disk
└─md0        9:0    0  5.2T  0 raid0 /raid
sdc          8:32   0  1.8T  0 disk
└─md0        9:0    0  5.2T  0 raid0 /raid
sdd          8:48   0  1.8T  0 disk
└─md0        9:0    0  5.2T  0 raid0 /raid
```

B.3. EL7-20.01 Only: Checking the Health of the DGX Station



Note: Starting with release EL7-20.02, the NVIDIA System Health Checker (`nvhealth`) tool is replaced by NVIDIA System Management (NVSM). For information about how to use NVSM to perform this task, see [Show Health in NVIDIA System Management User Guide](#).

The DGX Station provides the NVIDIA System Health Checker (`nvhealth`) tool to exercise the system and verify its health. The output of `nvhealth` is an itemized list of checks and their status, typically Healthy or Unhealthy. On a healthy system, all checks should return Healthy. You should investigate any checks that return Unhealthy to determine their root cause and resolve them.

To check the health of the DGX Station, run the following command:

```
$ sudo nvhealth [-k output-file]
```


output-file

The name and the path of the file in which the raw state of the system is written. The `nvhealth` command displays this file name at the end of the output from the command.

If you omit the output file, the information is written to the file `/tmp/nvhealth-log.random-string.jsonl`, for example, `/tmp/nvhealth-log.6wf3WriAC3.jsonl`.

**Note:**

If you run the `nvhealth` command while the RAID array is being rebuilt after a change in RAID level to RAID 5, `nvhealth` reports the status of the RAID volume as unhealthy. To avoid this potentially misleading result, wait until RAID array is rebuilt before running `nvhealth`.

To check the progress of the rebuild and show the percentage complete and an estimate of the time to completion, run this command:

```
# cat /proc/mdstat

Personalities : [raid6] [raid5] [raid4] [linear] [multipath] [raid0] [raid1]
[raid10]
md0 : active raid5 sdb[0] sdc[1] sdd[2]
      181764096 blocks super 1.2 level 5, 512k chunk, algorithm 2 [4/3] [UUU_]
      [==>.....]    recovery = 17.2% (10426232/60588032)
      finish=45.8min speed=18238K/sec
```

B.4. EL7-20.01 Only: Collecting Information for Troubleshooting the DGX Station



Note: Starting with release EL7-20.02, the tool to collect troubleshooting information (`nvsysinfo`) tool is replaced by NVIDIA System Management (NVSM). For information about how to use NVSM to perform this task, see [Dump Health in NVIDIA System Management User Guide](#).

To help diagnose and resolve issues, the DGX Station provides a tool to collect troubleshooting information for NVIDIA Support Enterprise Services.

The tool verifies basic functionality and performance of the DGX Station and collects the following information in an xz-compressed tar archive:

- ▶ Log files
- ▶ Hardware inventory
- ▶ SW inventory

To collect information for troubleshooting the DGX Station, run the following command:

```
sudo nvsysinfo [-o output-file]
```


output-file

The path of the file in which the information is written.

If you omit the output file, the name of the file to which the information is written is `/tmp/nvsysinfo-host-name-timestamp.tar.xz`.

Use any method that is convenient for you to send the file to NVIDIA Support Enterprise Services. For example, send the file as an e-mail attachment.

Appendix C. Expanding the DGX Station RAID Array

After adding SSDs to the DGX Station, you must expand the RAID array to add the new SSDs to the array. The procedure for expanding the RAID array is the same for all supported RAID levels.

Add the extra SSDs to the DGX Station as explained in [DGX Station User Guide](#).

Because expanding a RAID array risks loss of data, ensure that you have a backup of data on the array that you want to preserve.

This task requires `sudo` privileges.

Use standard Linux operating system commands to expand the DGX Station RAID array.

1. Obtain the device IDs of the SSDs that were added by searching for 1.8 T drives that aren't mounted in the output from the `lsblk` (<http://manpages.ubuntu.com/manpages/bionic/man8/lsblk.8.html>) command.

```
$ lsblk
```

In the following example, the device IDs of the SSDs that were added are `sde`, `sdf`, `sdg`, and `sdh`.

```
$ lsblk
NAME        MAJ:MIN RM  SIZE RO TYPE  MOUNTPOINT
sda          8:0    0   1.8T 0 disk
├─sda1       8:1    0   487M 0 part  /boot/efi
├─sda2       8:2    0   1.8T 0 part  /
sdb          8:16   0   1.8T 0 disk
├─md0        9:0    0   5.2T 0 raid0 /raid
sdc          8:32   0   1.8T 0 disk
├─md0        9:0    0   5.2T 0 raid0 /raid
sdd          8:48   0   1.8T 0 disk
├─md0        9:0    0   5.2T 0 raid0 /raid
sde          8:64   0   1.8T 0 disk
sdf          8:80   0   1.8T 0 disk
sdg          8:96   0   1.8T 0 disk
sdh          8:112  0   1.8T 0 disk
```

2. Add the SSDs that you added to the DGX Station to the RAID array.

```
$ sudo mdadm --add raid-device-path ssd-device-paths
```

raid-device-path

The device path of the RAID array, for example, `/dev/md0`.

ssd-device-paths

A space-separated list of the device paths of the SSDs that you added, in which each path is of the form `/dev/device-id`.

This example adds the SSDs with device paths `/dev/sde`, `/dev/sdf`, `/dev/sdg`, and `/dev/sdh` to the RAID array with device path `/dev/md0`.

```
sudo mdadm --add /dev/md0 /dev/sde /dev/sdf /dev/sdg /dev/sdh
```

3. Increase the number of devices in the RAID array to 7.

```
$ sudo mdadm --grow --raid-devices=7 raid-device-path
```

raid-device-path

The device path of the RAID array, for example, `/dev/md0`.



Note: Increasing the number of devices in the RAID array to 7 may require several hours or even longer. If the system crashes, is shut down, or is rebooted while the number of devices on the array is being increased, all data that was on the array is erased.

This example increases the number of devices in the RAID array with device path `/dev/md0` to 7.

```
$ sudo mdadm --grow --raid-devices=7 /dev/md0
```

4. Resize the file system that resides on the RAID array to use the additional physical space in the array.

```
$ sudo resize2fs raid-device-path
```

raid-device-path

The device path of the RAID array, for example, `/dev/md0`.

This example resizes the file system that resides on the RAID array with device path `/dev/md0`.

```
$ sudo resize2fs /dev/md0
```

The RAID array is expanded with its existing RAID level. The data on the array is preserved, even if the array is a RAID 0 array.

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, 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.