



# DGX Software for Red Hat Enterprise Linux 8

Release Notes

# Table of Contents

<b>Chapter 1. DGX Software For Red Hat Enterprise Linux 8 Overview.....</b>	<b>1</b>
1.1. Current Versions.....	1
1.2. Installing the DGX Software Stack on Red Hat Enterprise Linux 8.....	1
1.3. Upgrading the DGX Software Stack and Red Hat Enterprise Linux 8.....	2
1.3.1. Upgrading the Software without Moving to a New Driver Branch.....	2
1.3.2. Upgrading and Moving to a New Driver Branch on non-NVSwitch Systems.....	2
1.3.3. Updating and Moving to a New Driver Branch on NVSwitch Systems.....	3
<b>Chapter 2. Version EL8-22.08.....</b>	<b>5</b>
<b>Chapter 3. Version EL8-22.06.....</b>	<b>9</b>
<b>Chapter 4. Version EL8-22.04.....</b>	<b>12</b>
<b>Chapter 5. Version EL8-22.02.....</b>	<b>15</b>
<b>Chapter 6. Version EL8-21.08.....</b>	<b>18</b>
<b>Chapter 7. Version EL8-20.11.....</b>	<b>22</b>
<b>Chapter 8. Known Issues.....</b>	<b>25</b>
8.1. [DGX-1, DGX-2]: DGCM Update Error.....	25
8.2. [DGX-1, DGX-2]: GPU MIG Partitions do not return output fields.....	25
8.3. [DGX-1, DGX-2]: Log displays CEC error.....	26
8.4. [DGX-1]: NVSM show controllers SerialNumber shows `NOT_SET` .....	26
8.5. [DGX-1, DGX-2]: nysys fails to launch.....	26
8.6. [DGX-1, DGX-2]: nvsm dump health Does not Generate sosreport.....	27
8.7. [DGX-2]: No Rebuild Function for RAID 0 if Volume is not md1.....	27
8.8. [DGX-2]: Storage Alerts Persist from Previous RAID Configuration.....	28
<b>Chapter 9. Rotating the GPG Key.....</b>	<b>29</b>

---

# Chapter 1. DGX Software For Red Hat Enterprise Linux 8 Overview

NVIDIA provides a NVIDIA® DGX™ software stack targeted for installation on DGX systems that have been user-installed with Red Hat Enterprise Linux. The software stack provides the same features and functionality that are provided by the original DGX OS Server and DGX OS Desktop software built on the Ubuntu operating system. See also the [DGX Software on Red Hat Enterprise Linux 8 Installation Guide](#).

## 1.1. Current Versions

The following are the current versions available.

Product	Current Version
DGX Station	<a href="#">EL8-22.08</a>
DGX-2, DGX-1	<a href="#">EL8-22.08</a>
DGX A100	<a href="#">EL8-22.08</a>
DGX Station A100	<a href="#">EL8-22.08</a>

## 1.2. Installing the DGX Software Stack on Red Hat Enterprise Linux 8



**WARNING:** A recently released RHEL kernel upgrade is incompatible with the MOFED driver. Customers should refrain from upgrading their systems until the next MOFED release, and wait for the MOFED update before installing RHEL if IB is required.



**Important:** Before you install or perform the upgrade, refer to the EL8-22.08 section for information about the latest release.

To install the software on a fresh DGX system, see the [DGX Software for Red Hat Enterprise Linux 8 - Installation Guide](#).

## 1.3. Upgrading the DGX Software Stack and Red Hat Enterprise Linux 8

This section provides information about how to update your DGX system while remaining on the same GPU driver branch and how to update your DGX system while switching to a different GPU driver branch.



**Important:** Here is some important information you need to know **before** upgrading:

- ▶ An in-place upgrade from RHEL 7 to RHEL 8 with the DGX software stack installed is **not supported**.
- ▶ **Before** you install or perform the upgrade, refer to the section in this release notes for the latest RHEL version.
- ▶ Upgrading to a different driver package can result in the server failing to boot. Follow the instructions to first uninstall the current driver.

Ensure that you are prepared to restore the `GRUB_CMDLINE_LINUX` setting as directed in the instructions in this section.

### 1.3.1. Upgrading the Software without Moving to a New Driver Branch

To update your DGX system with the latest RHEL-8 updates, run the following command:

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

### 1.3.2. Upgrading and Moving to a New Driver Branch on non-NVSwitch Systems

This procedure applies to the DGX-1, DGX Station, and DGX Station A100 systems.



**Important:** Before you install or perform the upgrade, refer to the EL8-22.08 section for information about the latest release.

1. Preserve the `GRUB_CMDLINE_LINUX` setting.

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

**Example:**

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

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

2. Issue the following to remove the current driver package and install the new driver package.

```
sudo dnf remove -y nv-persistence-mode libnvidia-nscq-<current driver version>
sudo dnf module remove --all -y nvidia-driver
sudo dnf module reset -y nvidia-driver
sudo dnf module install -y nvidia-driver:<new driver version>/{(default,src)}
sudo dnf install -y nv-persistence-mode libnvidia-nscq-<new driver version>
sudo dnf update -y --nobest
```

3. Restore the GRUB\_CMDLINE\_LINUX setting.

In the `/etc/default/grub` file, remove extra instances of `GRUB_CMDLINE_LINUX` and manually edit the file to restore the original setting (except for the blacklist parameter).

**Example:**

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

4. Reboot the system.

```
sudo reboot
```

### 1.3.3. Updating and Moving to a New Driver Branch on NVSwitch Systems

This procedure applies to the DGX-2 or DGX A100 systems.



**Important:** Before you install or perform the upgrade, refer to the EL8-22.08 section for information about the latest release.

1. Preserve the GRUB\_CMDLINE\_LINUX setting.

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

**Example:**

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

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

2. Issue the following to remove the current driver package and install the new driver package.

```
sudo dnf remove -y nv-persistence-mode nvidia-fm-enable
sudo dnf module remove --all -y nvidia-driver
sudo dnf module reset -y nvidia-driver
sudo dnf module install -y nvidia-driver:<new driver version>/{fm,src}
sudo dnf install -y nv-persistence-mode nvidia-fm-enable
sudo dnf update -y --nobest
```

3. Restore the GRUB\_CMDLINE\_LINUX setting.

In the `/etc/default/grub` file, remove any extra instances of `GRUB_CMDLINE_LINUX` and manually edit the file to restore the original setting (except for the blacklist parameter).

**Example:**

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

4. Reboot the system.

```
sudo reboot
```

---

# Chapter 2. Version EL8-22.08



**ATTENTION:** If your system is running a version **earlier** than EL8-22.05, you need to update the keys on the system. Refer to [Rotating the GPG Key](#) for more information about how to rotate the keys

The DGX Software for Red Hat Enterprise Linux 8, EL8-22.08, is available.

EL8-22.08 supports all DGX products - DGX A100, DGX-2, DGX-1, DGX Station, and DGX Station A100.



**Important:** Installing or updating to EL8-22.08 also updates the installed Red Hat Enterprise Linux 8 distribution to the latest version.

- ▶ NVIDIA GPUDirect Storage (GDS) v1.1 does not support Red Hat Enterprise Linux 8.5.  
If you are using GDS 1.1, contact NVIDIA Enterprise Support before performing the upgrade.
- ▶ If you need to use the Mellanox OpenFabrics Enterprise Distribution for Linux (`MLNX_OFED`), before you install or update to EL8-22.08, ensure that there is a `MLNX_OFED` package version available that supports the latest Red Hat Enterprise Linux 8 version.  
  
Refer to the [DGX Software for Red Hat Enterprise Linux 8 Installation Guide](#) for instructions.

## Change Highlights

- ▶ Updated R450, R470 GPU drivers (see Software Contents below for versions)
  - ▶ **Attention:** R515 GPU driver is currently not supported.
- ▶ Updated NVSM to 22.06.02
- ▶ Updated DCGM to 2.4.7
- ▶ Updated MLNX OFED to 5.4-3.5.8.0
- ▶ Updated NCCL to 2.15.1
- ▶ Updated cuDNN to 8.4.1
- ▶ Updated docker-ce: 20.10.18

## Software Contents

The following table provides version information for software included in the DGX Software Stack for Red Hat Enterprise Linux 8.



**Note:** Unlike the DGX OS shipped with the NVIDIA DGX system, the DGX software stack for Red Hat does not include the Mellanox OpenFabrics Enterprise Distribution (MLNX\_OFED) for Linux. When using MLNX\_OFED with Red Hat, ensure you install a supported MLNX\_OFED kernel version to avoid incompatibilities with the Red Hat distribution kernel. Refer to the [DGX Software for Red Hat Enterprise Linux 8 Installation Guide](#) for instructions.

Table 1. Contents of the Repositories

Component	Version	Additional Information
OS	RHEL 8.6	
Kernel	4.18.0-372.13.1 or later	
CUDA Toolkit	11.4	Refer to the <a href="#">NVIDIA CUDA Toolkit Release Notes</a> .  <div data-bbox="1039 1024 1421 1438"> <p><b>Note:</b> CUDA 11.4 has been qualified with Red Hat Enterprise Linux 8.4 and older. For newer releases, please refer to <a href="#">Installing Required Components</a> for installation instructions of the driver.</p> </div>
GPU Driver	<b>R450: 450.203.08</b> <b>R470: 470.141.10</b> R510: 510.85.02	Refer to the <a href="#">NVIDIA Data Center GPU documentation</a>  <b>Attention:</b> R515 GPU driver is currently not supported.
NCCL	<a href="#">2.15.1</a>	
cuDNN	<a href="#">8.4.1</a>	
DCGM	<a href="#">2.4.7</a>	Refer to the <a href="#">DCGM Release Notes</a> .



Component	Version	Additional Information
Mellanox OFED	<b>MLNX 5.4-3.5.8.0</b>	Refer to <a href="#">MLNX 5.4-3.5.8.0</a>
MLNX FW	ConnectX-4 12.28.2006 ConnectX-5 16.31.2006 <b>ConnectX-6 20.31.2354</b> <b>ConnectX-7 28.34.4000</b>	
NVSM	22.06.02	Refer to the <a href="#">NVIDIA System Management Documentation</a> .
Docker Engine	docker-ce: 20.10.18  <b>Note:</b> If necessary, the following components require separate installation via <code>sudo apt install</code> :  <ul style="list-style-type: none"> <li>▶ docker-ce-rootless-extras 20.10.17</li> <li>▶ docker-scan-plugin 0.9.0</li> </ul>	Refer to <a href="#">v20.10.17</a>
NVIDIA Container Toolkit	nvidia-container-toolkit: 1.10.0-1 nvidia-docker2: 2.11.0-1 libnvidia-container1: 1.10.0-1 libnvidia-container-tools: 1.10.0-1	Refer to the <a href="#">NVIDIA Container Toolkit</a> documentation.
NGC CLI	2.2.0-1	Refer to the <a href="#">NGC CLI Documentation</a>
GPUDirect Storage (GDS)	v1.0	Refer to <a href="#">GDS Documentation</a>
MIG Configuration Tool	nvidia-mig-manager 0.4.3	Refer to the following NVIDIA mig-parted github pages: <a href="https://github.com/NVIDIA/mig-parted">https://github.com/NVIDIA/mig-parted</a> and <a href="https://github.com/NVIDIA/mig-parted/tree/master/deployments/systemd">https://github.com/NVIDIA/mig-parted/tree/master/deployments/systemd</a>
nvipmitool	1.0.6.0	
nvidia-peer-memory/nvidia-peer-memory DKMS	1.3.0	

## Compatibility

NVIDIA has validated and tested **DGX Software version EL8-22.08** on the following systems:

- ▶ Linux Distribution and kernel:
  - ▶ Red Hat Enterprise Linux 8.6
  - ▶ Rocky Linux 8
  - ▶ Kernel **4.18.0-372.13.1**
- ▶ NVIDIA DGX systems
  - ▶ NVIDIA DGX A100 with Red Hat Enterprise Linux 8.6 and Rocky Linux 8
  - ▶ NVIDIA DGX-2 with Red Hat Enterprise Linux 8.6 and Rocky Linux 8
  - ▶ NVIDIA DGX-1 (V100) with Red Hat Enterprise Linux 8.6 and Rocky Linux 8
  - ▶ NVIDIA DGX Station with Red Hat Enterprise Linux 8.6 and Rocky Linux 8
  - ▶ NVIDIA DGX Station A100 with Red Hat Enterprise Linux 8.6 and Rocky Linux 8
- ▶ 22.08 Deep Learning Framework containers
- ▶ NVIDIA GPUDirect Storage v1.0 (refer to the [GDS documentation](#) for additional information)
- ▶ MLNX OFED version **5.4-3.5.8.0**
- ▶ ConnectX Firmware: see table 1 above

## Update Instructions

See the section [Installing and Updating the Software](#) for instructions.

---

# Chapter 3. Version EL8-22.06



**ATTENTION:** If your system is running a version **earlier** than EL8-22.05, you need to update the keys on the system. Refer to [Rotating the GPG Key](#) for more information about how to rotate the keys

The DGX Software for Red Hat Enterprise Linux 8, EL8-22.06, is available.

EL8-22.06 supports all DGX products - DGX A100, DGX-2, DGX-1, DGX A100, DGX Station, and DGX Station A100.



**Important:** Installing or updating to EL8-22.06 also updates the installed Red Hat Enterprise Linux 8 distribution to the latest version.

- ▶ NVIDIA GPUDirect Storage (GDS) v1.1 does not support Red Hat Enterprise Linux 8.5.  
If you are using GDS 1.1, contact NVIDIA Enterprise Support before performing the upgrade.
- ▶ If you need to use the Mellanox OpenFabrics Enterprise Distribution for Linux (`MLNX_OFED`), before you install or update to EL8-22.06, ensure that there is a `MLNX_OFED` package version available that supports the latest Red Hat Enterprise Linux 8 version.

Refer to the [DGX Software for Red Hat Enterprise Linux 8 Installation Guide](#) for instructions.

## Change Highlights

- ▶ Updated R450 and R470 GPU drivers (see Software Contents below for versions)
- ▶ Updated NVSM to 22.03.05
- ▶ Updated DCGM to 2.3.6
- ▶ Updated NCCL to 2.12.12
- ▶ Updated cuDNN to 8.3.3
- ▶ Updated docker-ce: 20.10.16

## Software Contents

The following table provides version information for software included in the DGX Software Stack for Red Hat Enterprise Linux 8.



**Note:** Unlike the DGX OS shipped with the NVIDIA DGX system, the DGX software stack for Red Hat does not include the Mellanox OpenFabrics Enterprise Distribution (`MLNX_OFED`) for Linux. When using `MLNX_OFED` with Red Hat, ensure you install a supported `MLNX_OFED` kernel version to avoid incompatibilities with the Red Hat distribution kernel. Refer to the [DGX Software for Red Hat Enterprise Linux 8 Installation Guide](#) for instructions.

Table 2. Contents of the Repositories

Component	Versions
GPU Driver (R470)	<a href="#">470.129.06</a>
GPU Driver (R450)	<a href="#">450.191.01</a>
NCCL	<a href="#">2.12.12</a>
cuDNN	<a href="#">8.3.3</a>
NVIDIA System Management (NVSM)	<a href="#">22.03.05</a>
Data Center GPU Management (DCGM)	<a href="#">2.3.6</a>
DGX Station Theme	nv-yaru-theme: 20.10-1
CUDA Toolkit	CUDA 11.4
Docker Engine	20.10.16
nvidia-container-runtime	3.7.0-1
NGC CLI	2.2.0-1
nvidia-mig-manager	0.1.2-1

## Compatibility

NVIDIA has validated and tested **DGX Software version EL8-22.06** on the following systems:

- ▶ Linux Distribution and kernel:
  - ▶ Red Hat Enterprise Linux 8.6
  - ▶ Rocky Linux 8.6
  - ▶ Kernel 4.18.0-372.9.1
- ▶ NVIDIA DGX systems
  - ▶ NVIDIA DGX A100 with Red Hat Enterprise Linux 8.6 and Rocky Linux 8

- ▶ NVIDIA DGX-2 with Red Hat Enterprise Linux 8.6 and Rocky Linux 8
- ▶ NVIDIA DGX-1 (V100) with Red Hat Enterprise Linux 8.6 and Rocky Linux 8
- ▶ NVIDIA DGX Station with Red Hat Enterprise Linux 8.6 and Rocky Linux 8
- ▶ NVIDIA DGX Station A100 with Red Hat Enterprise Linux 8.6 and Rocky Linux 8
- ▶ 22.04 Deep Learning Framework containers
- ▶ NVIDIA GPUDirect Storage v1.0 (refer to the [GDS documentation](#) for additional information)
- ▶ MLNX OFED version 5.4-3.4.0.0
- ▶ ConnectX Firmware
  - ▶ ConnectX-4: 12.28.2006
  - ▶ ConnectX-5: 16.31.2006
  - ▶ ConnectX-6: 20.31.2006

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.

### Update Instructions

See the section [Installing and Updating the Software](#) for instructions.

---

# Chapter 4. Version EL8-22.04



**Important:** The features and component versions in EL8-22.04 are identical to the versions in EL8-22.02. In EL8-22.04, the GPG keys that are used to sign the packages and metadata in those repositories need to be rotated. Refer to [Rotating the GPG Key](#) for more information about how to rotate the keys.

EL8-22.04 supports all DGX products - DGX A100, DGX-2, DGX-1, DGX Station, and DGX Station A100.



**Important:** Installing or updating to EL8-22.04 also updates the installed Red Hat Enterprise Linux 8 distribution to the latest version.

- ▶ NVIDIA GPUDirect Storage (GDS) v1.1 does not support Red Hat Enterprise Linux 8.5. If you are using GDS 1.1, contact NVIDIA Enterprise Support before performing the upgrade.
- ▶ If you require use of the Mellanox OpenFabrics Enterprise Distribution for Linux (MLNX\_OFED), then before installing or updating to EL8-22.04, be sure that there is a MLNX\_OFED package version available that supports the latest Red Hat Enterprise Linux 8 version. Refer to the [DGX Software for Red Hat Enterprise Linux 8 Installation Guide](#) for instructions.

## Software Contents

The following table provides version information for software included in the DGX Software Stack for Red Hat Enterprise Linux 8.



**Note:** Unlike the DGX OS shipped with the NVIDIA DGX system, the DGX software stack for Red Hat does not include the Mellanox OpenFabrics Enterprise Distribution (MLNX\_OFED) for Linux. When using MLNX\_OFED with Red Hat, ensure you install a supported MLNX\_OFED kernel version to avoid incompatibilities with the Red Hat distribution kernel. Refer to the [DGX Software for Red Hat Enterprise Linux 8 Installation Guide](#) for instructions.

Table 3. Contents of the Repositories

Component	Versions
GPU Driver (R470)	<a href="#">470.103.01</a>
GPU Driver (R450)	<a href="#">450.172.01</a>
NCCL	<a href="#">2.11.4</a>
cuDNN	<a href="#">8.3.2</a>
NVIDIA System Management (NVSM)	<a href="#">21.09.14</a>
Data Center GPU Management (DCGM)	<a href="#">2.3.2</a>
DGX Station Theme	nv-yaru-theme: 20.10-1
CUDA Toolkit	CUDA 11.4
Docker Engine	20.10.11
nvidia-container-runtime	3.7.0-1
NGC CLI	2.2.0
nvidia-mig-manager	0.1.2-1

## Compatibility

NVIDIA has validated and tested **DGX Software version EL8-22.04** on the following systems:

- ▶ Linux Distribution and kernel:
  - ▶ Red Hat Enterprise Linux 8.5
  - ▶ Kernel 4.18.0-xxx
- ▶ NVIDIA DGX systems
  - ▶ NVIDIA DGX A100 with Red Hat Enterprise Linux 8.5
  - ▶ NVIDIA DGX-2 with Red Hat Enterprise Linux 8.5
  - ▶ NVIDIA DGX-1 (V100) with Red Hat Enterprise Linux 8.5
  - ▶ NVIDIA DGX Station with Red Hat Enterprise Linux 8.5
  - ▶ NVIDIA DGX Station A100 with Red Hat Enterprise Linux 8.5
- ▶ 21.10 Deep Learning Framework containers
- ▶ NVIDIA GPUDirect Storage v1.0 (refer to the [GDS documentation](#) for additional information)
- ▶ MLNX OFED version 5.4-3.1.0.0
- ▶ ConnectX Firmware
  - ▶ ConnectX-4: 12.28.2006
  - ▶ ConnectX-5: 16.31.1014
  - ▶ ConnectX-6: 20.31.1014

## Update Instructions

See the section [Installing and Updating the Software](#) for instructions.



---

# Chapter 5. Version EL8-22.02

The DGX Software for Red Hat Enterprise Linux 8, EL8-22.02, is available.

EL8-22.02 supports all DGX products - DGX A100, DGX-2, DGX-1, DGX Station, and DGX Station A100.



**Important:** Installing or updating to EL8-22.02 also updates the installed Red Hat Enterprise Linux 8 distribution to the latest version.

- ▶ NVIDIA GPUDirect Storage (GDS) v1.1 does not support Red Hat Enterprise Linux 8.5. If you are using GDS 1.1, contact NVIDIA Enterprise Support before performing the upgrade.
- ▶ If you require use of the Mellanox OpenFabrics Enterprise Distribution for Linux (MLNX\_OFED), then before installing or updating to EL8-22.02, be sure that there is a MLNX\_OFED package version available that supports the latest Red Hat Enterprise Linux 8 version. Refer to the [DGX Software for Red Hat Enterprise Linux 8 Installation Guide](#) for instructions.

## Change Highlights

- ▶ Updated R450 and R470 GPU drivers (see Software Contents below for versions)
- ▶ Updated NVSM to 21.09.14.
- ▶ Updated DCGM to 2.3.2

## Software Contents

The following table provides version information for software included in the DGX Software Stack for Red Hat Enterprise Linux 8.



**Note:** Unlike the DGX OS shipped with the NVIDIA DGX system, the DGX software stack for Red Hat does not include the Mellanox OpenFabrics Enterprise Distribution (MLNX\_OFED) for Linux. When using MLNX\_OFED with Red Hat, ensure you install a supported MLNX\_OFED kernel version to avoid incompatibilities with the Red Hat distribution kernel. Refer to the [DGX Software for Red Hat Enterprise Linux 8 Installation Guide](#) for instructions.

Table 4. Contents of the Repositories

Component	Versions
GPU Driver (R470)	<a href="#">470.103.01</a>
GPU Driver (R450)	<a href="#">450.172.01</a>
NCCL	<a href="#">2.11.4</a>
cuDNN	<a href="#">8.3.2</a>
NVIDIA System Management (NVSM)	<a href="#">21.09.14</a>
Data Center GPU Management (DCGM)	<a href="#">2.3.2</a>
DGX Station Theme	nv-yaru-theme: 20.10-1
CUDA Toolkit	CUDA 11.4
Docker Engine	20.10.11
nvidia-container-runtime	3.7.0-1
NGC CLI	2.2.0
nvidia-mig-manager	0.1.2-1

## Compatibility

NVIDIA has validated and tested **DGX Software version EL8-22.02** on the following systems:

- ▶ Linux Distribution and kernel:
  - ▶ Red Hat Enterprise Linux 8.6
  - ▶ Rocky Linux 8
  - ▶ Kernel 4.18.0-xxx
- ▶ NVIDIA DGX systems
  - ▶ NVIDIA DGX A100 with Red Hat Enterprise Linux 8.5 and CentOS 8
  - ▶ NVIDIA DGX-2 with Red Hat Enterprise Linux 8.5 and CentOS 8
  - ▶ NVIDIA DGX-1 (V100) with Red Hat Enterprise Linux 8.5 and CentOS 8
  - ▶ NVIDIA DGX Station with Red Hat Enterprise Linux 8.5 and CentOS 8
  - ▶ NVIDIA DGX Station A100 with Red Hat Enterprise Linux 8.5 and CentOS 8
- ▶ 21.10 Deep Learning Framework containers
- ▶ NVIDIA GPUDirect Storage v1.0 (refer to the [GDS documentation](#) for additional information)
- ▶ MLNX OFED version 5.4-3.1.0.0
- ▶ ConnectX Firmware
  - ▶ ConnectX-4: 12.28.2006
  - ▶ ConnectX-5: 16.31.1014
  - ▶ ConnectX-6: 20.31.1014

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.

### Update Instructions

See the section [Installing and Updating the Software](#) for instructions.

---

# Chapter 6. Version EL8-21.08

The DGX Software for Red Hat Enterprise Linux 8, EL8-21.08, is available.

EL8-21.08 supports all DGX products - DGX A100, DGX-2, DGX-1, DGX Station, and DGX Station A100.



**Important:** Installing or updating to EL8-21.08 also updates the installed Red Hat Enterprise Linux 8 distribution to the latest version.

- ▶ NVIDIA GPUDirect Storage (GDS) v1.1 does not support Red Hat Enterprise Linux 8.6. If you are using GDS 1.1, contact NVIDIA Enterprise Support before performing the upgrade.
- ▶ If you require use of the Mellanox OpenFabrics Enterprise Distribution for Linux (MLNX\_OFED), then before installing or updating to EL8-21.08, be sure that there is a MLNX\_OFED package version available that supports the latest Red Hat Enterprise Linux 8 version. Refer to the [DGX Software for Red Hat Enterprise Linux 8 Installation Guide](#) for instructions.

## Change Highlights

- ▶ Support for DGX Station A100
- ▶ Added the NVIDIA GPU Driver Release 470
- ▶ Validated with NVIDIA `MLNX_OFED` 5.4
- ▶ Docker CE updated to 20.10
- ▶ Additional updates

Upgrading or updating a system replaces the installed packages with the latest versions available from NVIDIA at the time the system is updated. These packages include security updates and corrections for other high-impact bugs, with focus on maintaining stability and compatibility with earlier versions. The following table lists the updates that have been made to the NVIDIA repository since the initial release of EL8-21.08.

Table 5.

Update Date	Package	Version
December 14, 2021	R470 GPU Driver	470.82.01
	NVIDIA System Management (NVSM)	21.09.10
	Docker Engine CE	docker-ce-* 20.10.11
		<b>Note:</b> If needed, the following components require separate installation via <code>sudo dnf install</code> .
		docker-ce-rootless-extras 20.10.11
		docker-scan-plugin 0.9.0
	NVIDIA Container Stack	nvidia-docker2-2.8.0-1
		nvidia-container-runtime-3.7.0-1
		nvidia-container-toolkit-1.7.0-1
		libnvidia-container-tools-1.7.0-1
	libnvidia-container1-1.7.0-1	

## Software Contents

The following table provides version information for software included in the DGX Software Stack for Red Hat Enterprise Linux 8.



**Note:** Unlike the DGX OS shipped with the NVIDIA DGX system, the DGX software stack for Red Hat does not include the Mellanox OpenFabrics Enterprise Distribution (`MLNX_OFED`) for Linux. When using `MLNX_OFED` with Red Hat, ensure you install a supported `MLNX_OFED` kernel version to avoid incompatibilities with the Red Hat distribution kernel. Refer to the [DGX Software for Red Hat Enterprise Linux 8 Installation Guide](#) for instructions.

Table 6. Contents of the Repositories

Component	Versions
GPU Driver (R470)	<a href="#">470.57.02</a>

Component	Versions
GPU Driver (R450)	<a href="#">450.142.00</a>
NCCL	<a href="#">2.10.3</a>
cuDNN	<a href="#">8.2.2</a>
NVIDIA System Management (NVSM)	21.07.14
Data Center GPU Management (DCGM)	2.2.8
DGX Station Theme	nv-yaru-theme: 20.10-1
CUDA Toolkit	CUDA 11.4
Docker Engine	20.10.07
nvidia-container-runtime	3.5.0-1
NGC CLI	2.2.0
nvidia-mig-manager	0.1.2-1

## Compatibility

NVIDIA has validated and tested the **DGX Software version EL8-21.08**, December 14th update, on the following systems:

- ▶ Linux Distribution and kernel:
  - ▶ Red Hat Enterprise Linux 8.6
  - ▶ Rocky Linux 8.6
  - ▶ Kernel 4.18.0-305
- ▶ NVIDIA DGX systems
  - ▶ NVIDIA DGX A100 with Red Hat Enterprise Linux 8.6 and Rocky Linux 8
  - ▶ NVIDIA DGX-2 with Red Hat Enterprise Linux 8.6 and Rocky Linux 8
  - ▶ NVIDIA DGX-1 (V100) with Red Hat Enterprise Linux 8.6 and Rocky Linux 8
  - ▶ NVIDIA DGX Station with Red Hat Enterprise Linux 8.6 and Rocky Linux 8
  - ▶ NVIDIA DGX Station A100 with Red Hat Enterprise Linux 8.6 and Rocky Linux 8
- ▶ 21.10 Deep Learning Framework containers
- ▶ NVIDIA GPUDirect Storage v1.0 (refer to the [GDS documentation](#) for additional information)
- ▶ MLNX OFED version 5.4-3.1.0.0
- ▶ ConnectX Firmware
  - ▶ ConnectX-4: 12.28.2006
  - ▶ ConnectX-5: 16.31.1014
  - ▶ ConnectX-6: 20.31.1014

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.

## Update Instructions

See the section [Installing and Updating the Software](#) for instructions.

## Fixed Issues

- ▶ [\[DGX-2\]: No Rebuild Function for RAID 0 if Volume is not md1](#)
- ▶ [\[DGX-1, DGX-2\]: nvsm dump health Does not Generate sosreport](#)
- ▶ [\[DGX-2\]: Storage Alerts Persist from Previous RAID Configuration](#)

---

# Chapter 7. Version EL8-20.11

The DGX Software for Red Hat Enterprise Linux 8, EL8-20.11, is available.

EL8-20.11 supports the following DGX products - DGX A100, DGX-2, DGX-1, and DGX Station.



**Important:** Installing or updating to EL8-20.11 also updates the installed Red Hat Enterprise Linux 8 distribution to the latest version. If you require use of the Mellanox OpenFabrics Enterprise Distribution for Linux (MLNX\_OFED), then before installing or updating to EL8-20.11, be sure that there is a MLNX\_OFED package version available that supports the latest Red Hat Enterprise Linux 8 version.

- ▶ To check the latest Red Hat Enterprise Linux 8 version, visit <https://access.redhat.com/articles/3078>
- ▶ To check the MLNX\_OFED package OS support, visit <https://docs.mellanox.com/category/mlnxofedib>, click the latest MLNX\_OFED software version and then use the side menu to navigate to Release Notes->General Support in MLNX\_OFED and view Supported Operating Systems.

See also the [Compatibility](#) section for more information.

## Change Highlights

- ▶ Initial release of the DGX for Red Hat Enterprise Linux 8 software stack
- ▶ Additional updates

Upgrading or updating a system replaces the installed packages with the latest versions available from NVIDIA at the time the system is updated. These packages include security updates and corrections for other high-impact bugs, with focus on maintaining stability and compatibility with earlier versions. The following table lists the updates that have been made to the NVIDIA repository since release of EL8-20.11.



Table 7.

Update Date	Package	Version	Description
June 9, 2021	NVIDIA System Management (NVSM)	20.09.26	Added check for homogeneity of PSUs for DGX A100.

## Software Contents

The following table provides version information for software included in the DGX Software Stack for Red Hat Enterprise Linux 8.



**Note:** Unlike the DGX OS shipped with the NVIDIA DGX system, the DGX software stack for Red Hat does not include the Mellanox OpenFabrics Enterprise Distribution (MLNX\_OFED) for Linux. When using MLNX\_OFED with Red Hat, ensure you install a supported MLNX\_OFED kernel version to avoid incompatibilities with the Red Hat distribution kernel. Refer to the [DGX Software for Red Hat Enterprise Linux 8 Installation Guide](#) for instructions.

Table 8. Contents of the Repositories

Component	Versions in the Release 450 Driver Package
GPU Driver	<a href="#">450.80.02</a>
NVIDIA System Management (NVSM)	20.09.26
Data Center GPU Management (DCGM)	2.0.13
DGX Station Theme	dgxstation-desktop - 19.10-0 dgx-gnome - 19.10-0
CUDA Toolkit	CUDA 11.2
Docker Engine	19.03.13
nvidia-container-runtime	3.4.0-1

## Compatibility

NVIDIA has validated and tested the **DGX Software version EL8-20.11** with the following:

- ▶ Linux Distribution
  - ▶ Red Hat Enterprise Linux 8.3
  - ▶ CentOS 8.3
- ▶ NVIDIA DGX systems

- ▶ NVIDIA DGX A100
- ▶ NVIDIA DGX-2
- ▶ NVIDIA DGX-1 (V100)
- ▶ NVIDIA DGX Station



**Important:** Currently, the Mellanox cards are not supported with Red Hat Enterprise Linux/CentOS 8.4. To stay on Release 8.3:

- ▶ During the initial installation, select a Red Hat Enterprise Linux ISO image for version 8.3 to install.
- ▶ When performing an update, as part of the initial installation or after an installation, issue the following command to pin the release to 8.3 and then perform the update:

```
sudo subscription-manager release --set=8.3
sudo dnf update -y --nobest
```

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.

## Update Instructions

See the section [Installing and Updating the Software](#) for instructions.

## Known Issues

- ▶ [\[DGX-1, DGX-2\]: nvsm dump health Does not Generate sosreport](#)
- ▶ [\[DGX-2\]: No Rebuild Function for RAID 0 if Volume is not md1](#)
- ▶ [\[DGX-2\]: Storage Alerts Persist from Previous RAID Configuration](#)

---

# Chapter 8. Known Issues

See the sections for specific versions to see which issues are open in those versions.

## 8.1. [DGX-1, DGX-2]: DCGM Update Error

### Issue

When attempting to update DCGM, the following error message may be displayed: There was an internal error during the test: 'Couldn't find the ubergemm executable which is required; the install may have failed.'

### Explanation

This is a known issue and is anticipated to be fixed in the next release.

## 8.2. [DGX-1, DGX-2]: GPU MIG Partitions do not return output fields

### Issue

When enabling MIG and creating a MIG partition for the GPU, there is no output returned for non-device specific fields: `dcgmi dmon -e 1,2,3,4,5`

### Explanation

This issue affects EL 8 with:

- ▶ Driver Version: 470.141.03
- ▶ CUDA Version: 11.4.152
- ▶ DCGM: 2.4.5

## 8.3. [DGX-1, DGX-2]: Log displays CEC error

### Issue

DGX A100 Firmware Update Container log may show error messages such as "Unable to send RAW command (channel=0x0 netfn=0x3c lun=0x0 cmd=0xf rsp=0xd3): Destination unavailable"

This error will be displayed when running supported commands and may be safely ignored.

## 8.4. [DGX-1]: NVSM show controllers SerialNumber shows `NOT\_SET`

### Issue

After rebooting, `nvsm show controllers` may display a blank serial number.

### Explanation

This issue is specific to the DGX-1 platform with the MegaRAID controller and can be remedied by restarting the nvsm service after 30 minutes. To restart the service, run `systemctl restart nvsm`

## 8.5. [DGX-1, DGX-2]: nsys fails to launch

### Issue

When attempting to install a different package version from the CUDA networking repository, `nsys` will not launch.

More specifically, when installing CUDA Toolkit there are some `nsight-systems` packages with different versions and the most recent `nsight-systems-2022.1.3-2022.1.3.3\_1c7b5f7-0.x86\_64.rpm` will be installed by default.

### Workaround

Download and install this package from CUDA repo, which resolves the path to fix the issue:

```
$ wget https://developer.download.nvidia.com/compute/cuda/repos/rhel8/x86_64/nsight-systems-2021.3.2-2021.3.2.4_027534f-0.x86_64.rpm
$ sudo dnf install ./nsight-systems-2021.3.2-2021.3.2.4_027534f-0.x86_64.rpm
```

## 8.6. [DGX-1, DGX-2]: nvsm dump health Does not Generate sosreport

### Issue

[Fixed in EL8-21.08]

After running `nvsm dump health`, the log file reports

```
INFO: Could not find sosreport output file
```

Analysis of the log files reveals that information is missing for components that are installed on the system; such as InfiniBand cards.

### Explanation

The sosreport is not getting collected. This will be resolved in a later software release.

## 8.7. [DGX-2]: No Rebuild Function for RAID 0 if Volume is not md1

### Issue

[Fixed in EL8-21.08]

If the RAID 0 volumes are a designation other than `md1`, such as `md128`, then there is no `Rebuild` under `Target`.

Example:

```
$ nvsm show /systems/localhost/storage/volumes/md128
Properties:
  CapacityBytes = 61462699992678
  Encrypted = False
  Id = md128
  Drives = [ nvme10n1, nvme11n1, nvme12n1, nvme13n1, nvme14n1, nvme15n1, nvme16n1,
nvme17n1, nvme2n1, nvme3n1, nvme4n1, nvme5n1, nvme6n1, nvme7n1, nvme8n1, nvme9n1 ]
  Name = md128
  Status_Health = OK
  Status_State = Enabled
  VolumeType = RAID-0
Targets:
  encryption
Verbs:
  cd
  show
```

## Explanation

This occurs if the drives are not configured in the default configuration of two OS drives in a RAID 1 configuration and storage drives used for caching in a RAID 0 configuration. The `nvsm rebuild` command does not support non-standard drive configurations.

## 8.8. [DGX-2]: Storage Alerts Persist from Previous RAID Configuration

### Issue

[fixed with EL8-21.08] After switching to a custom drive configuration, such as by adding or removing storage drives, any NVSM storage alerts from the previous configuration will still be reported even though the current drive status is healthy.

For example, alerts can appear for `md0` even though the current RAID drive name is `md125` and healthy.

### Explanation

To configure NVSM to support custom drive partitioning, perform the following.

1. Edit `/etc/nvsm/nvsm.config` and set the `"use_standard_config_storage"` parameter to false.

```
"use_standard_config_storage":false
```

2. Remove the NVSM database.

```
$ sudo rm /var/lib/nvsm/sqlite/nvsm.db
```

3. Restart NVSM.

```
$ systemctl restart nvsm
```

---

# Chapter 9. Rotating the GPG Key

NVIDIA constantly evaluates and improves security implementations. As part of these improvements, we are rolling out changes to harden the security and reliability of our repositories. These changes require rotating the GPG keys that are used to sign the packages and metadata in those repositories. This section provides information about how to rotate the GPG keys on your system.

1. Directly install the `nvidia-repo-setup` package:

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

2. Manually revoke the previous DGX and CUDA GPG keys.

```
$ sudo rpm -e gpg-pubkey-629c85f2-57571711  
$ sudo rpm -e gpg-pubkey-7fa2af80-576db785
```

3. Clean up old repository metadata.

```
$ sudo dnf clean metadata
```

OTA updates can now occur as normal.

**Copyright**

© 2022 NVIDIA Corporation & Affiliates. All rights reserved.

