




# NVIDIA DGX OS SERVER VERSION 3.1.7

DA-08260-317\_v03.1 | June 2018

**Release Notes and Update Guide**



# TABLE OF CONTENTS

<b>NVIDIA DGX OS Server Release Notes for Version 3.1.7</b> .....	<b>3</b>
Update Advisement .....	4
About Release 3.1 .....	4
Highlights and Changes in Version 3.1.7 .....	5
Known Issues .....	5
Ubuntu / Linux Kernel Issues .....	7
DGX OS Server Software Content .....	9
Version Reference .....	10
Pascal .....	10
Volta (16 GB) .....	10
Volta (32 GB) .....	10
Re-Imaging Your DGX-1 .....	11
<b>Updating to Version 3.1.7</b> .....	<b>12</b>
Update Path Instructions .....	12
Connecting to the DGX-1 Console .....	13
Verifying the DGX-1 Connection to the Repositories .....	14
Updating from 1.x to 2.x .....	14
Update Instructions .....	14
Recovering from an Interrupted Update .....	17
Updating from 2.x to 3.1.7 .....	18
Update Instructions .....	18
Verifying the nvidia-peer-memory Module .....	21
Enabling Dynamic DNS Updates .....	22
Recovering from an Interrupted Update .....	23
Updating from 3.1.x to 3.1.7 .....	24
Update Instructions .....	24
Verifying the nvidia-peer-memory Module .....	25
Recovering from an Interrupted or Failed Update .....	25
<b>Appendix A. Third Party License Notice</b> .....	<b>26</b>

# NVIDIA DGX OS SERVER RELEASE NOTES FOR VERSION 3.1.7

This document describes version 3.1.7 of the NVIDIA® DGX™ OS Server Release 3.1 software and update package.

DGX OS Server v3.1.7 is provided as an “over the network” update, and requires an internet connection and ability to access the NVIDIA public repository. See the chapter [Updating to Version 3.1.7](#) for instructions on performing the update.

DGX OS Server v3.1.7 is also provided by Enterprise Support as an ISO image in the event a system needs to be re-imaged.

If your DGX-1 is not connected to a network with internet access, refer to the DGX-1 User Guide for instructions on how to install the software on air-gapped systems.

## UPDATE ADVISEMENT

### ▶ DGX OS Server software

NVIDIA recommends updating the DGX OS Server software on their DGX-1 systems from version 1.x or 2.x to version 3.1.7. See the Highlights section for details of version 3.1.7.

### ▶ NVIDIA Docker Containers

In conjunction with DGX OS Server v3.1.7, customers should update their NVIDIA Docker containers to the latest container release<sup>1</sup>.

### ▶ Ubuntu Security Updates

Customers are responsible for keeping the DGX-1 up to date with the latest Ubuntu security updates using the 'apt-get upgrade' procedure. See the [Ubuntu Wiki Upgrades](#) web page for more information.

## ABOUT RELEASE 3.1

The following are the primary features of the DGX OS Server Release 3.1:

- ▶ Supports DGX-1 using NVIDIA Pascal as well as Volta GPUs.
- ▶ Ubuntu 16.04 LTS
  - Initialization daemon changed from **Upstart** to **systemd**.
  - Updated network interface naming policy.
 

Policy now uses predictable names, rather than the native naming scheme used in previous releases. The first and second Ethernet interfaces, enumerated as em1 and em2 in previous releases, will now enumerate as enp1s0f0 and enp1s0f1 respectively.
- ▶ NVIDIA GPU Driver Release 384
  - Supports the NVIDIA Tesla™ V100 GPUs.
  - Supports CUDA 9.0
- ▶ CUDA drivers and diagnostic packages updated to Release 384
- ▶ Mellanox drivers updated to 4.0.
- ▶ Docker CE and the Docker Engine Utility for NVIDIA GPUs are pre-installed, and the docker daemon automatically launched.

---

<sup>1</sup> See the *NVIDIA Deep Learning Frameworks* documentation website (<http://docs.nvidia.com/deeplearning/dgx/index.htm>) for information on the latest container releases as well as instructions for how to access them.

## HIGHLIGHTS AND CHANGES IN VERSION 3.1.7

- ▶ NVIDIA GPU Driver Version 384.145  
Includes security updates for driver components.
- ▶ Added support for updating to the NVIDIA Container Runtime for Docker
- ▶ Updated `nvsysinfo` to version 18.04-4
- ▶ Updated `nvhealth` to version 18.04-4

## KNOWN ISSUES

- ▶ [Script Cannot Recreate RAID Array After Re-inserting a Known Good SSD](#)
- ▶ [Software Power Cap Not Reported Correctly by `nvidia-smi`](#)
- ▶ [GPUs Cannot be Reset While the System is Running](#)
- ▶ [Apparmor Profile May not Work with Some Containers](#)

### Script Cannot Recreate RAID Array After Re-inserting a Known Good SSD

#### Issue

When a good SSD is removed from the DGX-1 RAID 0 array and then re-inserted, the script to recreate the array fails.

#### Workaround

After re-inserting the SSD back into the system, the RAID controller sets the array to offline and marks the re-inserted SSD as `Unconfigured_Bad (UBad)`. The script will fail when attempting to rebuild an array when one or more of the SSDs are marked `Ubad`.

To recreate the array in this case,

1. Set the drive back to a good state.

```
# sudo /opt/MegaRAID/storcli/storcli64
/c0/e<enclosure_id>/s<drive_slot> set good
```

2. Run the script to recreate the array.

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

## Software Power Cap Not Reported Correctly by nvidia-smi

### Issue

On DGX-1 systems with Pascal GPUs, nvidia-smi does not report Software Power Cap as "Active" when clocks are throttled by power draw.

### Workaround

This issue is with nvidia-smi reporting and not with the actual functionality. This will be fixed in a future release.

## GPUs Cannot be Reset While the System is Running

### Issue

You will not be able to reset the GPUs while the system is running.

### Workaround

If an issue occurs which causes the GPUs to hang or if they need to be reset, you must reboot the system.

## Apparmor Profile May not Work with Some Containers

### Issue

Apparmor is enabled in this version of the DGX OS Server, with Docker generating a default profile. The default profile may or may not work with your containers.

### Workaround

If there is a conflict with your containers, then either

- ▶ Disable Apparmor, or
- ▶ Provide a custom Apparmor profile and include it in the docker run command.

## UBUNTU / LINUX KERNEL ISSUES

The following are known issues with either the Ubuntu OS or the Linux kernel that affect the DGX-1.

- ▶ [System May Slow Down When Using mpirun](#)
- ▶ [FS-Cache Assertion Error Leading to System Panic May Occur](#)
- ▶ [Network Performance Drop](#)

### System May Slow Down When Using mpirun

#### Issue

Customers running Message Passing Interface (MPI) workloads may experience the OS becoming very slow to respond. When this occurs, a log message similar to the following would appear in the kernel log:

```
kernel BUG at /build/linux-fQ94TU/linux-4.4.0/fs/ext4/inode.c:1899!
```

#### Explanation

Due to the current design of the Linux kernel, the condition may be triggered when `get_user_pages` is used on a file that is on persistent storage. For example, this can happen when `cudaHostRegister` is used on a file path that is stored in an ext4 filesystem. DGX systems implement `/tmp` on a persistent ext4 filesystem.

#### Workaround



**NOTE:** If you performed this workaround on a previous DGX OS software version, you do not need to do it again after updating to the latest DGX OS version.

In order to avoid using persistent storage, MPI can be configured to use shared memory at `/dev/shm` (this is a temporary filesystem).

If you are using Open MPI, then you can solve the issue by configuring the Modular Component Architecture (MCA) parameters so that `mpirun` uses the temporary file system in memory.

For details on how to accomplish this, see the Knowledge Base Article [DGX System Slows Down When Using mpirun](#) (requires login to [NVIDIA Enterprise Support](#)).

## FS-Cache Assertion Error Leading to System Panic May Occur

### Issue

An issue in the Linux kernel can, under some workloads, cause a kernel thread to crash due to an FS-Cache service assertion failure.

This can be confirmed by examining the kernel logs (`/var/log/kern.log*`).

### Example:

```
Mar 27 11:19:42 dev-dgx01 kernel: [514599.193456] FS-Cache:
Mar 27 11:19:42 dev-dgx01 kernel: [514599.193536] FS-Cache: Assertion
failed
Mar 27 11:19:42 dev-dgx01 kernel: [514599.193605] FS-Cache: 6 == 5 is
false
Mar 27 11:19:42 dev-dgx01 kernel: [514599.193767] -----[ cut
here ]-----
Mar 27 11:19:42 dev-dgx01 kernel: [514599.193843] kernel BUG at
/build/linux-3phnTq/linux-4.4.0/fs/fscache/operation.c:494!
```

### Workaround



**NOTE:** If you performed this workaround on a previous DGX OS software version, you do not need to do it again after updating to the latest DGX OS version.

The FS-Cache service can be disabled to prevent system panics as follows.

1. Remove FS-Cache options from NFS or CIFS volumes.

Edit the `/etc/fstab` file to ensure that when mounting an NFS or CIFS volume, the "fsc" option is not used.

For example, remove the highlighted portion:

```
nfs-server.domain:/volume /mnt nfs
ro,noatime,rsize=32768,wsize=32768,nolock,tcp,intr,fsc,nofail,nfsver
s=3 0 0
```

2. Remount the NFS volume without FS-Cache.

```
$ sudo umount /mnt
$ sudo mount /mnt
```

3. Disable the FS-Cache service.

- a) Stop the FS-cache service.

```
$ sudo service cachefilesd stop
```



b) Edit `/etc/default/cachefilesd` and set "RUN=no".

For more background information, see the Knowledge Base Article [DGX-1 System Panic Due to FS-Cache Assertion Failure](#) (requires login to [NVIDIA Enterprise Support](#)).

## Network Performance Drop

### Issue

An issue with the Ubuntu kernel since 4.4.0-116 results in slower network performance when running server-side UDP workloads.

### Details

As of the DGX OS Server 3.1.7 release, the Ubuntu kernel (4.4.0-124) is still subject to this issue. A later kernel version may resolve the issue, at which point an over-the-network update of the DGX OS software will incorporate the fix.

## DGX OS SERVER SOFTWARE CONTENT

The following table provides version information for software included in the DGX OS Server update.

Component	Version	Notes
DGX OS Server	3.1.7	Software versions included in this DGX OS Server update package and ISO image
GPU Driver	384.145	
Docker Engine Utility for NVIDIA GPUs <sup>2</sup>	1.0.1-4	
Ubuntu	16.04 LTS	
Ubuntu kernel	4.4.0-124 or later	Version included in the ISO image For over-the-network updates, the Ubuntu update will incorporate the kernel version available at the time of the update.
Docker CE	18.03-1	Version included in the ISO image and incorporated in the over-the-network update.  As of 5/7/2018 the release of Docker CE installed on your DGX-1 is the release that is provided from a repository maintained by NVIDIA.

<sup>2</sup> NVIDIA Container Runtime for Docker is available as an optional upgrade to replace the Docker Engine Utility for NVIDIA GPUs. For more information, see [Upgrading to the NVIDIA Container Runtime for Docker](#).

## VERSION REFERENCE

The following table shows the firmware and BIOS versions for the DGX-1 hardware at the time of this release. Information provided for reference purposes.

### Pascal

Component	Version	Notes
DGX-1 BMC	3.20.30	Released versions for DGX-1 hardware (Pascal) at the time of this software release. Information provided for reference purposes.
DGX-1 SBIOS	S2W_3A05	
DGX-1 VBIOS	86.00.41.00.05	

### Volta (16 GB)

Component	Version	Notes
DGX-1 BMC	3.20.30	Released versions for DGX-1 hardware (Volta) at the time of this software release. Information provided for reference purposes.
DGX-1 SBIOS	S2W_3A05	
DGX-1 VBIOS	88.00.18.00.01	

### Volta (32 GB)

Component	Version	Notes
DGX-1 BMC	3.20.30	Released versions for DGX-1 hardware (Volta) at the time of this software release. Information provided for reference purposes.
DGX-1 SBIOS	S2W_3A05	
DGX-1 VBIOS	88.00.43.00.04	

## RE-IMAGING YOUR DGX-1

The recommended method for updating your DGX-1 software is by performing the over-the-network update as described in the next chapter.

If necessary, you can re-image the system using the ISO image. Refer to the *DGX-1 User Guide* (<http://docs.nvidia.com/dgx/dgx1-user-guide/index.html>) for instructions on how to re-image the system with the new ISO image.

The re-imaging process creates a fresh installation of the DGX OS. The standard software installation overwrites any data or file systems that may exist on the OS disk as well as the RAID disks. All the disks are re-partitioned, and the RAID array is mounted as `/raid`.



**WARNING:** This process destroys all data and software customizations that you have made on the DGX-1. Be sure to back up any data that you want to preserve, and push any Docker images that you want to keep to a trusted registry.

After re-imaging the system, follow the instructions in the User Guide for first time setup of the DGX-1.

# UPDATING TO VERSION 3.1.7

These instructions explain how to update the DGX-1 software through an internet connection to the NVIDIA public repository. The process updates a DGX-1 system image to the latest QA'd versions of the entire DGX-1 software stack, including the drivers.

Perform the updates using commands on the DGX-1 console.

## UPDATE PATH INSTRUCTIONS

The update instructions depend on the software version currently installed on the DGX-1. Follow the instructions corresponding to your current DGX-1 software.

### Updating from Version 1.x

- a) [Update from Version 1.x to 2.x](#)
- b) [Update from 2.x to 3.1.7](#)

### Updating from Version 2.x

Follow instructions at [Updating from 2.x to 3.1.7](#).

### Updating from Version 3.1.x

Follow instructions at [Updating from 3.1.x to 3.1.7](#).

## CONNECTING TO THE DGX-1 CONSOLE

Connect to the DGX-1 console using either a direct connection or a remote connection through the BMC.

**NOTE:** SSH can be used to perform the update. However, if the Ethernet port is configured for DHCP, there is the potential that the IP address can change after the DGX-1 is rebooted during the update, resulting in loss of connection. If this happens, connect using either a direct connection or through the BMC to continue the update process.

**WARNING:** Connect directly to the DGX-1 Console if the DGX-1 is connected to a 172.17.xx.xx subnet.

DGX OS Server software version 3.1.7 installs Docker CE which uses the 172.17.xx.xx subnet by default for Docker containers. If the DGX-1 system is on the same subnet, you will not be able to establish a network connection to the DGX-1.

Refer to the DGX-1 User Guide for instructions on how to change the default Docker network settings after performing the update.

### Direct Connection

1. Connect a display to the VGA connector and a keyboard to any one of the USB ports.
2. Power on the DGX-1.

### Remote Connection through the BMC

This method requires that you have the BMC login credentials.

1. Make sure you have connected the IPMI port on the DGX-1 to your LAN.
2. Open a Java-enabled browser within your LAN and go to `http://<IPMI IP Address>/`. Use Firefox or Internet Explorer. Google Chrome is not officially supported by the BMC.

Make sure popups are allowed for the BMC address.

3. Log in.
 

If the administrator has not manually created a password, then the username that was created during the initial DGX-1 setup is used for both the BMC username and BMC password.
4. From the top menu, click **Remote Control** and then select **Console Redirection**.
5. Click **Java Console** to open the popup window.
6. If necessary, power on the DGX-1 using the power button icon on the upper right corner of the BMC window.

## VERIFYING THE DGX-1 CONNECTION TO THE REPOSITORIES

Before attempting to perform the update, verify that the DGX-1 network connection can access the public repositories and that the connection is not blocked by a firewall or proxy.

Enter the following on the DGX-1 system.

```
$ wget -O f1-changelogs http://changelogs.ubuntu.com/meta-release-lts
$ wget -O f2-archive
http://archive.ubuntu.com/ubuntu/dists/xenial/Release
$ wget -O f3-usarchive
http://us.archive.ubuntu.com/ubuntu/dists/xenial/Release
$ wget -O f4-security
http://security.ubuntu.com/ubuntu/dists/xenial/Release
$ wget -O f5-download
https://download.docker.com/linux/ubuntu/dists/xenial/Release
$ wget -O f6-international
http://international.download.nvidia.com/dgx/repos/dists/xenial/Release
```

All the `wget` commands should be successful and there should be six files in the directory with non-zero content.

## UPDATING FROM 1.X TO 2.X

These instructions serve as the first phase of the update to DGX OS Server software version 3.1.7 from version 1.x. If you plan to use version 2.x without updating further, then refer to the Release Notes for version 2.x for the proper instructions.

See the section [Connecting to the DGX-1 Console](#) for guidance on connecting to the console to perform the update

### Update Instructions

1. If you have not already done so, verify that your DGX-1 system can access the public repositories as explained in [Verifying the DGX-1 Connection to the Repositories](#).
2. Update the list of available packages and their versions.

```
$ sudo apt-get update
```

3. Install CURL, if not already installed.

```
$ sudo apt-get install curl
```

- Verify that the DGX-1 public repository is configured.

```
$ dpkg -l dgx1-repo-ubuntu1404
```

Expected output:

```
ii dgx1-repo-ubuntu1404 1.1-1 amd64 dgx1 repository configuration files
```

If this output appears, then proceed directly to step 5.

If the `dgx1-repo-ubuntu1404` package is not present, then install the package manually as follows:

- Download the debian repo using `wget`.

```
$ sudo wget
http://international.download.nvidia.com/dgx1/repos/pool/multiverse/
d/dgx1-repo/dgx1-repo-ubuntu1404_1.1-1_amd64.deb -O /tmp/dgx1-repo-
ubuntu1404_1.1-1_amd64.deb
```

- Install the package.

```
$ sudo dpkg -i /tmp/dgx1-repo-ubuntu1404_1.1-1_amd64.deb
```

- Verify that the package installed.

```
$ dpkg -l dgx1-repo-ubuntu1404
```

Expected output:

```
ii dgx1-repo-ubuntu1404 1.1-1 amd64 dgx1 repository configuration files
```

- Get the new package list.

```
$ sudo apt-get update
```

- Install the update script.

```
$ sudo apt-get -y --force-yes install dgx1-ota-update
```

- Run the update script.

The script automatically performs several steps:

- Removes packages (nvidia-361, ar-mgr, isert-dkms, etc).
- Installs the update meta package (dgx1-ota-update-meta).
- Reboots the DGX-1 upon completion.

To run the script, enter the following:

```
$ sudo /usr/bin/dgxl-ota-update.sh
```

Enter y at the prompt:

```
Upgrade NVIDIA DGX1 from 1.0.x to 2.1.4? (y/NO)-:y
```

- You may be presented with a configuration choice such as the following:

```
Setting up dgx-limits (1.0-1) ...

Configuration file '/etc/security/limits.d/dgxl- limits.conf'

==> File on system created by you or by a script.
==> File also in package provided by package maintainer.
   What would you like to do about it ? Your options are:
   Y or I      : install the package maintainer's version
   N or O      : keep your currently-installed version
   D           : show the differences between the versions
   Z           : start a shell to examine the situation

The default action is to keep your current version.

*** dgxl-limits.conf (Y/I/N/O/D/Z) [default=N] ?
```

Press **Y** to all such messages to “install the package maintainer’s version”.

- If you see the following message:

```
A new version of configuration file /etc/default/grub is available, but
the version installed currently has been locally modified.
What do you want to do about modified configuration file
grub?
  install the package maintainer's version
  keep the local version currently installed
  show the differences between the versions
  show a side-by-side difference between the versions
  show a 3-way difference between available versions
  do a 3-way merge between available versions
  start a new shell to examine the situation
```

Select **install the package maintainer’s version**.

- You may see the following error message during DKMS configuration for the `nv_peer_mem` module while running the script:

```
DKMS: install completed.
modprobe: ERROR: could not insert 'nv_peer_mem': Exec format error
The update to version 3.1.7 will remedy this error.
```

8. When completed, press y at the prompt to reboot the system.

```
Reboot now to finalize 2.1.4? (y/NO) : y
```

9. Wait for the system to come back up, then confirm that the Linux kernel version is 4.4.0-124 or later.

```
$ uname -a
```

Expected output:



```
Linux jws-1 4.4.0-124-generic ...
```

10. Confirm GPU driver version is 384.145.

```
$ nvidia-smi
```

Expected output (first line)

```
NVIDIA-SMI 384.145 Driver Version: 384.145
```

## Recovering from an Interrupted Update

If the script is interrupted during the update, such as from a loss of power or loss of network connection, you can attempt to recover as follows:

1. Restore power or restore the network connection, whichever caused the interruption.
  - If the system encounters a kernel panic after you restore power and reboot the DGX-1, you will not be able to perform the over-the-network update. You will need to re-image the DGX-1 with the latest image (see the [DGX-1 User Guide](#) for instructions) and then perform the network update.
  - If you are successfully returned to the Linux command line, continue the following steps.

1. Reset the script.

```
$ sudo dpkg --configure -a
```

2. Continue following the instructions from step 7 in the [Updating from Version 1.x](#) update instructions.

## UPDATING FROM 2.X TO 3.1.7

See the section [Connecting to the DGX-1 Console](#) for guidance on connecting to the console to perform the update.

### Update Instructions

1. If you have not already done so, verify that your DGX-1 system can access the public repositories as explained in [Verifying the DGX-1 Connection to the Repositories](#).
2. Update the list of available packages and their versions.

```
$ sudo apt-get update
```

3. Install any updates.

```
$ sudo apt-get upgrade -y
```

4. Install dgx-release-upgrade.

```
$ sudo apt-get install -y dgx-release-upgrade
```

5. Begin the update process.

```
$ sudo dgx-release-upgrade
```

If you are using a proxy server, then add the **-E** option to keep your proxy environment variables.

Example:

```
$ sudo -E dgx-release-upgrade
```

6. At the prompt to confirm starting the upgrade, press **Y** to begin.

```
Do you want to start the upgrade?
...
```

```
Installing the upgrade can take several hours. Once the download has
finished, the process cannot be canceled.
```

```
Continue [yN] Details [d]
```

7. At the prompt whether to restart services during the package upgrades without asking, select **Yes**.

```
Restart services during package upgrades without
asking? <Yes> <No>
```

8. After starting the update process, respond to the presented options as follows:

- Select “**keep the local version currently installed**” if there is a new grub package and you see the following text:

```
A new version of configuration file /etc/default/grub is available,
but the version installed currently has been locally modified.
```

```
What do you want to do about modified configuration file grub?
  Install the package maintainer's version
  keep the local version currently installed ← Select
  show the differences between the versions
  show a side-by-side difference between the versions
  show a 3-way difference between available versions
  do a 3-way merge between available versions
  start a new shell to examine the situation
```

The local version contains the changes that have been made on the DGX-1.

- Press **Y** if prompted about InfiniBand configuration choices.

```
Configuration file
`./usr/src/mlnx-ofed-kernel-4.0/ofed_scripts/ib_ipoib.conf'

==> File on system created by you or by a script.
==> File also in package provided by package maintainer.

What would you like to do about it ? Your options are:
  Y or I  : install the package maintainer's version
  N or O  : keep your currently-installed version
  D       : show the differences between the versions
  Z       : start a shell to examine the situation

The default action is to keep your current version.
*** ib_ipoib.conf (Y/I/N/O/D/Z) [default=N] ?
```

- Press **Y** if prompted about docker.list configuration choices.

```
Configuration file '/etc/apt/sources.list.d/docker.list'

==> File on system created by you or by a script.
==> File also in package provided by package maintainer.

What would you like to do about it ? Your options are:
  Y or I  : install the package maintainer's version
  N or O  : keep your currently-installed version
  D       : show the differences between the versions
  Z       : start a shell to examine the situation

The default action is to keep your current version.
*** docker.list (Y/I/N/O/D/Z) [default=N] ?
```

9. Press **Y** to proceed with the final reboot.

```
Restart required
```

```
To finish the upgrade, a restart is required. Upon the next boot,
the system will continue to install several packages in the
background.
```

```
*** This can take another ~7 minutes to complete. DO NOT shutdown or
reboot the system during this period ***.
```

```
If you select 'y' the system will be restarted.
Continue [yN]
```

After this reboot, the update process will take several minutes to perform some final installation steps.

#### 10. Confirm the Linux kernel version.

```
$ uname -a
```

Expected output: (or later version)

```
Linux jws-1 4.4.0-104-generic #116~16.04.1-Ubuntu SMP Mon Aug 14
16:07:05 UTC 2017 x86_64 x86_64 x86_64 GNU/Linux
```

#### 11. Confirm the CUDA driver version.

```
$ nvidia-smi
```

Expected output, first line

```
NVIDIA-SMI 384.145 Driver Version: 384.145
```

## Verifying the nvidia-peer-memory Module

1. Make sure the nvidia-peer-memory module is installed.

```
$ lsmod | grep nv_peer_mem
```

Expected output:

```
nv_peer_mem      16384  0
nvidia           11911168  30 nv_peer_mem,nvidia_modeset,nvidia_uvm
ib_core         143360  13
rdma_cm,ib_cm,ib_sa,iw_cm,nv_peer_mem,mlx4_ib,mlx5_ib,ib_mad,ib_ucm
,ib_umad,ib_uverbs,rdma_ucm,ib_ipoib
```

- If the expected output appears, then no further action is needed.
- If there is no output, then continue the steps to install the nvidia-peer-memory module.

2. Install the module.

```
$ sudo apt-get install --reinstall mlnx-ofed-kernel-dkms nvidia-
peer-memory-dkms
```

Expected output:

```
DKMS: install completed.
Processing triggers for initramfs-tools (0.103ubuntu4.2) ...
update-initramfs: Generating /boot/initrd.img-4.4.0-64-generic
```

3. Add the module to the Linux kernel.

```
$ sudo modprobe nv_peer_mem
```

There is no expected output for this command.

4. Repeat step 1 to verify that the nvidia-peer-memory module is installed.

## Enabling Dynamic DNS Updates

If your network is configured for DHCP, then dynamic DNS updates need to be enabled. You may have already enabled dynamic DNS updates on the DGX OS Server 2.x software, in which case no further action is needed.

To verify, check whether `/etc/resolv.conf` is a link to `/run/resolvconf/resolv.conf`.

```
$ ls -l /etc/resolv.conf
```

Expected output:

```
lrwxrwxrwx 1 root root 29 Dec 1 21:19 /etc/resolv.conf ->
../run/resolvconf/resolv.conf
```

- ▶ If the expected output appears, then no further action is needed.
- ▶ If this does not appear, then enable dynamic DNS updates as follows:
  1. Launch the Resolvconf Reconfigure package.

```
$ sudo dpkg-reconfigure resolvconf
```

The Configuring resolvconf screen appears.

2. Select **<Yes>** when asked whether to prepare `/etc/resolv.conf` for dynamic updates.
3. Select **<No>** when asked whether to append original file to dynamic file.
4. Select **<OK>** at the Reboot recommended screen.
  - You do not need to reboot.
  - You are returned to the command line.
5. Bring down the interface, where `<network-interface>` is the name of your primary network interface (em1, em2, enp1s0f0, or enp1s0f1).

```
$ sudo ifup <network-interface>
```

Expected output:

```
ifdown: interface <network interface> not configured
```

6. Bring up the interface.

```
$ sudo ifup <network interface>
```

Expected output (last line):

```
...
bound to <IP address> -- renewal in ...
```

## Recovering from an Interrupted Update

If the script is interrupted during the update, such as from a loss of power or loss of network connection, restore power or restore the network connection, whichever caused the interruption.

- ▶ If you have *not* rebooted the system and are successfully returned to the Linux command line, restart the update by entering the following, replacing `*` with the actual string of characters that identify the directory where the tarball was extracted:

```
$ cd /tmp/ubuntu-release-upgrade-*/
$ sudo ./xenial
```

- ▶ If you rebooted the system and are successfully returned to the Linux command line, you will need to download and extract the tarball manually.

```
$ sudo apt-get install -y wget
$ wget
http://international.download.nvidia.com/dgx/repos/release_updates/xenial.tar.gz
$ tar -zxvf xenial.tar.gz -C <extract_path>
```

Then restart the failed update:

```
$ cd <extracted_path>
$ sudo ./xenial
```

- ▶ If the system encounters a kernel panic after you restore power and reboot the DGX-1, you will not be able to perform the over-the-network update. You will need to re-image the DGX-1 with the latest image (see the [DGX-1 User Guide](#) for instructions) and then perform the network update.

## UPDATING FROM 3.1.X TO 3.1.7

See the section [Connecting to the DGX-1 Console](#) for guidance on connecting to the console to perform the update.



**CAUTION:** These instructions update all software for which updates are available from your configured software sources, including applications that you installed yourself. If you want to prevent an application from being updated, you can instruct the Ubuntu package manager to keep the current version. For more information, see [Introduction to Holding Packages](#) on the Ubuntu Community Help Wiki.

### Update Instructions

1. If you have not already done so, verify that your DGX-1 system can access the public repositories as explained in [Verifying the DGX-1 Connection to the Repositories](#).
2. Update the list of available packages and their versions.

```
$ sudo apt update
```

3. Review the packages that will be updated.

```
$ sudo apt full-upgrade -s
```

To prevent an application from being updated, instruct the Ubuntu package manager to keep the current version. See [Introduction to Holding Packages](#).

4. Upgrade to version 3.1.7.

```
$ sudo apt full-upgrade
```

- Answer any questions that appear.
    - Most questions require a Yes or No response. When asked to select the grub configuration to use, select the current one on the system.
    - Other questions will depend on what other packages were installed before the update and how those packages interact with the update.
  - If a message appears indicating that `nvidia-docker.service` failed to start, you can disregard it and continue with the next step. The service will start normally at that time.
5. Reboot the system.



## Verifying the nvidia-peer-memory Module

1. Make sure the nvidia-peer-memory module is installed.

```
$ lsmod | grep nv_peer_mem
```

Expected output:

```
nv_peer_mem      16384  0
nvidia           11911168  30 nv_peer_mem,nvidia_modeset,nvidia_uvm
ib_core         143360  13
rdma_cm,ib_cm,ib_sa,iw_cm,nv_peer_mem,mlx4_ib,mlx5_ib,ib_mad,ib_ucm
,ib_umad,ib_uverbs,rdma_ucm,ib_ipoib
```

- If the expected output appears, then no further action is needed.
- If there is no output, then continue the steps to install the nvidia-peer-memory module.

2. Install the module.

```
$ sudo apt-get install --reinstall mlnx-ofed-kernel-dkms nvidia-
peer-memory-dkms
```

Expected output:

```
DKMS: install completed.
Processing triggers for initramfs-tools (0.103ubuntu4.2) ...
update-initramfs: Generating /boot/initrd.img-4.4.0-64-generic
```

3. Add the module to the Linux kernel.

```
$ sudo modprobe nv_peer_mem
```

There is no expected output for this command.

4. Repeat step 1 to verify that the nvidia-peer-memory module is installed.

## Recovering from an Interrupted or Failed Update

If the script is interrupted during the update, such as from a loss of power or loss of network connection, then restore power or restore the network connection, whichever caused the interruption.

- ▶ If the system encounters a kernel panic after you restore power and reboot the DGX-1, you will not be able to perform the over-the-network update. You will need to re-image the DGX-1 with the latest image (see the [DGX-1 User Guide](#) for instructions) and then perform the network update.
- ▶ If you are successfully returned to the Linux command line, continue following the instructions from step 2 in the [Updating from Version 3.1.x to 3.1.6](#) update instructions.

# APPENDIX A. THIRD PARTY LICENSE NOTICE

This NVIDIA product contains third party software that is being made available to you under their respective open source software licenses. Some of those licenses also require specific legal information to be included in the product. This section provides such information.

## libressl (OpenSSL)

LibreSSL <http://www.libressl.org/> version 2.4.1 is provided under the following terms:

```
* Copyright (c) 1998-2011 The OpenSSL Project. All rights reserved.
*
* Redistribution and use in source and binary forms, with or without
* modification, are permitted provided that the following conditions *
* are met:
*
* 1. Redistributions of source code must retain the above copyright
* notice, this list of conditions and the following disclaimer.
*
* 2. Redistributions in binary form must reproduce the above copyright
* notice, this list of conditions and the following disclaimer in
* the documentation and/or other materials provided with the
* distribution.
*
* 3. All advertising materials mentioning features or use of this
* software must display the following acknowledgment:
* "This product includes software developed by the OpenSSL Project
* for use in the OpenSSL Toolkit. (http:// www.openssl.org/)"
*
```

```

* 4. The names "OpenSSL Toolkit" and "OpenSSL Project" must not be used
to * endorse or promote products derived from this software without
* prior written permission. For written permission, please contact
* openssl-core@openssl.org.
*
* 5. Products derived from this software may not be called "OpenSSL"
* nor may "OpenSSL" appear in their names without prior written *
permission of the OpenSSL Project.
*
* 6. Redistributions of any form whatsoever must retain the following
* acknowledgment:
* "This product includes software developed by the OpenSSL Project
* for use in the OpenSSL Toolkit (http:// www.openssl.org/)"
*
* THIS SOFTWARE IS PROVIDED BY THE OpenSSL PROJECT ``AS IS'' AND ANY
* EXPRESSED OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE
* IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR
* PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE OpenSSL PROJECT OR
* ITS CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL,
* SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT
* NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES;
* LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION)
* HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT,
* STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE)
* ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED
* OF THE POSSIBILITY OF SUCH DAMAGE.

```

## Mellanox (OFED)

MLNX OFED (<http://www.mellanox.com/>) is provided under the following terms:

```

Copyright (c) 2006 Mellanox Technologies.
All rights reserved.

```

```

Redistribution and use in source and binary forms, with or without
modification, are permitted provided that the following conditions are
met:

```

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

```

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS
IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED
TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A
PARTICULAR PURPOSE ARE DISCLAIMED.

```

```

IN NO EVENT SHALL THE COPYRIGHT HOLDER OR CONTRIBUTORS BE LIABLE FOR
ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL
DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS
OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION)

```

HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

## Notice

THE INFORMATION IN THIS DOCUMENT AND ALL OTHER INFORMATION CONTAINED IN NVIDIA DOCUMENTATION REFERENCED IN THIS DOCUMENT IS PROVIDED "AS IS." NVIDIA MAKES NO WARRANTIES, EXPRESSED, IMPLIED, STATUTORY, OR OTHERWISE WITH RESPECT TO THE INFORMATION FOR THE PRODUCT, AND EXPRESSLY DISCLAIMS ALL IMPLIED WARRANTIES OF NONINFRINGEMENT, MERCHANTABILITY, AND FITNESS FOR A PARTICULAR PURPOSE. Notwithstanding any damages that customer might incur for any reason whatsoever, NVIDIA's aggregate and cumulative liability towards customer for the product described in this document shall be limited in accordance with the NVIDIA terms and conditions of sale for the product.

THE NVIDIA PRODUCT DESCRIBED IN THIS DOCUMENT IS NOT FAULT TOLERANT AND IS NOT DESIGNED, MANUFACTURED OR INTENDED FOR USE IN CONNECTION WITH THE DESIGN, CONSTRUCTION, MAINTENANCE, AND/OR OPERATION OF ANY SYSTEM WHERE THE USE OR A FAILURE OF SUCH SYSTEM COULD RESULT IN A SITUATION THAT THREATENS THE SAFETY OF HUMAN LIFE OR SEVERE PHYSICAL HARM OR PROPERTY DAMAGE (INCLUDING, FOR EXAMPLE, USE IN CONNECTION WITH ANY NUCLEAR, AVIONICS, LIFE SUPPORT OR OTHER LIFE CRITICAL APPLICATION). NVIDIA EXPRESSLY DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY OF FITNESS FOR SUCH HIGH RISK USES. NVIDIA SHALL NOT BE LIABLE TO CUSTOMER OR ANY THIRD PARTY, IN WHOLE OR IN PART, FOR ANY CLAIMS OR DAMAGES ARISING FROM SUCH HIGH RISK USES.

NVIDIA makes no representation or warranty that the product described in this document will be suitable for any specified use without further testing or modification. Testing of all parameters of each product is not necessarily performed by NVIDIA. It is customer's sole responsibility to ensure the product is suitable and fit for the application planned by customer and to do 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 does not accept any 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.

Other than the right for customer to use the information in this document with the product, no other license, either expressed or implied, is hereby granted by NVIDIA under this document. Reproduction of information in this document is permissible only if reproduction is approved by NVIDIA in writing, is reproduced without alteration, and is accompanied by all associated conditions, limitations, and notices.

## Trademarks

NVIDIA, the NVIDIA logo, and DGX-1 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

© 2018 NVIDIA Corporation. All rights reserved.