




# NVIDIA DGX OS SERVER VERSION 4.0.2

DA-08260-402\_v01 | September 2018

**Release Notes and Update Guide**



# TABLE OF CONTENTS

<b>NVIDIA DGX OS Server Release Notes for Version 4.0.2</b> .....	<b>3</b>
Update Advisement .....	4
About Release 4.0 .....	4
Highlights in Version 4.0.2 .....	5
Version History .....	5
Version 4.0.2 .....	5
Version 4.0.1 .....	5
Known Software Issues .....	6
Known DGX-2 System Issues .....	8
Known Ubuntu / Linux Kernel Issues .....	10
DGX OS Server Software Content .....	13
DGX Server Firmware Version Reference .....	13
DGX-2 Firmware .....	13
<b>Appendix A. Third Party License Notice</b> .....	<b>14</b>

# NVIDIA DGX OS SERVER RELEASE NOTES FOR VERSION 4.0.2

This document describes version 4.0.2 of the NVIDIA® DGX™ OS Server Release 4.0 software and update package.

DGX OS Server v4.0.2 is provided as an ISO image to the factory and the ISO image is also available from NVIDIA Enterprise Support in the event the server needs to be re-imaged. The software is also provided as an “over-the-network” update.

Refer to the *DGX-2 User Guide* (<http://docs.nvidia.com/dgx/dgx2-user-guide/index.html>) for instructions on

- ▶ How to re-image the system with the ISO image,
- ▶ How to update the software through an internet connection
- ▶ How to install the software on air-gapped systems

## UPDATE ADVISEMENT

### ▶ NVIDIA Docker Containers

In conjunction with DGX OS Server v4.0.2, 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 server up to date with the latest Ubuntu security updates using the 'apt-get dist-upgrade' procedure. See the [Ubuntu Wiki Upgrades](#) web page for more information. Also, the Ubuntu Security Notice site (<https://usn.ubuntu.com/>) lists known Common Vulnerabilities and Exposures (CVEs), including those that can be resolved by updating the DGX OS software.

## ABOUT RELEASE 4.0

The following are the primary features of the DGX OS Server Release 4.0:

- ▶ Supports the NVIDIA DGX-2 System.
- ▶ Ubuntu 18.04 LTS
- ▶ NVIDIA GPU Driver Release 410
  - Supports CUDA 10.0
- ▶ Docker CE and the NVIDIA Container Runtime for Docker are pre-installed, and the docker daemon automatically launched.
- ▶ New NVSM health monitoring software framework
  - Replaces `nvsysinfo` and `nvhealth`.

---

<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 <https://docs.nvidia.com/deeplearning/dgx/user-guide/index.html> for instructions on how to access them.

## HIGHLIGHTS IN VERSION 4.0.2

- ▶ Supports NVIDIA DGX-2 System only  
DGX OS version 4.0.2 does not support DGX-1 servers.
- ▶ NVIDIA GPU Driver Version 410.48

## VERSION HISTORY

### Version 4.0.2

- ▶ Updated NVIDIA GPU driver to version 410.48.

### Version 4.0.1

- ▶ Initial release for the NVIDIA DGX-2 System
- ▶ NVIDIA GPU driver version 410.47
- ▶ See [About Release 4.0](#) for additional features list

## KNOWN SOFTWARE ISSUES

The following are known issues with the software.

- ▶ [Health Monitor Processes Listen on IPv4 and IPv6 Interfaces](#)
- ▶ [GPUs Cannot be Reset While the System is Running](#)
- ▶ [Apparmor Profile May not Work with Some Containers](#)

### Health Monitor Processes Listen on IPv4 and IPv6 Interfaces

#### Issue

Two processes within the NVIDIA Health Monitor (NVSM) framework - "mosquitto" and "nvsm\_apis" - listen on all IPv4 and IPv6 interfaces, on ports 273 (nvsm\_apis) and 1883 (mosquitto).

#### Workaround

If this violates your organization's security policies, you can disable the associated services.

#### List of services:

- mosquito
- nvsm-apis
- nvsm-apis-plugin-environment
- nvsm-apis-mqtt
- nvsm-apis-plugin-memory
- nvsm-apis-mongodb
- nvsm-apis
- nvsm-apis-selwatcher
- nvsm-storage-dshm
- nvsm-env-dshm
- nvsm-sys-dshm

#### Instructions for removing

Perform the following for each service

```
sudo systemctl stop <service>
sudo systemctl disable <service>
```



**NOTE:** This disables all NVSM health monitoring functionality.

## 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.

## KNOWN DGX-2 SYSTEM ISSUES

The following are known issues specific to the DGX-2 server.

- ▶ [Hot-plugging of Storage NVMe is not Supported](#)
- ▶ [Storage NVMe Removal Results in Two Removed Devices](#)
- ▶ [System BIOS Password Feature is Ineffective](#)
- ▶ [Micron Storage NVMe Temperature Monitoring Does Not Work](#)
- ▶ [BMC SNMP Community String Limitations](#)
- ▶ [Some BMC Dashboard Quick Links Appear Erroneously](#)

### Hot-plugging of Storage NVMe is not Supported

#### Issue

Hot-plugging or hot-swapping one of the storage non-volatile memory express (NVMe) drive might result in system instability.

#### Resolution

Turn off the system before removing and replacing any of the storage NVMe drives. This will be resolved in a future software update.

### Storage NVMe Removal May Result in Removal of Different NVMe

#### Issue

When attempting to remove access to an NVMe drive using the following command,

```
echo 1 > /sys/class/nvme/nvmeX/device/remove
```

where **X** specifies which NVMe drive to remove, other NVMe drives may get removed.

#### Resolution

This is the result of an issue in the NVMe driver, and will be resolved in a future software update. To work around, shut down the system before removing the NVMe drive.



## System BIOS Password Feature is Ineffective

### Issue

After setting a password for accessing the system BIOS settings, a user can access the settings by pressing return when prompted to enter a password,

### Resolution

This issue will be resolved in a future software update.

## Micron Storage NVMe Temperature Monitoring Does Not Work

### Issue

Some Micron storage NVMe drives shipped with the server do not have temperature monitoring enabled.

### Workaround

Enable temperature monitoring as follows:

1. From the command line, check which NVMe drives need to have temperature monitoring enabled.

```
sudo /opt/MicronTechnology/MicronMSECLI/msecli -L
```

2. Note the Device Name that does not show a temperature reading, then enter the following for each affected NVMe drive.

```
sudo /opt/MicronTechnology/MicronMSECLI/msecli -M -k 1 -n <Device Name>
```

## BMC SNMP Community String Limitations

### Issue

The DGX-2 BMC has the following SNMP Community String limitations:

- No support for SNMPv3
- No SNMP configuration controls in the BMC dashboard
- No support for setting RO and RW permissions from the command line ipmitool.

### Resolution

This will be resolved in a future BMC firmware release.

## Some BMC Dashboard Quick Links Appear Erroneously

### Issue

On the BMC dashboard, the following Quick Links appear by mistake and should not be used.

- Maintenance->Firmware Update
- Settings->NvMeManagement->NvMe P3700Vpd Info

### Resolution

These quick links will be removed from the menu in a future BMC firmware release.

## KNOWN UBUNTU / LINUX KERNEL ISSUES

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

- ▶ [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 the [NVIDIA Enterprise Support portal](#)).

## 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 4.0.2 release, the Ubuntu kernel (4.15.0-29) 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 ISO image.

Component	Version
DGX OS Server	4.0.2
GPU Driver	410.48
NVIDIA Container Runtime for Docker	2.0.3
Ubuntu	18.04 LTS
Ubuntu kernel	4.15.0-29 or later
Docker CE	18.06.0-ce
NVIDIA System Health Monitor (NVSM)	nvsm-cli 18.08.2 nvsm-dshm 18.08.4 nvsm-apis 18.08.3
Data Center GPU Management (DCGM)	1.5.2
Mellanox OFED	MLNX 4.4-2.0.7.0

## DGX SERVER FIRMWARE VERSION REFERENCE

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

### DGX-2 Firmware

Component	Version	Notes
BMC	V0.96.01	Released versions for DGX-2 at the time of this software release. Information provided for reference purposes.
SBIOS	V0.14	
VBIOS	88.00.6B.00.01	
InfiniBand FW	16.23.1020	
PSU FW	2.1	
SSD (Samsung OS drive)	CXV8601Q	
SSD (Micron storage drive)	101008R0	

# 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.

## msecli

The msecli utility (<https://www.micron.com/products/solid-state-storage/storage-executive-software>) is provided under the following terms:

Micron Technology, Inc. Software License Agreement

PLEASE READ THIS LICENSE AGREEMENT ("AGREEMENT") FROM MICRON TECHNOLOGY, INC. ("MTI") CAREFULLY: BY INSTALLING, COPYING OR OTHERWISE USING THIS SOFTWARE AND ANY RELATED PRINTED MATERIALS ("SOFTWARE"), YOU ARE ACCEPTING AND AGREEING TO THE TERMS OF THIS AGREEMENT. IF YOU DO NOT AGREE WITH THE TERMS OF THIS AGREEMENT, DO NOT INSTALL THE SOFTWARE.

LICENSE: MTI hereby grants to you the following rights: You may use and make one

(1) backup copy the Software subject to the terms of this Agreement. You must maintain all copyright notices on all copies of the Software. You agree not to modify, adapt, decompile, reverse engineer, disassemble, or otherwise translate the Software. MTI may make changes to the Software at any time without notice to you.

In addition MTI is under no obligation whatsoever to update, maintain,

or provide new versions or other support for the Software.

**OWNERSHIP OF MATERIALS:** You acknowledge and agree that the Software is proprietary property of MTI (and/or its licensors) and is protected by United States copyright law and international treaty provisions. Except as expressly provided herein, MTI does not grant any express or implied right to you under any patents, copyrights, trademarks, or trade secret information. You further acknowledge and agree that all right, title, and interest in and to the Software, including associated proprietary rights, are and shall remain with MTI (and/or its licensors). This Agreement does not convey to you an interest in or to the Software, but only a limited right to use and copy the Software in accordance with the terms of this Agreement. The Software is licensed to you and not sold.

**DISCLAIMER OF WARRANTY:** THE SOFTWARE IS PROVIDED "AS IS" WITHOUT WARRANTY OF ANY KIND. MTI EXPRESSLY DISCLAIMS ALL WARRANTIES EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO, NONINFRINGEMENT OF THIRD PARTY RIGHTS, AND ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR ANY PARTICULAR PURPOSE. MTI DOES NOT WARRANT THAT THE SOFTWARE WILL MEET YOUR REQUIREMENTS, OR THAT THE OPERATION OF THE SOFTWARE WILL BE UNINTERRUPTED OR ERROR-FREE. FURTHERMORE, MTI DOES NOT MAKE ANY REPRESENTATIONS REGARDING THE USE OR THE RESULTS OF THE USE OF THE SOFTWARE IN TERMS OF ITS CORRECTNESS, ACCURACY, RELIABILITY, OR OTHERWISE. THE ENTIRE RISK ARISING OUT OF USE OR PERFORMANCE OF THE SOFTWARE REMAINS WITH YOU. IN NO EVENT SHALL MTI, ITS AFFILIATED COMPANIES OR THEIR SUPPLIERS BE LIABLE FOR ANY DIRECT, INDIRECT, CONSEQUENTIAL, INCIDENTAL, OR SPECIAL DAMAGES (INCLUDING, WITHOUT LIMITATION, DAMAGES FOR LOSS OF PROFITS, BUSINESS INTERRUPTION, OR LOSS OF INFORMATION) ARISING OUT OF YOUR USE OF OR INABILITY TO USE THE SOFTWARE, EVEN IF MTI HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. Because some jurisdictions prohibit the exclusion or limitation of liability for consequential or incidental damages, the above limitation may not apply to you.

**TERMINATION OF THIS LICENSE:** MTI may terminate this license at any time if you are in breach of any of the terms of this Agreement. Upon termination, you will immediately destroy all copies the Software.

**GENERAL:** This Agreement constitutes the entire agreement between MTI and you regarding the subject matter hereof and supersedes all previous oral or written communications between the parties. This Agreement shall be governed by the laws of the State of Idaho without regard to its conflict of laws rules.

**CONTACT:** If you have any questions about the terms of this Agreement, please contact MTI's legal department at (208) 368-4500.

By proceeding with the installation of the Software, you agree to the terms of this Agreement. You must agree to the terms in order to install and use the Software.

## 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, DGX, DGX-1, and DGX-2 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.