



## **Deploying NVIDIA Converged Accelerator**

# Table of contents

## Configuring Operation Mode

---

Standard Mode

---

BlueField-X Mode

---

## Verifying Configured Operational Mode

---

## Verifying GPU Ownership

---

## GPU Firmware

---

Updating GPU Firmware

---

Get GPU Firmware

---

## **Note**

This page is relevant for A30 and A100 converged accelerator cards. For instructions on the NVIDIA® H20 NVL16 device (configured to run in NIC mode only), refer to [Deploying BlueField Software from Host](#).

## **Info**

It is recommended to upgrade your BlueField product to the latest software and firmware versions available to benefit from new features and latest bug fixes.

This section assumes that you have installed the BlueField OS BFB on your NVIDIA® Converged Accelerator using any of the following guides:

- [Deploying BlueField Software from Host](#)
- [Deploying BlueField Software from BlueField BMC](#)
- [Deploying BlueField Software Using PXE](#)

NVIDIA® CUDA® (GPU driver) must be installed to use the GPU. For information on how to install CUDA on your Converged Accelerator, refer to [NVIDIA CUDA Installation Guide for Linux](#).

## **Configuring Operation Mode**

After installing the BFB, you may now select the mode you want your NVIDIA Converged Accelerator to operate in.

- Standard (default) – the NVIDIA® BlueField® and the GPU operate separately (GPU is owned by the host)
- BlueField-X – the GPU is exposed to BlueField and is no longer visible on the host (GPU is owned by BlueField)

## **Note**

It is important to know your device name (e.g., `mt41686_pciconf0`).

MST tool is necessary for this purpose which is installed by default on the DPU.

Run:

```
mst status -v
```

Example output:

```
MST modules:
-----
    MST PCI module is not loaded
    MST PCI configuration module loaded
PCI devices:
-----
DEVICE_TYPE          MST
PCI      RDMA        NET
NUMA
BlueField2(rev:1)
/dev/mst/mt41686_pciconf0.1  3b:00.1  mlx5_1
net-ens1f1                   0

BlueField2(rev:1)          /dev/mst/mt41686_pciconf0
3b:00.0  mlx5_0           net-ens1f0
0
```

## **BlueField-X Mode**

1. Run the following command from the host:

```
mlxconfig -d /dev/mst/<device-name> s  
PCI_DOWNSTREAM_PORT_OWNER[4]=0xF
```

2. Perform a BlueField system-level reset for the `mlxconfig` settings to take effect.

## Standard Mode

To return BlueField from BlueField-X mode to Standard mode:

1. Run the following command from the host:

```
mlxconfig -d /dev/mst/<device-name> s  
PCI_DOWNSTREAM_PORT_OWNER[4]=0x0
```

2. Perform a BlueField system-level reset for the `mlxconfig` settings to take effect.

## Verifying Configured Operational Mode

Use the following command from the host or BlueField:

```
$ sudo mlxconfig -d /dev/mst/<device-name> q  
PCI_DOWNSTREAM_PORT_OWNER[4]
```

- Example of Standard mode output:

```
Device #1:
```

```
-----  
[...]  
Configurations:                                Next Boot  
          PCI_DOWNSTREAM_PORT_OWNER[4]  
DEVICE_DEFAULT(0)
```

- Example of BlueField-X mode output:

```
Device #1:  
-----  
[...]  
Configurations:                                Next Boot  
          PCI_DOWNSTREAM_PORT_OWNER[4]         EMBEDDED_CPU(15)
```

## Verifying GPU Ownership

The following are example outputs for when BlueField is configured to BlueField-X mode.

The GPU is no longer visible from the host:

```
root@host:~# lspci | grep -i nv  
None
```

The GPU is now visible from BlueField:

```
ubuntu@bf:~$ lspci | grep -i nv  
06:00.0 3D controller: NVIDIA Corporation GA20B8 (rev a1)
```

# GPU Firmware

## Get GPU Firmware

smbpbi: (See SMBPBI spec)

```
root@bf:~# i2cset -y 3 0x4f 0x5c 0x05 0x08 0x00 0x80 s
```

```
root@bf:~# i2cget -y 3 0x4f 0x5c ip 5
```

```
5: 0x04 0x05 0x08 0x00 0x5f
```

```
root@bf:~# i2cget -y 3 0x4f 0x5d ip 5
```

```
5: 0x04 0x39 0x32 0x2e 0x30
```

```
root@bf:~#
```

```
root@bf:~#
```

```
root@bf:~# i2cset -y 3 0x4f 0x5c 0x05 0x08 0x01 0x80 s
```

```
root@bf:~# i2cget -y 3 0x4f 0x5c ip 5
```

```
5: 0x04 0x05 0x08 0x01 0x5f
```

```
root@bf:~# i2cget -y 3 0x4f 0x5d ip 5
```

```
5: 0x04 0x30 0x2e 0x36 0x42
```

```
root@bf:~# i2cset -y 3 0x4f 0x5c 0x05 0x08 0x02 0x80 s
```

```
root@bf:~# i2cget -y 3 0x4f 0x5c ip 5
```

```
5: 0x04 0x05 0x08 0x02 0x5f
```

```
root@bf:~# i2cget -y 3 0x4f 0x5d ip 5
```

```
5: 0x04 0x2e 0x30 0x30 0x2e
```

```
root@bf:~# i2cset -y 3 0x4f 0x5c 0x05 0x08 0x03 0x80 s
```

```
root@bf:~# i2cget -y 3 0x4f 0x5c ip 5
```

```
5: 0x04 0x05 0x08 0x03 0x5f
```

```
root@bf:~# i2cget -y 3 0x4f 0x5d ip 5
```

```
5: 0x04 0x30 0x31 0x00 0x00
```

```
root@bf:~#
```

```
39 32 2e 30 30 2e 36 42 2e 30 30 2e 30 31 00 00 92.00.6B.00.01
```

## Updating GPU Firmware

```
root@bf:~# scp root@10.23.201.227:~/<path-to-fw-  
bin>/1004_0230_891__92006B0001-dbg-ota.bin /tmp/gpu_images/  
root@10.23.201.227's password:  
1004_0230_891__92006B0001-dbg-ota.bin  
100% 384KB 384.4KB/s 00:01
```

```
root@bf:~# cat /tmp/gpu_images/progress.txt  
TaskState="Running"  
TaskStatus="OK"  
TaskProgress="50"
```

```
root@bf:~# cat /tmp/gpu_images/progress.txt  
TaskState="Running"  
TaskStatus="OK"  
TaskProgress="50"
```

```
root@bf:~# cat /tmp/gpu_images/progress.txt  
TaskState=Firmware update succeeded.  
TaskStatus=OK  
TaskProgress=100
```

## Notice

This document is provided for information purposes only and shall not be regarded as a warranty of a certain functionality, condition, or quality of a product. NVIDIA Corporation (“NVIDIA”) makes no representations or warranties, expressed or implied, as to the accuracy or completeness of the information contained in this document and assumes no responsibility for any errors contained herein. NVIDIA shall have no liability for the consequences or use of such information or for any infringement of patents or other rights of third parties that may result from its use. This document is not a commitment to develop, release, or deliver any Material (defined below), code, or functionality.

NVIDIA reserves the right to make corrections, modifications, enhancements, improvements, and any other changes to this document, at any time without notice.

Customer should obtain the latest relevant information before placing orders and should verify that such information is current and complete.

NVIDIA products are sold subject to the NVIDIA standard terms and conditions of sale supplied at the time of order acknowledgement, unless otherwise agreed in an individual sales agreement signed by authorized representatives of NVIDIA and customer (“Terms of Sale”). NVIDIA hereby expressly objects to applying any customer general terms and conditions with regards to the purchase of the NVIDIA product referenced in this document. No contractual obligations are formed either directly or indirectly by this document.

NVIDIA products are not designed, authorized, or warranted to be suitable for use in medical, military, aircraft, space, or life support equipment, nor in applications where failure or malfunction of the NVIDIA product can reasonably be expected to result in personal injury, death, or property or environmental damage. NVIDIA accepts no liability for inclusion and/or use of NVIDIA products in such equipment or applications and therefore such inclusion and/or use is at customer’s own risk.

NVIDIA makes no representation or warranty that products based on this document will be suitable for any specified use. Testing of all parameters of each product is not necessarily performed by NVIDIA. It is customer’s sole responsibility to evaluate and determine the applicability of any information contained in this document, ensure the product is suitable and fit for the application planned by customer, and perform the necessary testing for the application in order to avoid a default of the application or the product. Weaknesses in customer’s product designs may affect the quality and reliability of the NVIDIA product and may result in additional or different conditions and/or requirements beyond those contained in this document. NVIDIA accepts no liability related to any default, damage, costs, or problem which may be based on or attributable to: (i) the use of the NVIDIA product in any manner that is contrary to this document or (ii) customer product designs.

No license, either expressed or implied, is granted under any NVIDIA patent right, copyright, or other NVIDIA intellectual property right under this document. Information published by NVIDIA regarding third-party products or services does not constitute a license from NVIDIA to use such products or services or a warranty or endorsement thereof. Use of such information may require a license from a third party under the patents or other intellectual property rights of the third party, or a license from NVIDIA under the patents or other intellectual property rights of NVIDIA.

Reproduction of information in this document is permissible only if approved in advance by NVIDIA in writing, reproduced without alteration and in full compliance with all applicable export laws and regulations, and accompanied by all associated conditions, limitations, and notices.

THIS DOCUMENT AND ALL NVIDIA DESIGN SPECIFICATIONS, REFERENCE BOARDS, FILES, DRAWINGS, DIAGNOSTICS, LISTS, AND OTHER DOCUMENTS (TOGETHER AND SEPARATELY, “MATERIALS”) ARE BEING PROVIDED “AS IS.” NVIDIA MAKES NO WARRANTIES, EXPRESSED, IMPLIED, STATUTORY, OR OTHERWISE WITH RESPECT TO THE MATERIALS, AND EXPRESSLY DISCLAIMS ALL IMPLIED WARRANTIES OF NONINFRINGEMENT, MERCHANTABILITY, AND FITNESS FOR A PARTICULAR PURPOSE. TO THE EXTENT NOT PROHIBITED BY LAW, IN NO EVENT WILL NVIDIA BE LIABLE FOR ANY DAMAGES, INCLUDING WITHOUT LIMITATION ANY DIRECT, INDIRECT, SPECIAL, INCIDENTAL, PUNITIVE, OR CONSEQUENTIAL DAMAGES, HOWEVER CAUSED AND REGARDLESS OF THE THEORY OF LIABILITY, ARISING OUT OF

ANY USE OF THIS DOCUMENT, EVEN IF NVIDIA HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. Notwithstanding any damages that customer might incur for any reason whatsoever, NVIDIA's aggregate and cumulative liability towards customer for the products described herein shall be limited in accordance with the Terms of Sale for the product.

## **Trademarks**

NVIDIA and the NVIDIA logo are trademarks and/or registered trademarks of NVIDIA Corporation in the U.S. and other countries. Other company and product names may be trademarks of the respective companies with which they are associated.

© Copyright 2025, NVIDIA. PDF Generated on 11/20/2025