



Release Notes Change Log History

Table of contents

[Changes and New Features in 4.9.0](#)

[Changes and New Features in 4.8.0](#)

[Changes and New Features in 4.7.0](#)

[Changes and New Features in 4.6.0](#)

[Changes and New Features in 4.5.0](#)

[Changes and New Features in 4.2.0](#)

[Changes and New Features in 4.0.3](#)

[Changes and New Features in 4.0.2](#)

[Changes and New Features in 3.9.3](#)

[Changes and New Features in 3.9.2](#)

[Changes and New Features in 3.9.0](#)

[Changes and New Features in 3.8.5](#)

[Changes and New Features in 3.8.0](#)

Changes and New Features in 4.9.0

- In the [devlink](#) and [mlxdevm](#) APIs, added support for setting a maximum number of completion event queues for SFs. Setting I/O completion EQs causes the netdev of the SF to use an equal number of TX and RX channels as I/O completion EQs.
- Added the ability to designate a part of BlueField-3's [L3 cache](#) as exclusive to NIC.

Note

Allocating more exclusive cache to NIC improves networking performance at the expense of applications running on the BlueField-3 Arm OS.

- Enhanced firmware reset flow for Sync1 utilizing community-accepted hot reset kernel flow
- BlueField B3220 [DDR speed](#) can be configured to either 5200Mhz or 5600Mhz using the Arm UEFI menu if the hardware revision supports these speeds.

Info

DDR speed default remains 5200Mhz to maintain backward compatibility.

- Added support to the `bfcfg` script to dump extra system configuration information. The `-l` parameter can be used along with `-d` to set the dump level.

Changes and New Features in 4.8.0

- Added support for bf.cfg file of size up to 128 KB

- Added support for transferring ownership of RShim driver from the host to the BlueField BMC
- Introduced new behavior for 150W platforms where, to ensure proper operation of the BlueField, ATF will not boot BlueField-3 platforms if the ATX +12V is not connected

Changes and New Features in 4.7.0

- Added support for new BlueField reset and reboot procedures for loading new firmware and firmware configuration changes which replace previous need for server power cycle
- Updated the default operation mode of SuperNICs to NIC mode (from DPU mode). This is relevant to the following SKUs:
 - 900-9D3B4-00CC-EA0
 - 900-9D3B4-00SC-EA0
 - 900-9D3B4-00CV-EA0
 - 900-9D3B4-00SV-EA0
 - 900-9D3B4-00EN-EA0
 - 900-9D3B4-00PN-EA0
 - 900-9D3D4-00EN-HA0
 - 900-9D3D4-00NN-HA0

Note

When upgrading one of these SuperNICs to 2.7.0, if its mode of operation was changed at any point in the past, then the last configured mode of operation will remain unchanged. Otherwise, the SuperNIC will rise in NIC operation mode.

- Installing the BFB Bundle now performs NIC firmware update by default
- Added ability to install NIC firmware and BMC software in NIC mode in NVIDIA® BlueField®-3.

(i) Note

It is important to note the following:

- During BFB Bundle installation, Linux is expected to boot to upgrade NIC firmware and BMC software
- As Linux is booting during BFB Bundle installation, it is expected for the mlx5 core driver to timeout on the BlueField Arm
- During the BFB Bundle installation, it is expected for the mlx5 driver to error messages on the x86 host. These prints may be ignored as they are resolved by a mandatory, post-installation power cycle.
- It is mandatory to power cycle the host after the installation is complete for the changes to take effect

- Software packaging – new BlueField firmware bundle package (`bf-fwbundle-<version>.prod.bfb`), a smaller image for Day 2 upgrades, without the OS and DOCA runtime. Includes ATF, UEFI, nic-fw, bmc-fw, and eROT only.
- Improved BlueField BMC robustness –
 - Report LLDP for L2 discovery via Redfish
 - Improved BlueField DPU debuggability
- Increased support for virtio-net VF devices on BlueField-3 networking platforms to 2K
- Reduced power consumption for BlueField NIC mode

- RAS
 - Report DDR Error to OS, including both single-bit ECC error and UCE error
 - Support error injection in processors, memory, and PCIe devices

Changes and New Features in 4.6.0

- Updated minimum [UEFI password requirements](#)
- Included DPU BMC firmware as part of the BFB image
- Added virtio-net support for plugging/unplugging parallel devices
- Implemented virtio debug enhancements

Changes and New Features in 4.5.0

- Added Redfish support for configuring all UEFI secure boot settings (disable, enable, enroll user keys, etc.) at scale, remotely, and securely
- For FHHL DPUs, added support for performing PCIe bifurcation configuration via MFT tool

Note

Only a subset of configurations are supported.

- Updated the print of the manufacturing (MFG) setting, `MFG_00B_MAC`, displayed by the command `bfcfg -d` to appear in lower-case to align with standard Linux tools

Changes and New Features in 4.2.0

Note

Upgrading to this BSP version installs a new version of Ubuntu GRUB. This version of GRUB revokes the old UEFI secure boot certificates and install new ones. The new certificates will not validate older images and boot will fail. Therefore, to roll back to older software versions, users must disable UEFI secure boot.

- BFB installation chooses the on-chip NVMe (`/dev/nvme0n1`) by default for the EFI system partition and Linux rootfs installation and can be overloaded with `device=/dev/mmcblk0` in `bf.cfg` to push together with the BFB.

Info

Installing on NVMe causes DPU booting to stay at the UEFI shell when changing to Livefish mode.

Info

A previously installed OS on the eMMC device stays intact. Only the EFI boot entry is updated to boot from the SSD device.

Changes and New Features in 4.0.3

- BlueField-3 tuning update for power and performance

Changes and New Features in 4.0.2

- BlueField-3 power-capping and thermal-throttling
- Added Linux `fscck` to boot flow
- Log PCIe errors (to RShim log)

- Halt uncorrectable double-bit ECC error on DDR

Changes and New Features in 3.9.3

- Added support for live migration of VirtIO-net and VirtIO-blk VFs from one VM to another. Requires working with the new [vDPA driver](#).
- OS configuration – enabled tmpfs in `/tmp`

Changes and New Features in 3.9.2

- Added support for Arm host
- Enroll new NVIDIA certificates to DPU UEFI database

Warning

Important: User action required! See known issue [#3077361](#) for details.

Changes and New Features in 3.9.0

Note

This is the last release to offer GA support for first-generation NVIDIA® BlueField® DPUs.

- Added support for [NIC mode](#) of operation
- Added [password protection](#) to change boot parameters in GRUB menu
- Added IB support for DOCA runtime and dev environment
- Implemented RShim PF interrupts

- Virtio-net-controller is split to 2 processes for fast recovery after service restart
- Added support for [live virtio-net controller upgrade](#) instead of performing a full restart
- Expanded BlueField-2 PCIe bus number range to 254 (0-253)
- Added a new CAP field, `log_max_queue_depth` (value can be set to `2K`/`4K`), to indicate the maximal NVMe SQ and CQ sizes supported by firmware. This can be used by NVMe controllers or by non-NVMe drivers which do not rely on NVMe CAP field.
- Added ability for the RShim driver to still work when the host is in secure boot mode
- Added `bfb-info` command which provides the breakdown of the software components bundled in the BFB package
- Added support for [rate limiting VF groups](#)

Changes and New Features in 3.8.5

- PXE boot option is enabled automatically and is available for the ConnectX and OOB network interfaces
- Added Vendor Class option "BF2Client" in DHCP request for PXE boot to identify card
- Updated the "force PXE" functionality to continue to retry PXE boot entries until successful. A configuration called "boot override retry" has been added. With this configured, UEFI does not rebuild the boot entries after all boot options are attempted but loops through the PXE boot options until booting is successful. Once successful, the boot override entry configuration is disabled and would need to be reenabled for future boots.
- Added ability to change the CPU clock dynamically according to the temperature and other sensors of the DPU. If the power consumption reaches close to the maximum allowed, the software module decreases the CPU clock rate to ensure that the power consumption does not cross the system limit.

Note

This feature is relevant only for OPNs MBF2H516C-CESOT, MBF2M516C-EECOT, MBF2H516C-EESOT, and MBF2H516C-CECOT.

- Bug fixes

Changes and New Features in 3.8.0

- Added ability to perform warm reboot on BlueField-2 based devices
- Added support for DPU BMC with OpenBMC
- Added support for [NVIDIA Converged Accelerator](#) (900-21004-0030-000)

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 NON-INFRINGEMENT, 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 03/09/2025