



Known Issues

Ref #	Issue
41868	<p>Description: Disabling boot options applies the new setting after the first reboot of the BlueField Arm. However, during the second reboot, some boot option settings are unexpectedly re-enabled. This issue specifically affects UEFI boot options.</p> <p>Workaround: N/A</p> <p>Keyword: UEFI; UEFI boot options</p> <p>Reported in version: 4.9.1</p>
4157867	<p>Description: If the SSD is power cycled many times, the file system may be corrupted.</p> <p>Workaround: Re-install the BFB or power cycle the BlueField.</p> <p>Keyword: eMMC; file corruption</p> <p>Reported in version: 4.9.0</p>
4052874	<p>Description: <code>UefiSignatureOwner</code> field is not supported, if this field was populated with data, an exception occurs.</p> <p>Workaround: N/A</p> <p>Keyword: Redfish; UEFI</p> <p>Reported in version: 4.9.0</p>
4049034	<p>Description: On openEuler 22.03 SP3 and openEuler 20.03 SP1, it is not possible to do <code>yum update</code> after BFB installation.</p> <p>Workaround: Attempt the following:</p> <ol style="list-style-type: none"> To update DOCA only, disable ALL repos: <pre>dnf --disablerepo='*' --enablerepo='kubernetes,doca' -y update</pre> To update OS components, exclude <code>rdma-core</code> using one of the following options: <ul style="list-style-type: none"> Option 1 – Disable <code>rdma-core</code> from <code>dnf.conf</code> using <code>excludepkgs=rdma-core*oe2203sp3*</code> Option 2 – Disable <code>rdma-core</code> from <code>openEuler.repo</code> file under OS and everything using <code>exclude=rdma-core*</code>

Ref #	Issue
	<ul style="list-style-type: none"> Option 3 - <code>yum update -x rdma-core</code>
	Keyword: openEuler
	Reported in version: 4.9.0
4119852	Description: Redfish communication may be skipped upon reboot immediately following the installation of DOCA BFB-bundle.
	Workaround: Reboot (<code>SW_RESET</code>) the BlueField Arm to fix Redfish communication.
	Keyword: RF; installation; connection
	Reported in version: 4.9.0
3935124	Description: The number of rules users can offload depends on the amount of freed up memory on their system.
	Workaround: N/A
	Keyword: Rule; offload
	Reported in version: 4.9.0
4098782	Description: BlueField does not support the <code>fru print</code> format because it uses ipmitool version 1.8.18.
	Workaround: Read the information from the host using ipmitool version 1.8.19.
	Keyword: IPMI
	Reported in version: 4.9.0
4129715	Description: Compiling Rocky 9.2 may fail when using GCC with the <code>native</code> arch flag.
	Workaround: Upgrade to toolset 13 (gcc 13).
	Keyword: Linux; GCC
	Reported in version: 4.9.0
39576	Description: Failure to load kernel modules during BFB installation after an upgrade to 2.7.0/4.7.0 (or newer) release from 2.6.0/4.6.0 (or older) release, which results in a downgrade to 2.6.0/4.6.0 (or older) release from 2.7.0/4.7.0 (or newer) release.
	Workaround 1: Before the upgrade that follows a downgrade, delete the expired certificate.

Ref #	Issue
80	<p>Workaround 2: After the upgrade that follows a downgrade, re-install the BFB.</p> <p>Keyword: Software</p> <p>Reported in version: 4.8.0</p>
3859084	<p>Description: Running <code>bfm-install</code> in remote mode may fail if the remote server lacks an up-to-date <code>nc</code> software that supports TCP server functionality. This issue is commonly observed with servers such as BlueField BMC.</p> <p>Workaround: N/A</p> <p>Keyword: bfm-install; TCP</p> <p>Reported in version: 4.8.0</p>
3876636	<p>Description: RDMA devices are missing after installing BFB in NIC mode.</p> <p>Workaround: An additional BlueField reboot is required</p> <p>Keyword: RDMA; reboot; NIC mode</p> <p>Reported in version: 4.8.0</p>
3960613	<p>Description: When configuring NVMe-oF target offload on Arm, harmless CMA allocation errors appear at the kernel log.</p> <p>Workaround: Increase the CMA limit or cancel its use using the kernel's CMD line parameters:</p> <ul style="list-style-type: none"> • Add the parameter <code>cma=256M</code> to increase the CMA limit to 256MB; or • Add the parameter <code>cma=0</code> to disable the use of CMA <p>Keyword: NVMe-oF target offload; CMA</p> <p>Reported in version: 4.8.0</p>
396477	<p>Description: At times, BlueField may boot into grub shell after BFB installation.</p> <p>Workaround: Power-cycle the BlueField.</p> <p>Keyword: BFB installation; grub shell</p> <p>Reported in version: 4.8.0</p>

Ref #	Issue
3837649	<p>Description: Possible stack trace in <code>mlxbf_i2c</code> can occur intermittently when booting BFB.</p> <p>Workaround: Reboot the BlueField.</p> <p>Keyword: Boot</p> <p>Reported in version: 4.8.0</p>
3964584	<p>Description: RShim driver is not enabled by default after installation.</p> <p>Workaround: Enable RShim driver manually by running <code>systemctl enable rshim</code>. Then either reboot or run <code>systemctl start rshim</code> to start the RShim driver.</p> <p>Keyword: RShim driver</p> <p>Reported in version: 4.8.0</p>
3914629	<p>Description: When "Force PXE" is set right before installing a BFB image via BMC RShim, BlueField PXE boot fails to boot from NET-OOB-IPV4.</p> <p>Workaround: N/A</p> <p>Keyword: PXE boot</p> <p>Reported in version: 4.7.0</p>
3239320	<p>Description: Resetting hugepage size to 0 on Rocky Linux 8.6 using the <code>sysctl</code> tool fails.</p> <p>Workaround: Use the following command instead:</p> <pre data-bbox="219 1398 1463 1608">echo 0 > /sys/kernel/mm/hugepages/hugepages- <Size>/nr_hugepages</pre> <p>Keyword: Hugepage; sysctl</p> <p>Reported in version: 4.7.0</p>
3859	<p>Description: Reloading MLNX_OFED drivers with the command <code>/etc/init.d/openibd restart</code> fails when the NVMe driver is installed and in use.</p>

Ref #	Issue
13	Workaround: Reboot the machine to load all the MLNX_OFED drivers.
	Keyword: NVMe; driver
	Reported in version: 4.7.0
3748649	Description: With the numbering of CPUs in an 8-core configuration, the kernel is expected to assign virtual CPU ID numbers from 0-7, where N is the number of cores enabled. With CTyunOS, however, the numbering is unexpected.
	Workaround: N/A
	Keyword: CTyunOS; CPU numbering
	Reported in version: 4.7.0
3756748	Description: When performing BFB push repeatedly, BlueField-3 may in rare instances fail to boot with the message "PSC error -60" appearing in the RShim log sometimes.
	Workaround: Reset the card or repeat the operation (bfb push).
	Keyword: BFB Push; FW Reset
	Reported in version: 4.7.0
3665070	Description: Virtio-net controller fails to load if DPA_AUTHENTICATION is enabled.
	Workaround: N/A
	Keyword: Virtio-net; DPA
	Reported in version: 4.7.0
3862683	Description: Creating VFs and hotplug PFs in parallel can lead to controller crash.
	Workaround: Create VFs followed by hotplug PF or vice versa.
	Keyword: Virtio-net emulation
	Reported in version: 4.7.0
38440	Description: On CentOS 7.6 with kernel 4.19, bringing up OVS bridge interface causes call traces:
	WARNING: CPU: 5 PID: 14339 at kernel/rcu/tree_plugin.h:342 rcu_note_context_switch+0x48/0x538

Ref #	Issue
66	<p>Workaround: Do not bring UP OVS bridge interfaces.</p> <p>Keyword: CentOS; kernel; rcu_note_context_switch</p> <p>Reported in version: 4.7.0</p>
3844705	<p>Description: In OpenEuler 20.03, the Linux Kernel version 4.19.90 is affected by an issue that impacts the discard/trim functionality for the DPU eMMC device which may cause degraded performance of the DPU eMMC over time.</p> <p>Workaround: Upgrade to Linux Kernel version 5.10 or later.</p> <p>Keyword: eMMC discard; trim functionality</p> <p>Reported in version: 4.7.0</p>
387725	<p>Description: During BFB installation in NIC mode on BlueField-3, too much information is added into RShim log which fills it, causing the Linux installation progress log to not appear in the RShim log.</p> <pre>echo "DISPLAY_LEVEL 2" > /dev/rshim0/misc cat /dev/rshim0/misc</pre> <p>Workaround: Monitor the BlueField-3 Arm's UART console to check whether BFB installation has completed or not for NIC mode.</p> <pre>[13:58:39] INFO: Installation finished ... [14:01:53] INFO: Rebooting...</pre> <p>Keyword: NIC mode; BFB install</p> <p>Reported in version: 4.7.0</p>
385702	<p>Description: Trying to jump from a steering level in the hardware to a lower level using software steering is not supported on <code>rdma-core</code> lower than 48.x.</p> <p>Workaround: N/A</p> <p>Keyword: RDMA; SWS</p> <p>Reported in version: 4.7.0</p>

Ref #	Issue
	<p>Description: <code>mlxfwreset</code> could timeout on servers where the RShim driver is running and INTx is not supported. The following error message is printed:</p> <pre>BF reset flow encountered a failure due to a reset state error of negotiation timeout</pre> <p>Workaround: Set <code>PCIE_HAS_VFIO=0</code> and <code>PCIE_HAS_UIO=0</code> in <code>/etc/rshim.conf</code> and restart the RShim driver. Then re-run the <code>mlxfwreset</code> command.</p> <p>If host Linux kernel lockdown is enabled, then manually unbind the RShim driver before <code>mlxfwreset</code> and bind it back after <code>mlxfwreset</code>:</p> <pre>echo "DROP_MODE 1" > /dev/rshim0/misc mlxfwreset <arguments> echo "DROP_MODE 0" > /dev/rshim0/misc</pre> <p>Keyword: Timeout; <code>mlxfwreset</code>; INTx</p> <p>Reported in version: 4.7.0</p>
3 6 7 0 3 6 1	<p>Description: Rarely, the driver takes more than several minutes to load.</p> <p>Workaround: Re-run <code>/sbin/mlnx_bf_configure</code>.</p> <p>Keywords: Driver; boot</p> <p>Reported in version: 4.6.0</p>
3 7 4 6 8 6 6	<p>Description: The error message <code>IANA PEN registry open failed: No such file or directory</code> may appear when using ipmitool version 1.8.19-7. This message can be safely ignored.</p> <p>Workaround: N/A</p> <p>Keywords: IPMI; Debian</p> <p>Reported in version: 4.6.0</p>
3 7 5	<p>Description: UEFI synchronous exception is observed at address <code>0x479B7xxxx</code> where the UEFI module names are not printed. See the following example:</p>

Ref #	Issue
5143	<p>ERR[UEFI]: PC=0x479B78480 (B4000040 3900001F A94153F3 F94013F5) ERR[UEFI]: PC=0x479B78480 ERR[UEFI]: PC=0x479B7E684 ERR[UEFI]: PC=0x47A0E93F4 ERR[UEFI]: PC=0x47A0E9608</p> <p>Workaround: Run software reset or reinstall the BFB.</p> <p>Keywords: UEFI synchronous exception</p> <p>Reported in version: 4.6.0</p>
3772177	<p>Description: SSHing to the DPU with Debian 12 can print the following warning: <pre>-bash: warning: setlocale: LC_ALL: cannot change locale (en_US.UTF-8)</pre></p> <p>Workaround: Run: <pre>sudo dpkg-reconfigure locales</pre></p> <p>Keywords: Debian 12; locale; LC_ALL</p> <p>Reported in version: 4.6.0</p>
3704985	<p>Description: When the RShim driver is not running on the external host or when the <code>tmfifo_net0</code> interface is down on the DPU side, the following kernel warning may appear on the DPU side: <code>virtio_net virtio1 tmfifo_net0: TX timeout</code>.</p> <p>Workaround: N/A</p> <p>Keywords: RShim; log</p> <p>Reported in version: 4.6.0</p>
37675	<p>Description: On Debian 12, the first boot after BFB installation may fail with the following kernel panic:</p>

Ref #	Issue
80	<p data-bbox="277 352 1365 443">[end Kernel panic - not syncing: Attempted to kill init! exitcode=0x00000100]</p> <p data-bbox="217 527 1029 562">Workaround: Reset the DPU using the RShim interface:</p> <pre data-bbox="277 625 898 661">echo "SW_RESET 1" > /dev/rshim0/misc</pre> <p data-bbox="217 743 834 779">Keywords: Debian 12; Kernel panic; kill init</p> <p data-bbox="217 806 594 842">Reported in version: 4.6.0</p>
3771601	<p data-bbox="217 869 1354 953">Description: On Debian 12, <code>/etc/init.d/openibd restart</code> fails with the following error:</p> <pre data-bbox="277 1016 1273 1052">rmmod: ERROR: Module rdma_cm is in use by: nvme_rdma</pre> <p data-bbox="217 1136 480 1171">Workaround: Run:</p> <pre data-bbox="277 1234 1235 1270">modprobe -r nvme_rdma; /etc/init.d/openibd restart</pre> <p data-bbox="217 1352 850 1388">Keywords: Debian 12; openibd; nvme_rdma</p> <p data-bbox="217 1415 594 1451">Reported in version: 4.6.0</p>
3686053	<p data-bbox="217 1478 1463 1751">Description: BlueField-2 supports a total of 120GB of PCIe memory space. When the GPU is configured to be exposed to the BlueField, it requests 32GB of space for its BAR0. The Linux 5.15 kernel also attempts to reserve space for the total number of VFs, even if they are not enabled. By default, the A100 allows 20 VFs which each need 4GB of memory space. Because of PCIe memory alignment requirements and other small devices on the bus, this additional 80GB causes PCIe resource allocation to fail.</p> <p data-bbox="217 1772 1430 1856">Workaround: Add "pci=realloc=off" to the Linux command line. This will force Linux to accept the resource allocation done by UEFI and allow enumeration to succeed.</p> <p data-bbox="217 1877 667 1913">Keyword: VF; kernel; resources</p>

Ref #	Issue
	Reported in version: 4.6.0
3678069	<p>Description: If using DPUs with NVMe and mmcblk and configured to boot from mmcblk, users must create a <code>bf.cfg</code> file with <code>device=/dev/mmcblk0</code> before installing the <code>*.bfb</code>.</p> <p>Workaround: N/A</p> <p>Keyword: NVMe</p> <p>Reported in version: 4.6.0</p>
3747285	<p>Description: The ipmitool command to force PXE in BMC modifies both the IPMI and Redfish request settings. When Redfish is enabled in UEFI, Redfish takes priority, so all PXE boot entries are attempted and before regular boot continues.</p> <p>Workaround: Redfish must be disabled if IPMI force PXE retry behavior is expected.</p> <p>Keyword: PXE; retry; fail</p> <p>Reported in version: 4.6.0</p>
374529	<p>Description: When rebooting the DPU while the host side is running traffic over bond, TX timeout is likely to occur. This generates a TX timeout recovery flow that may conflict with host recovery attempts from the DPU reboot.</p> <p>Workaround: N/A</p> <p>Keyword: Bond; timeout</p> <p>Reported in version: 4.6.0</p>
37333713	<p>Description: CA certificates in the UEFI are stored in the database provided by the user. It is user responsibility to enroll the correct certificate. The user is the owner of the certificate and should make sure of its validity.</p> <p>Workaround: N/A</p> <p>Keyword: CA certificates; UEFI</p> <p>Reported in version: 4.6.0</p>
37337	<p>Description: CA certificates in the BMC are owned by the user who is required to enroll valid and correct certificates. If incorrect BMC CA certificates are enrolled, then DPU-BMC redfish communication will be invalid.</p> <p>Workaround: N/A</p> <p>Keyword: CA certificates; BMC</p>

Ref #	Issue
40	Reported in version: 4.6.0
366	Description: Running <code>systemctl restart openibd</code> on the DPU can result in openvswitch service crash.
66	Workaround: Run <code>/etc/init.d/openvswitch-switch start</code> .
57	Keyword: OVS fail; openibd
4	Reported in version: 4.6.0
3880194	<p>Description: <code>mlxbf-bootctl</code> command failed to install <code>default.bfb</code>.</p> <p>Workaround: The following are possible options –</p> <ul style="list-style-type: none"> • Boot the BFB file <code>doca_2.5.0_bsp_4.5.0_ubuntu_22.04-1.23-10.prod.bfb</code> to update the platform software • Download, compile, and install latest <code>mlxbf-bootctl</code> command from GitHub • Edit <code>default.bfb</code> by using the <code>mlx-mkbfb</code> command to incorporate the platform-specific images and filtering out unused images. Example for a BlueField-2 device: <pre style="border: 1px solid black; padding: 10px;"> \$ mlx-mkbfb -x default.bfb \$ mlx-mkbfb \ --bl2r-v1=dump-bl2r-v1 \ --bl2r-cert-v1=dump-bl2r-cert-v1 \ --bl2-v1=dump-bl2-v1 \ --bl2-cert-v1=dump-bl2-cert-v1 \ --bl31-v1=dump-bl31-v1 \ --bl31-cert-v1=dump-bl31-cert-v1 \ --bl31-key-cert-v1=dump-bl31-key-cert-v1 \ --bl33-v0=dump-bl33-v0 \ --bl33-cert-v1=dump-bl33-cert-v1 \ --bl33-key-cert-v1=dump-bl33-key-cert-v1 \ --boot-acpi-v0=dump-boot-acpi-v0 \ --boot-args-v0=dump-boot-args-v0 \ --boot-desc-v0=dump-boot-desc-v0 \ --boot-path-v0=dump-boot-path-v0 \ --ddr_ini-v1=dump-ddr_ini-v1 \ --ddr-cert-v1=dump-ddr-cert-v1 \ --ddr_ate_imem-v1=dump-ddr_ate_imem-v1 \ --ddr_ate_dmem-v1=dump-ddr_ate_dmem-v1 \ --snps_images-v1=dump-snps_images-v1 \ </pre>

Ref #	Issue
	<pre data-bbox="305 296 878 365">--trusted-key-cert-v1=dump-trusted-key-cert-v1 \ default_min.bfb</pre> <p data-bbox="217 432 651 468">Keywords: Software; upgrade</p> <p data-bbox="217 493 613 529">Discovered in version: 4.5.1</p>
3204153	<p data-bbox="217 554 1458 632">Description: On BlueField-2, the OOB may not get an IP address due to the interface being down.</p> <p data-bbox="217 657 1076 735">Workaround: restart auto-negotiation using the command <code>ethtool -r oob_net0</code>.</p> <p data-bbox="217 766 475 802">Keyword: OOB; IP</p> <p data-bbox="217 827 594 863">Reported in version: 4.5.0</p>
3601491	<p data-bbox="217 890 1433 968">Description: Symmetric pause must be enabled in the DHCP server for the OOB to be able to reliably get an IP address assigned.</p> <p data-bbox="217 999 472 1035">Workaround: N/A</p> <p data-bbox="217 1060 475 1096">Keyword: OOB; IP</p> <p data-bbox="217 1127 594 1163">Reported in version: 4.5.0</p>
367330	<p data-bbox="217 1190 1398 1268">Description: On Debian 12, Arm ports remain in Legacy mode after multiple Arm reboot iterations. The following error message appears in <code>/var/log/syslog</code>:</p> <pre data-bbox="217 1276 1463 1486">mInx_bf_configure[2601]: ERR: Failed to configure switchdev mode for 0000:03:00.0 after 61 retries</pre> <p data-bbox="217 1512 483 1547">Workaround: Run:</p> <pre data-bbox="217 1556 1463 1759">\$ echo SET_MODE_RETRY_NUM=300 >> /etc/mellanox/mInx-bf.conf \$ reboot</pre> <p data-bbox="217 1778 545 1814">Keyword: Debian; Arm</p> <p data-bbox="217 1839 594 1875">Reported in version: 4.5.0</p>

Ref #	Issue
3695543	<p>Description: PXE boot may fail after a firmware upgrade from 32.36.xxxx, 32.37.xxxx, to 32.38.xxxx and above.</p> <p>Workaround: Create <code>/etc/bf.cfg</code> with the following lines, then run <code>bfcfg</code> to recreate the PXE boot entries:</p> <pre data-bbox="224 495 1463 951"> B00T0=DISK B00T1=NET-NIC_P0-IPV4 B00T2=NET-NIC_P0-IPV6 B00T3=NET-NIC_P1-IPV4 B00T4=NET-NIC_P1-IPV6 B00T5=NET-00B-IPV4 B00T6=NET-00B-IPV6 </pre> <p>Keyword: MAC allocation; PXE boot</p> <p>Reported in version: 4.5.0</p>
3647476	<p>Description: Debian 12 OS does not support CT tunnel offload.</p> <p>Workaround: Recompile the kernel with <code>CONFIG_NET_TC_SKB_EXT</code> set.</p> <p>Keyword: Connection tracking; Linux</p> <p>Reported in version: 4.5.0</p>
3007696	<p>Description: When configuring a static IP address for <code>tmfifo_net0</code> interface in <code>/etc/network/interfaces</code>, the IP address is lost after restarting the RShim driver on Debian Linux.</p> <p>Workaround: Use netplan configuration. For example</p> <pre data-bbox="224 1598 1463 1944"> # cat /etc/netplan/tmfifo_net0.yaml network: version: 2 renderer: networkd ethernets: tmfifo_net0: </pre>

Ref #	Issue
	<pre>addresses : - 192.168.100.1/30 dhcp4: false</pre> <p>Then run "netplan apply".</p> <p>Keyword: IP address; tmfifo_net0; host</p> <p>Reported in version: 4.5.0</p>
3670628	<p>Description: When NIC subsystem is in recovery mode, the interface towards to NVMe is not accessible. Thus, the SSD boot device would not be available.</p> <p>Workaround: The admin must configure the Arm subsystem boot device to boot from the eMMC, for example.</p> <p>Keyword: mlxfwreset; RShim</p> <p>Reported in version: 4.5.0</p>
3702393	<p>Description: On rare occasions, the boot process part of SWRESET (via RShim) or FWRESET (via mlxfwreset) may result in a device hanging on the boot flow or cause the host server to reboot.</p> <p>Workaround: Perform graceful shutdown and then a power cycle.</p> <p>Keyword: mlxfwreset; RShim</p> <p>Reported in version: 4.5.0</p>
3665724	<p>Description: If the UEFI password is an empty string (""), then it cannot be changed via Redfish.</p> <p>Workaround: UEFI; password; Redfish</p> <p>Keyword: UEFI; password; Redfish</p> <p>Reported in version: 4.5.0</p>
3671185	<p>Description: XFRM rules must be deleted before driver restart or warm reboot are performed.</p> <p>Workaround: N/A</p> <p>Keyword: IPsec</p> <p>Reported in version: 4.5.0</p>

Ref #	Issue
36610	<p>Description: Installing BFB using <code>bfb-install</code> when <code>mlxconfig</code> <code>PF_TOTAL_SF</code> >1700, triggers server reboot immediately.</p> <p>Workaround: Change <code>PF_TOTAL_SF</code> to 0, perform graceful shutdown, then power cycle, and then install the BFB.</p> <p>Keyword: SF; PF_TOTAL_SF; BFB installation</p> <p>Reported in version: 4.2.2</p>
36054	<p>Description: Following a system power cycle, both the DPU and BMC boot independently which may lead to the DPU's UEFI boot process to complete before the BMC's. As a result, when attempting to establish Redfish communication, the BMC may not yet be prepared to respond.</p> <p>Workaround: Wait until the BMC is done booting before issuing a reset command to the DPU.</p> <p>Keyword: Power cycle; Redfish; boot</p> <p>Reported in version: 4.2.1</p>
3602044	<p>Description: When the public key is deleted while Redfish is enabled, UEFI secure boot is disabled and UEFI reverts to Setup Mode (i.e., the <code>SecureBootEnable</code> Redfish property is reset to <code>false</code>). If later, the public key is re-enrolled, the platform does not implement UEFI secure boot until the <code>SecureBootEnable</code> Redfish property is explicitly changed to <code>true</code>.</p> <p>Workaround: Set <code>SecureBootEnable</code> to true using the Redfish API.</p> <p>Keyword: Redfish; UEFI secure boot</p> <p>Reported in version: 4.2.1</p>
3592080	<p>Description: When using UEK8 on the host in DPU mode, creating a VF on the host consumes about 100MB memory on the DPU.</p> <p>Workaround: N/A</p> <p>Keyword: UEK; VF</p> <p>Reported in version: 4.2.1</p>
356	<p>Description: Downgrading BSP software from 4.2.0 fails if UEFI secure boot is enabled.</p>

Ref #	Issue
8341	<p>Workaround: Disable UEFI secure boot before downgrading.</p> <p>Keyword: Software; downgrade</p> <p>Reported in version: 4.2.0</p>
3566042	<p>Description: Virtio hotplug is not supported in GPU-HOST mode on the NVIDIA Converged Accelerator.</p> <p>Workaround: N/A</p> <p>Keyword: Virtio; Converged Accelerator</p> <p>Reported in version: 4.2.0</p>
354674	<p>Description: PXE boot over ConnectX interface might not work due to an invalid MAC address in the UEFI boot entry.</p> <p>Workaround: On the DPU, create <code>/etc/bf.cfg</code> file with the relevant PXE boot entries, then run the command <code>bfcfg</code>.</p> <p>Keyword: PXE; boot; MAC</p> <p>Reported in version: 4.2.0</p>
3546202	<p>Description: After rebooting a BlueField-3 DPU running Rocky Linux 8.6 BFB, the kernel log shows the following error:</p> <pre>[3.787135] mlxbf_gige MLNXBF17:00: Error getting PHY irq. Use polling instead</pre> <p>This message indicates that the Ethernet driver will function normally in all aspects, except that PHY polling is enabled.</p> <p>Workaround: N/A</p> <p>Keyword: Linux; PHY; kernel</p> <p>Reported in version: 4.2.0</p>
33064	<p>Description: When performing longevity tests (e.g., <code>mlxfwreset</code>, DPU reboot, burning of new BFBs), a host running an Intel CPU may observe errors related to "CPU 0: Machine Check Exception".</p> <p>Workaround: Add <code>intel_idle.max_cstate=1</code> entry to the kernel command line.</p>

Ref #	Issue
89	<p>Keyword: Longevity; mlxfwreset; DPU reboot</p> <p>Reported in version: 4.2.0</p>
3538486	<p>Description: When removing LAG configuration from the DPU, a kernel warning for <code>uverbs_destroy_ufile_hw</code> is observed if virtio-net-controller is still running.</p> <p>Workaround: Stop virtio-net-controller service before cleaning up bond configuration.</p> <p>Keyword: Virtio-net; LAG</p> <p>Reported in version: 4.2.0</p>
3462630	<p>Description: When trying to perform a PXE installation when UEFI Secure Boot is enabled, the following error messages may be observed:</p> <pre data-bbox="219 842 1463 1052"> error: shim_lock protocol not found. error: you need to load the kernel first. </pre> <p>Workaround: Download a Grub EFI binary from the Ubuntu website. For further information on Ubuntu UEFI Secure Boot PXE Boot, please visit Ubuntu's official website.</p> <p>Keyword: PXE; UEFI Secure Boot</p> <p>Reported in version: 4.0.2</p>
3412847	<p>Description: Socket-Direct is currently not supported on BlueField-3 devices.</p> <p>Workaround: N/A</p> <p>Keyword: Socket-Direct; support</p> <p>Reported in version: 4.0.2</p>
3448841	<p>Description: While running CentOS 8.2, switchdev Ethernet DPU runs in "shared" RDMA net namespace mode instead of "exclusive".</p> <p>Workaround: Use <code>ib_core</code> module parameter <code>netns_mode=0</code>. For example:</p> <pre data-bbox="219 1787 1463 1934"> echo "options ib_core netns_mode=0" >> /etc/modprobe.d/mlnx- </pre>

Ref #	Issue
	<p>bf.conf</p> <p>Keywords: RDMA; isolation; Net NS</p> <p>Reported in version: 4.0.2</p>
341338	<p>Description: Using <code>m1nx-sf</code> script, creating and deleting an SF with same ID number in a stressful manner may cause the setup to hang due to a race between create and delete commands.</p> <p>Workaround: N/A</p> <p>Keywords: Hang; <code>m1nx-sf</code></p> <p>Reported in version: 4.0.2</p>
3452740	<p>Description: Ovs-pki is not working due to two versions of OpenSSL being installed, causing the PKA engine to not load properly.</p> <p>Workaround: N/A</p> <p>Keywords: PKA; OpenSSL</p> <p>Reported in version: 4.0.2</p>
3273435	<p>Description: Changing the mode of operation between NIC and DPU modes results in different capabilities for the host driver which might cause unexpected behavior.</p> <p>Workaround: Reload the host driver or reboot the host.</p> <p>Keywords: Modes of operation; driver</p> <p>Reported in version: 4.0.2</p>
270803	<p>Description: When an NVMe controller, SoC management controller, and DMA controller are configured, the maximum number of VFs is limited to 124.</p> <p>Workaround: N/A</p> <p>Keywords: VF; limitation</p> <p>Reported in version: 4.0.2</p>
32644	<p>Description: When trying to change boot order using efibootmgr, BlueField fails to attempt PXE boot from <code>p0</code> even though efibootmgr returns a successful result.</p> <p>Workaround: Drop into the UEFI menu and regenerate all the EFI entries.</p>

Ref #	Issue
224	Keywords: PXE; efibootmgr
	Reported in version: 3.9.3.1
3188	Description: An Arm firmware update to the same version that is installed will fail and is not supported.
	Workaround: N/A
41	Keywords: Arm; firmware; update
5	Reported in version: 3.9.2
N/A	Description: The <code>BootOptionEnabled</code> attribute changes back to true after DPU-force reset.
	Workaround: N/A
	Keywords: Redfish; <code>BootOptionEnabled</code>
	Reported in version: 3.9.2
3012	Description: The command <code>ethtool -I --show-fec</code> is not supported by the DPU with kernel 5.4.
	Workaround: N/A
18	Keywords: Kernel; show-fec
2	Reported in version: 3.9.0
285	Description: After disabling SR-IOV VF on a virtio device, removing virtio-net/PCIe driver from guest OS may render the virtio controller unusable .
559	Workaround: Restart the virtio-net controller to recover it. To avoid this issue, monitor the log from controller and make sure VF resources are destroyed before unloading virtio-net/PCIe drivers.
86	Keywords: Virtio-net; VF
	Reported in version: 3.9.0
28634	Description: SA limit by packet count (hard and soft) are supported only on traffic originated from the ECPF. Trying to configure them on VF traffic removes the SA when hard limit is hit. However, traffic could still pass as plain text due to the tunnel offload used in such configuration.
	Workaround: N/A

Ref #	Issue
56	Keywords: ASAP2; IPsec Full Offload
	Reported in version: 3.9.0
2982184	Description: When multiple BlueField resets are issued within 10 seconds of each other, EEPROM error messages are displayed on the console and, as a result, the BlueField may not boot from the eMMC and may halt at the UEFI menu.
	Workaround: Power-cycle the BlueField to fix the EEPROM issue. Manual recovery of the boot options and/or SW installation may be needed.
	Keywords: Reset; EEPROM
	Reported in version: 3.9.0
2853408	Description: Some pre-OS environments may fail when sensing a hot plug operation during their boot stage.
	Workaround: Run " <code>mlxconfig -d <mst dev> set PF_LOG_BAR_SIZE=0</code> ".
	Keywords: BIOS; hot-plug; Virtio-net
	Reported in version: 3.9.0
2934833	Description: Running I/O traffic and toggling both physical ports status in a stressful manner on the receiving-end machine may cause traffic loss.
	Workaround: N/A
	Keywords: MLNX_OFED; RDMA; port toggle
	Reported in version: 3.8.5
2911425	Description: ProLiant DL385 Gen10 Plus server with BIOS version 1.3 hangs when large number of SFs (<code>PF_TOTAL_SF=252</code>) are configured.
	Workaround: Update the BIOS version to 2.4 which should correctly detect the PCIe device with the bigger BAR size.
	Keywords: Scalable functions; BIOS
	Reported in version: 3.8.5
N/A	Description: Only QP queues are supported for GGA accelerators from this version onward.
	Workaround: N/A
	Keywords: Firmware; SQ; QP

R ef #	Issue
	Reported in version: 3.8.0
2 8 4 6	Description: Setting <code>VHCA_TRUST_LEVEL</code> does not work when there are active SFs or VFs.
1 0	Workaround: N/A
8	Keywords: Firmware; SF; VF
	Reported in version: 3.8.0
2 7 5 0	Description: Some devlink commands are only supported by mlnx devlink (<code>/opt/mellanox/iproute2/sbin/devlink</code>). The default devlink from the OS may produce failure (e.g., <code>devlink port show -j</code>).
4	Workaround: N/A
9	Keywords: Devlink
9	Reported in version: 3.7.1
2 7	Description: Kernel upgrade is not currently supported on BlueField as there are out of tree kernel modules (e.g., ConnectX drivers that will stop working after kernel upgrade).
3 0 1 5	Workaround: Kernel can be upgraded if there is a matching DOCA repository that includes all the drivers compiled with the new kernel or as a part of the new BFB package.
7	Keywords: Kernel; upgrade
	Reported in version: 3.7.0
2 7 0	Description: Call traces are seen on the host when recreating VFs before the controller side finishes the deletion procedure.
6	Workaround: N/A
7 1	Keywords: Virtio-net controller
0	Reported in version: 3.7.0
2 6 8 5	Description: 3rd party (netkvm.sys) Virtio-net drivers for Windows do not support SR-IOV.
4	Workaround: N/A
	Keywords: Virtio-net; SR-IOV; WinOF-2

Ref #	Issue
78	Reported in version: 3.7.0
268	<p>Description: Once the contiguous memory pool, a limited resource, is exhausted, fallback allocation to other methods occurs. This process triggers <code>cma_alloc</code> failures in the dmesg log.</p>
45	Workaround: N/A
0	Keywords: Log; cma_alloc; memory
1	Reported in version: 3.7.0
259	<p>Description: ibdev2netdev tool is not supported for PCIe PF operating in switchdev mode or on SFs.</p>
0	Workaround: N/A
01	Keywords: ibdev2netdev
6	Reported in version: 3.6.0.11699
259	<p>Description: A "double free" error is seen when using the "curl" utility. This error is from libcrypto.so library which is part of the OpenSSL package. This happens only when OpenSSL is configured to use a dynamic engine (e.g. Bluefield PKA engine).</p>
0016	<p>Workaround: Set <code>OPENSSL_CONF=/etc/ssl/openssl.cnf.orig</code> before using the curl utility. For example:</p>
	<pre># OPENSSL_CONF=/etc/ssl/openssl.cnf.orig curl -0 https://tpo.pe/pathogen.vim</pre>
	<p>Note OPENSSL_CONF is aimed at using a custom config file for applications. In this case, it is used to point to a config file where dynamic engine (PKA engine) is not enabled.</p>
	Keywords: OpenSSL; curl

Ref #	Issue
	Reported in version: 3.6.0.11699
2407897	<p>Description: The host may crash when the number of PCIe devices overflows the PCIe device address. According to the PCIe spec, the device address space is 8 bits in total—device (5 bits) and function (3 bits)—which means that the total number of devices cannot be more than 256.</p> <p>The second PF maximum number of VFs is limited by the total number of additional PCIe devices that precedes it. By default, the preceding PCIe devices are 2 PFs + RShim DMA + 127 VFs of the first PF. This means that the maximum valid number of VFs for the second port will be 126.</p> <p>Workaround: Use the maximum allowed VFs on the 2nd PCIe PF of BlueField instead of the maximum of 127 VFs.</p> <p>Keywords: Emulated devices; VirtIO-net; VirtIO-blk; VFs; RShim</p> <p>Reported in version: 3.6.0.11699</p>
2445289	<p>Description: If secure boot is enabled, MFT cannot be installed on the BlueField DPU independently from BlueField drivers (MLNX_OFED).</p> <p>Workaround: N/A</p> <p>Keywords: MFT; secure boot</p> <p>Reported in version: 3.5.1.11601</p>
237021	<p>Description: Executing <code>sudo poweroff</code> on the Arm side causes the system to hang.</p> <p>Workaround: Perform graceful shutdown, then reboot your BlueField device or power cycle the server.</p> <p>Keywords: Hang; reboot</p> <p>Reported in version: 3.5.0.11563</p>
2350132	<p>Description: Boot process hangs at BIOS (version 1.2.11) stage when power cycling a server (model Dell PowerEdge R7525) after configuring "PCI_SWITCH_EMULATION_NUM_PORT" > 27.</p> <p>Workaround: N/A</p> <p>Keywords: Server; hang; power cycle</p> <p>Reported in version: 3.5.0.11563</p>

Ref #	Issue
2581408	<p>Description: On a BlueField device operating in Embedded CPU mode, PXE driver will fail to boot if the Arm side is not fully loaded and the OVS bridge is not configured.</p> <p>Workaround: Run warm reboot on the host side and boot again via the device when Arm is up and the OVS bridge is configured.</p> <p>Keywords: Embedded CPU; PXE; UEFI; Arm</p> <p>Reported in version: 2.5.0.11176</p>
1859322	<p>Description: On some setups, DPU does not power on following server cold boot when UART cable is attached to the same server.</p> <p>Workaround: As long as the RShim driver is loaded on the server and the RShim interface is visible, the RShim driver will detect this and auto-reset the card into normal state.</p> <p>Keywords: DPU; Arm; Cold Boot</p> <p>Reported in version: 2.4.0.11082</p>
189921	<p>Description: Driver restart fails when SNAP service is running.</p> <p>Workaround: Stop the SNAP services nvme_sf and nvme_snap@nvme0, then restart the driver. After the driver loads restart the services.</p> <p>Keywords: SNAP</p> <p>Reported in version: 2.2.0.11000</p>
19618	<p>Description: Defining namespaces with certain Micron disks (Micron_9300_MTFDHAL3T8TDP) using consecutive attach-ns commands can cause errors.</p> <p>Workaround: Add delay between attach-ns commands.</p> <p>Keywords: Micron; disk; namespace; attach-ns</p> <p>Reported in version: 2.2.0.11000</p>

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. PDF Generated on 03/09/2025