



Known Issues

Ref #	Issue
415	Description: If the SSD is power cycled many times, the file system may be corrupted.
7	Workaround: Soft reset the BlueField.
86	Keyword: eMMC; file corruption
7	Reported in version: 4.9.0
405	Description: <code>UefiSignatureOwner</code> field is not supported, if this field was populated with data, an exception occurs.
2	Workaround: N/A
87	Keyword: Redfish; UEFI
4	Reported in version: 4.9.0
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> Option 3 – <code>yum update -x rdma-core</code> <p>Keyword: openEuler</p> <p>Reported in version: 4.9.0</p>
41	Description: Redfish communication may be skipped upon reboot immediately following the installation of DOCA BFB-bundle.

Ref #	Issue
19852	<p>Workaround: Reboot (<code>SW_RESET</code>) the BlueField Arm to fix Redfish communication.</p> <p>Keyword: RF; installation; connection</p> <p>Reported in version: 4.9.0</p>
4128189	<p>Description: When updating the BFB image on a BlueField-2 system via HTTP/HTTPS using Redfish, the operation may fail if the system is overloaded.</p> <p>Workaround: Try using SCP instead.</p> <p>Keyword: Software; update; Redfish</p> <p>Reported in version: 4.9.0</p>
3935124	<p>Description: The number of rules users can offload depends on the amount of freed up memory on their system.</p> <p>Workaround: N/A</p> <p>Keyword: Rule; offload</p> <p>Reported in version: 4.9.0</p>
4098782	<p>Description: BlueField does not support the <code>fru print</code> format because it uses ipmitool version 1.8.18.</p> <p>Workaround: Read the information from the host using ipmitool version 1.8.19.</p> <p>Keyword: IPMI</p> <p>Reported in version: 4.9.0</p>
412715	<p>Description: Compiling Rocky 9.2 may fail when using GCC with the <code>native</code> arch flag.</p> <p>Workaround: Upgrade to toolset 13 (gcc 13).</p> <p>Keyword: Linux; GCC</p> <p>Reported in version: 4.9.0</p>
39576	<p>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.</p> <p>Workaround 1: Before the upgrade that follows a downgrade, delete the expired certificate.</p>

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>
32393320	<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="224 1100 1463 1310">echo 0 > /sys/kernel/mm/hugepages/hugepages- <Size>/nr_hugepages</pre> <p>Keyword: Hugepage; sysctl</p> <p>Reported in version: 4.7.0</p>
3859113	<p>Description: Reloading <code>MLNX_OFED</code> drivers with the command <code>/etc/init.d/openibd restart</code> fails when the NVMe driver is installed and in use.</p> <p>Workaround: Reboot the machine to load all the <code>MLNX_OFED</code> drivers.</p> <p>Keyword: NVMe; driver</p> <p>Reported in version: 4.7.0</p>
3748	<p>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.</p>

Ref #	Issue
649	<p>Workaround: N/A</p> <p>Keyword: CTyunOS; CPU numbering</p> <p>Reported in version: 4.7.0</p>
3756748	<p>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.</p> <p>Workaround: Reset the card or repeat the operation (bfb push).</p> <p>Keyword: BFB Push; FW Reset</p> <p>Reported in version: 4.7.0</p>
3665070	<p>Description: Virtio-net controller fails to load if DPA_AUTHENTICATION is enabled.</p> <p>Workaround: N/A</p> <p>Keyword: Virtio-net; DPA</p> <p>Reported in version: 4.7.0</p>
3862683	<p>Description: Creating VFs and hotplug PFs in parallel can lead to controller crash.</p> <p>Workaround: Create VFs followed by hotplug PF or vice versa.</p> <p>Keyword: Virtio-net emulation</p> <p>Reported in version: 4.7.0</p>
3844066	<p>Description: On CentOS 7.6 with kernel 4.19, bringing up OVS bridge interface causes call traces:</p> <pre>WARNING: CPU: 5 PID: 14339 at kernel/rcu/tree_plugin.h:342 rcu_note_context_switch+0x48/0x538</pre> <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>
3844	<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>

Ref #	Issue
4705	Workaround: Upgrade to Linux Kernel version 5.10 or later.
	Keyword: eMMC discard; trim functionality
	Reported in version: 4.7.0
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 data-bbox="219 604 1463 814"> echo "DISPLAY_LEVEL 2" > /dev/rshim0/misc cat /dev/rshim0/misc </pre>
387725	<p>Workaround: Monitor the BlueField-3 Arm's UART console to check whether BFB installation has completed or not for NIC mode.</p> <pre data-bbox="219 913 1463 1171"> [13:58:39] INFO: Installation finished ... [14:01:53] INFO: Rebooting... </pre>
	Keyword: NIC mode; BFB install
	Reported in version: 4.7.0
3855702	<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>
37433879	<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 data-bbox="219 1707 1463 1801"> BF reset flow encountered a failure due to a reset state error of negotiation timeout </pre>

Ref #	Issue
	<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; mlxfwreset; INTx</p> <p>Reported in version: 4.7.0</p>
3670361	<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>
374666	<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>
375143	<p>Description: UEFI synchronous exception is observed at address 0x479B7xxxx where the UEFI module names are not printed. See the following example:</p> <pre>ERR[UEFI]: PC=0x479B78480(B4000040 3900001F A94153F3 F94013F5) ERR[UEFI]: PC=0x479B78480 ERR[UEFI]: PC=0x479B7E684 ERR[UEFI]: PC=0x47A0E93F4</pre>

Ref #	Issue
	<p data-bbox="219 296 1463 394">ERR[UEFI] : PC=0x47A0E9608</p> <p data-bbox="219 415 1463 464">Workaround: Run software reset or reinstall the BFB.</p> <p data-bbox="219 474 1463 522">Keywords: UEFI synchronous exception</p> <p data-bbox="219 533 1463 581">Reported in version: 4.6.0</p>
3772177	<p data-bbox="219 604 1463 779">Description: SSHing to the DPU with Debian 12 can print the following warning: -bash: warning: setlocale: LC_ALL: cannot change locale (en_US.UTF-8)</p> <p data-bbox="219 789 1463 837">Workaround: Run:</p> <p data-bbox="219 848 1463 989">sudo dpkg-reconfigure locales</p> <p data-bbox="219 1010 1463 1058">Keywords: Debian 12; locale; LC_ALL</p> <p data-bbox="219 1068 1463 1117">Reported in version: 4.6.0</p>
3704985	<p data-bbox="219 1140 1463 1281">Description: When the RShim driver is not running on the external host or when the tmfifo_net0 interface is down on the DPU side, the following kernel warning may appear on the DPU side: virtio_net virtio1 tmfifo_net0: TX timeout.</p> <p data-bbox="219 1291 1463 1339">Workaround: N/A</p> <p data-bbox="219 1350 1463 1398">Keywords: RShim; log</p> <p data-bbox="219 1409 1463 1457">Reported in version: 4.6.0</p>
37667580	<p data-bbox="219 1486 1463 1556">Description: On Debian 12, the first boot after BFB installation may fail with the following kernel panic:</p> <p data-bbox="219 1566 1463 1766">[end Kernel panic - not syncing: Attempted to kill init! exitcode=0x00000100]</p> <p data-bbox="219 1787 1463 1927">Workaround: Reset the DPU using the RShim interface:</p>

Ref #	Issue
	<pre data-bbox="212 296 1469 457">echo "SW_RESET 1" > /dev/rshim0/misc</pre> <p data-bbox="212 464 1469 520">Keywords: Debian 12; Kernel panic; kill init</p> <p data-bbox="212 527 1469 583">Reported in version: 4.6.0</p>
3771601	<p data-bbox="212 590 1469 688">Description: On Debian 12, <code>/etc/init.d/openibd restart</code> fails with the following error:</p> <pre data-bbox="212 695 1469 856">rmmod: ERROR: Module rdma_cm is in use by: nvme_rdma</pre> <p data-bbox="212 863 1469 919">Workaround: Run:</p> <pre data-bbox="212 926 1469 1066">modprobe -r nvme_rdma; /etc/init.d/openibd restart</pre> <p data-bbox="212 1073 1469 1129">Keywords: Debian 12; openibd; nvme_rdma</p> <p data-bbox="212 1136 1469 1192">Reported in version: 4.6.0</p>
3686053	<p data-bbox="212 1199 1469 1493">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="212 1499 1469 1598">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="212 1604 1469 1661">Keyword: VF; kernel; resources</p> <p data-bbox="212 1667 1469 1717">Reported in version: 4.6.0</p>
36780	<p data-bbox="212 1724 1469 1885">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 data-bbox="212 1892 1469 1934">Workaround: N/A</p>

R ef #	Issue
6 9	Keyword: NVMe
	Reported in version: 4.6.0
3 7 4	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.
7 2	Workaround: Redfish must be disabled if IPMI force PXE retry behavior is expected.
8	Keyword: PXE; retry; fail
5	Reported in version: 4.6.0
3 7 4	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.
5	Workaround: N/A
5 2	Keyword: Bond; timeout
9	Reported in version: 4.6.0
3 7 3	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.
3 7	Workaround: N/A
1	Keyword: CA certificates; UEFI
3	Reported in version: 4.6.0
3 7 3	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.
3 7	Workaround: N/A
4	Keyword: CA certificates; BMC
0	Reported in version: 4.6.0
3 6 6	Description: Running <code>systemctl restart openibd</code> on the DPU can result in openvswitch service crash.
6 5	Workaround: Run <code>/etc/init.d/openvswitch-switch start</code> .

Ref #	Issue
74	<p>Keyword: OVS fail; openibd</p> <p>Reported in version: 4.6.0</p>
3204153	<p>Description: On BlueField-2, the OOB may not get an IP address due to the interface being down.</p> <p>Workaround: restart auto-negotiation using the command <code>ethtool -r oob_net0</code>.</p> <p>Keyword: OOB; IP</p> <p>Reported in version: 4.5.0</p>
3601491	<p>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>Workaround: N/A</p> <p>Keyword: OOB; IP</p> <p>Reported in version: 4.5.0</p>
367330	<p>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>mInx_bf_configure[2601]: ERR: Failed to configure switchdev mode for 0000:03:00.0 after 61 retries</pre> <p>Workaround: Run:</p> <pre>\$ echo SET_MODE_RETRY_NUM=300 >> /etc/mellanox/mlnx-bf.conf \$ reboot</pre> <p>Keyword: Debian; Arm</p> <p>Reported in version: 4.5.0</p>
3695	<p>Description: PXE boot may fail after a firmware upgrade from 32.36.xxxx, 32.37.xxxx, to 32.38.xxxx and above.</p>

Ref #	Issue
543	<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="215 394 1466 856"> 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="215 1493 1466 1942"> # cat /etc/netplan/tmfifo_net0.yaml network: version: 2 renderer: networkd ethernets: tmfifo_net0: addresses: - 192.168.100.1/30 </pre>

Ref #	Issue
	<pre>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>
3702933	<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>
366	<p>Description: Installing BFB using <code>bfb-install</code> when <code>mlxconfig</code> <code>PF_TOTAL_SF</code> >1700, triggers server reboot immediately.</p>

Ref #	Issue
661	Workaround: Change <code>PF_TOTAL_SF</code> to 0, perform graceful shutdown, then power cycle, and then install the BFB.
60	Keyword: SF; PF_TOTAL_SF; BFB installation
	Reported in version: 4.2.2
360	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.
525	Workaround: Wait until the BMC is done booting before issuing a reset command to the DPU.
4	Keyword: Power cycle; Redfish; boot
	Reported in version: 4.2.1
360	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> .
4	Workaround: Set <code>SecureBootEnable</code> to true using the Redfish API.
4	Keyword: Redfish; UEFI secure boot
	Reported in version: 4.2.1
359	Description: When using UEK8 on the host in DPU mode, creating a VF on the host consumes about 100MB memory on the DPU.
2	Workaround: N/A
0	Keyword: UEK; VF
8	Reported in version: 4.2.1
356	Description: Downgrading BSP software from 4.2.0 fails if UEFI secure boot is enabled.
8	Workaround: Disable UEFI secure boot before downgrading.
3	Keyword: Software; downgrade

Ref #	Issue
41	Reported in version: 4.2.0
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>
3546474	<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>
3306489	<p>Description: When performing longevity tests (e.g., mlxfwreset, 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> <p>Keyword: Longevity; mlxfwreset; DPU reboot</p> <p>Reported in version: 4.2.0</p>

Ref #	Issue
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="219 716 1463 926">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 1661 1463 1871">echo "options ib_core netns_mode=0" >> /etc/modprobe.d/mlnx-bf.conf</pre> <p>Keywords: RDMA; isolation; Net NS</p>

Ref #	Issue
	Reported in version: 4.0.2
341	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.
39	Workaround: N/A
38	Keywords: Hang; <code>m1nx-sf</code>
	Reported in version: 4.0.2
345	Description: Ovs-pki is not working due to two versions of OpenSSL being installed, causing the PKA engine to not load properly.
72	Workaround: N/A
40	Keywords: PKA; OpenSSL
	Reported in version: 4.0.2
327	Description: Changing the mode of operation between NIC and DPU modes results in different capabilities for the host driver which might cause unexpected behavior.
34	Workaround: Reload the host driver or reboot the host.
43	Keywords: Modes of operation; driver
5	Reported in version: 4.0.2
270	Description: When an NVMe controller, SoC management controller, and DMA controller are configured, the maximum number of VFs is limited to 124.
68	Workaround: N/A
80	Keywords: VF; limitation
3	Reported in version: 4.0.2
326	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.
42	Workaround: Drop into the UEFI menu and regenerate all the EFI entries.
22	Keywords: PXE; efibootmgr
4	Reported in version: 3.9.3.1

Ref #	Issue
318415	<p>Description: An Arm firmware update to the same version that is installed will fail and is not supported.</p> <p>Workaround: N/A</p> <p>Keywords: Arm; firmware; update</p> <p>Reported in version: 3.9.2</p>
N/A	<p>Description: The <code>BootOptionEnabled</code> attribute changes back to true after DPU-force reset.</p> <p>Workaround: N/A</p> <p>Keywords: Redfish; <code>BootOptionEnabled</code></p> <p>Reported in version: 3.9.2</p>
3012182	<p>Description: The command <code>ethtool -I --show-fec</code> is not supported by the DPU with kernel 5.4.</p> <p>Workaround: N/A</p> <p>Keywords: Kernel; show-fec</p> <p>Reported in version: 3.9.0</p>
2855986	<p>Description: After disabling SR-IOV VF on a virtio device, removing virtio-net/PCIe driver from guest OS may render the virtio controller unusable .</p> <p>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.</p> <p>Keywords: Virtio-net; VF</p> <p>Reported in version: 3.9.0</p>
2863456	<p>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.</p> <p>Workaround: N/A</p> <p>Keywords: ASAP2; IPsec Full Offload</p> <p>Reported in version: 3.9.0</p>

Ref #	Issue
2982184	<p>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.</p> <p>Workaround: Power-cycle the BlueField to fix the EEPROM issue. Manual recovery of the boot options and/or SW installation may be needed.</p> <p>Keywords: Reset; EEPROM</p> <p>Reported in version: 3.9.0</p>
2853408	<p>Description: Some pre-OS environments may fail when sensing a hot plug operation during their boot stage.</p> <p>Workaround: Run "<code>mlxconfig -d <mst dev> set PF_LOG_BAR_SIZE=0</code>".</p> <p>Keywords: BIOS; hot-plug; Virtio-net</p> <p>Reported in version: 3.9.0</p>
2934833	<p>Description: Running I/O traffic and toggling both physical ports status in a stressful manner on the receiving-end machine may cause traffic loss.</p> <p>Workaround: N/A</p> <p>Keywords: MLNX_OFED; RDMA; port toggle</p> <p>Reported in version: 3.8.5</p>
2911425	<p>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.</p> <p>Workaround: Update the BIOS version to 2.4 which should correctly detect the PCIe device with the bigger BAR size.</p> <p>Keywords: Scalable functions; BIOS</p> <p>Reported in version: 3.8.5</p>
N/A	<p>Description: Only QP queues are supported for GGA accelerators from this version onward.</p> <p>Workaround: N/A</p> <p>Keywords: Firmware; SQ; QP</p> <p>Reported in version: 3.8.0</p>

Ref #	Issue
2846108	<p>Description: Setting <code>VHCA_TRUST_LEVEL</code> does not work when there are active SFs or VFs.</p> <p>Workaround: N/A</p> <p>Keywords: Firmware; SF; VF</p> <p>Reported in version: 3.8.0</p>
2750499	<p>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>).</p> <p>Workaround: N/A</p> <p>Keywords: Devlink</p> <p>Reported in version: 3.7.1</p>
2730157	<p>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).</p> <p>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.</p> <p>Keywords: Kernel; upgrade</p> <p>Reported in version: 3.7.0</p>
2706710	<p>Description: Call traces are seen on the host when recreating VFs before the controller side finishes the deletion procedure.</p> <p>Workaround: N/A</p> <p>Keywords: Virtio-net controller</p> <p>Reported in version: 3.7.0</p>
26845478	<p>Description: 3rd party (netkvm.sys) Virtio-net drivers for Windows do not support SR-IOV.</p> <p>Workaround: N/A</p> <p>Keywords: Virtio-net; SR-IOV; WinOF-2</p> <p>Reported in version: 3.7.0</p>

Ref #	Issue
2684501	<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> <p>Workaround: N/A</p> <p>Keywords: Log; cma_alloc; memory</p> <p>Reported in version: 3.7.0</p>
259016	<p>Description: ibdev2netdev tool is not supported for PCIe PF operating in switchdev mode or on SFs.</p> <p>Workaround: N/A</p> <p>Keywords: ibdev2netdev</p> <p>Reported in version: 3.6.0.11699</p>
2590016	<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> <p>Workaround: Set <code>OPENSSL_CONF=/etc/ssl/openssl.cnf.orig</code> before using the curl utility. For example:</p> <pre data-bbox="224 1209 1463 1413"># OPENSSL_CONF=/etc/ssl/openssl.cnf.orig curl -0 https://tpo.pe/pathogen.vim</pre> <div data-bbox="224 1472 1463 1734" style="background-color: #ffffcc; padding: 10px;"> <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> </div> <p>Keywords: OpenSSL; curl</p> <p>Reported in version: 3.6.0.11699</p>

Ref #	Issue
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>
2377021	<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>
258	<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>

Ref #	Issue
1408	<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>
191618	<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 2024. PDF Generated on 11/25/2024