



## Known Issues

# Table of contents

DOCA Framework Known Issues

---

DOCA-Host and DOCA Drivers Known Issues

---

BlueField-3 Firmware Known Issues

---

BlueField-2 Firmware Known Issues

---

## DOCA Framework Known Issues

| Reference | Description   |
|-----------|---|
| 4297489   | Description: Due to incompatibility between DPA and host libraries, a DPA device application must be recompiled after updating DOCA to a newer version.   |
|           | Workaround: N/A   |
|           | Keyword: DPA; host library; update  |
|           | Reported in version: 2.10.0   |
| 4287011   | Description: Disabling OVS CT (using <code>ovs-vsctl set o . other_config:hw-offload-ct-size=0</code> ) and attempting to offload CT rules is not supported and could lead to OVS crashes.  |
|           | Workaround: N/A   |
|           | Keyword: OVS  |
|           | Reported in version: 2.10.0   |
| 4270602   | Description: UEFI/ATF firmware does not upgrade as part of the Linux Standard Tool process when Secure Boot is disabled.  |
|           | Workaround: Remove PK key and initiate UEFI/ATF firmware upgrade again. To remove the PK key, use the UEFI menu to navigate to Device Manager → Secure Boot Configuration → Custom Secure Boot Options → PK Options → Delete Signature. |
|           | Keyword: UEFI/ATF; PK; Secure Boot; EFI Capsule Authentication  |
|           | Reported in version: 2.10.0   |
| 4200690   | Description: The fTPM trusted application is signed for testing purposes only (i.e., not securely) with a development key.  |
|           | Workaround: N/A   |
|           | Keyword: fTPM over OP-TEE   |
|           | Reported in version: 2.10.0   |
| 3987526   | Description: OVS-DOCA offload of meter with sFlow is not supported and may cause OVS application to crash.  |
|           | Workaround: N/A   |

| Reference | Description   |
|-----------|---|
|           | Keyword: OVS-DOCA; meter; sFlow   |
|           | Reported in version: 2.9.0  |
| N/A       | Description: Applications using DPA might not work with older firmware versions .   |
|           | Workaround: Full upgrade of all DOCA 2.9.0 components including the firmware (i.e., doca-host and BF-Bundle) .  |
|           | Keyword: DPA; backward compatibility  |
|           | Reported in version: 2.9.0  |
| N/A       | Description: Applications using FlexIO SDK API may have missing symbols during runtime.   |
|           | Workaround: Re-compile FlexIO-based applications with the DOCA 2.9.0 release.   |
|           | Keyword: FlexIO; backward compatibility   |
|           | Reported in version: 2.9.0  |
| 4095728   | Description: Corrupt create repo causes doca-kernel repo to not contain the repo data.  |
|           | Workaround: If repo data is missing after installing the doca-kernel repo, run <code>createrepo --help</code> . If no output is generated, then the createrepo is corrupted and must be removed and reinstalled.  |
|           | Keyword: Kernel; repo   |
|           | Reported in version: 2.9.0  |
| 4049034   | Description: On openEuler 22.03 SP3 and openEuler 20.03 SP1, it is not possible to do yum update after BFB installation.  |
|           | Workaround: To perform yum update with either openEuler 22.03 SP3 and openEuler 20.03 SP1, follow these procedures depending on the use case: <ul style="list-style-type: none"> <li>1. To update DOCA only, disable ALL repos: <pre>dnf --disablerepo='*' --enablerepo='kubernetes,doca' -y update</pre> </li> <li>2. To update OS components, exclude <code>rdma-core</code> using one of the following options:</li> </ul> |

| Reference           | Description   |
|---------------------|---|
|                     | <ul style="list-style-type: none"> <li>◦ Option 1 – disable <code>rdma-core</code> from <code>dnf.conf</code> using <code>excludepkgs=rdma-core*oe2203sp3*</code></li> <li>◦ Option 2 – disable <code>rdma-core</code> from <code>openEuler.repo</code> file under OS and everything using <code>exclude=rdma-core*</code></li> <li>◦ Option 3 – run <pre style="background-color: #f0f0f0; padding: 10px; border: 1px solid #ccc;">yum update -x rdma-core</pre> </li> </ul> |
|                     | Keyword: openEuler  |
|                     | Reported in version: 2.9.0  |
| 40<br>46<br>18<br>0 | Description: PCIe data IDs that require <code>Node</code> , <code>PCIe index</code> , and <code>Depth</code> parameters in <code>doca_telemetry_diag</code> , the only valid values are 0, 0, 0.  |
|                     | Workaround: N/A   |
|                     | Keyword: DOCA Telemetry   |
|                     | Reported in version: 2.9.0  |
| 41<br>29<br>71<br>5 | Description: Compiling Rocky 9.2 may fail when using GCC with the "native" arch flag.   |
|                     | Keyword: Upgrade to toolset 13 (gcc 13).  |
|                     | Keyword: Linux; GCC   |
|                     | Reported in version: 2.9.0  |
| 40<br>35<br>55<br>3 | Description: <code>oper_sample_period</code> does not always reflect the correct sample period. In some cases, it will reflect the <code>admin_sample_period</code> instead.  |
|                     | Workaround: N/A   |
|                     | Keyword: Core   |
|                     | Reported in version: 2.8.0  |
| 40<br>23<br>25<br>7 | Description: If RDMA samples are compiled with memory sanitizer enabled, "read memory leak" errors are printed when running the samples with the RDMA CM flag and when running the client before the server.  |

| Reference                 | Description  |
|---------------------------|--|
|                           | <p>Workaround: Make sure to start the RDMA Server before RDMA Client.</p> <p>Keyword: DOCA RDMA; samples</p> <p>Reported in version: 2.8.0</p>   |
| 40<br>21<br>75<br>2<br>40 | <p>Description: In all RDMA samples, if an error occurs in any of the following functions:</p> <ul style="list-style-type: none"> <li>• Exporting RDMA/MMAP/Sync event</li> <li>• Connecting RDMA</li> <li>• Writing or reading the descriptors</li> </ul> <p>An error is printed but the sample resumes and might:</p> <ol style="list-style-type: none"> <li>1. Fail later, or be in busy-wait state indefinitely; and/or</li> <li>2. Result in access to an unknown address, causing an address sanitizer violation.</li> </ol> |
| 21<br>74<br>8             | <p>Workaround for 1: Either:</p> <ul style="list-style-type: none"> <li>• Follow the error logs to verify no errors occurred in the relevant function. And if it did, stop the sample.</li> <li>• Fix the issue locally.</li> </ul> <p>Workaround for 2: The mentioned address sanitizer violation shall be ignored in case of an error in a relevant function.</p> <p>Keyword: DOCA RDMA; samples</p> <p>Reported in version: 2.8.0</p>   |
| 40<br>22<br>56<br>3       | <p>Description: OVS-DOCA connection tracking with E2E enabled is not supported.</p> <p>Workaround: N/A</p> <p>Keyword: OVS-DPDK; connection tracking; E2E</p> <p>Reported in version: 2.8.0</p>  |
| 38<br>37<br>25<br>5       | <p>Description: When running Arm shutdown from the host OS it is expected to get the message <code>-E- Failed to send Register MRSI</code>. This message should be ignored.</p> <p>Workaround: Wait 2 more minutes before rebooting the host. Before proceeding with host OS reboot, it is recommended to query the operational state of the</p>   |

| Reference           | Description  |
|---------------------|--|
|                     | <p>BlueField Arm cores from the BlueField BMC to verify that shutdown state has been reached. Run the following command:</p> <pre data-bbox="235 373 1463 575">ipmitool -C 17 -I lanplus -H &lt;bmc_ip&gt; -U root -P &lt;password&gt; raw 0x32 0xA3</pre> <p>Expected output is "06".</p> <p>Keyword: Host OS; reboot; error</p> <p>Reported in version: 2.7.0</p>  |
| 38<br>44<br>70<br>5 | <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 BlueField eMMC device which may cause degraded performance of the BlueField 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: 2.7.0</p>  |
| 38<br>77<br>72<br>5 | <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="235 1220 1463 1421">echo "DISPLAY_LEVEL 2" &gt; /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 data-bbox="235 1528 1463 1787">[13:58:39] INFO: Installation finished ... [14:01:53] INFO: Rebooting...</pre> <p>Keyword: NIC mode; BFB install</p> <p>Reported in version: 2.7.0</p> |

| Reference | Description  |
|-----------|--|
| 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: 2.7.0</p>   |
| 3855485   | <p>Description: When enabling the <code>PCI_SWITCH_EMULATION_ENABLE</code> NVconfig, the mlx devices, and potentially the RShim devices disappear. Also, looking at the kernel logs using <code>dmesg</code> shows the following messages:</p> <pre>pci 0000:29:00.0: BAR 0: no space for [mem size 0x0200 0000 64bit pref] pci 0000:29:00.0: BAR 2: no space for [mem size 0x0080 0000 64bit pref] ...</pre> <p>Workaround: N/A</p> <p>Keyword: NVconfig; RShim; dmsg</p> <p>Reported in version: 2.7.0</p> |
| 3831230   | <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 BlueField eMMC device which may cause degraded performance of BlueField 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: 2.7.0</p>  |
| 3743879   | <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>.</p>  |



| Reference           | Description  |
|---------------------|--|
|                     | <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" &gt; /dev/rshim0/misc mlxfwreset &lt;arguments&gt; echo "DROP_MODE 0" &gt; /dev/rshim0/misc </pre> <p>Keyword: Timeout; mlxfwreset; INTx</p> <p>Reported in version: 2.7.0</p> |
| 36<br>78<br>06<br>9 | <p>Description: If using BlueField with NVMe and mmcbld and configured to boot from mmcbld, users must create <code>bf.cfg</code> file with <code>device=/dev/mmcbld0</code>, then install the <code>*.bfb</code> as normal.</p> <p>Workaround: N/A</p> <p>Keyword: NVMe</p> <p>Reported in version: 2.5.0</p>   |
| 36<br>80<br>53<br>8 | <p>Description: When using strongSwan or OVS-IPsec as explained in the <a href="#">NVIDIA BlueField DPU BSP</a>, the IPsec Rx data path is not offloaded to hardware and occurs in software running on the Arm cores. As a result, bandwidth performance is substantially low.</p> <p>Workaround: N/A</p> <p>Keyword: IPsec</p> <p>Reported in version: 2.5.0</p>  |
| N/A                 | <p>Description: Execution unit partitions are still not implemented and would be added in a future release.</p> <p>Workaround: N/A</p> <p>Keyword: EU tool</p> <p>Reported in version: 2.5.0</p>   |

| Reference | Description   |
|-----------|---|
| 366160    | <p>Description: Installing BFB using <code>bfb-install</code> when <code>mlxconfig</code> <code>PF_TOTAL_SF</code> &gt; 1700, triggers server reboot immediately.</p> <p>Workaround: Change <code>PF_TOTAL_SF</code> to 0, perform a <a href="#">graceful shutdown</a>, power cycle, then installing BFB.</p> <p>Keyword: SF; <code>PF_TOTAL_SF</code>; BFB installation</p> <p>Reported in version: 2.2.1</p>  |
| 3594836   | <p>Description: When enabling Flex IO SDK tracer at high rates, a slow-down in processing may occur and/or some traces may be lost.</p> <p>Workaround: Keep tracing limited to ~1M traces per second to avoid a significant processing slow-down. Use tracer for debug purposes and consider disabling it by default.</p> <p>Keyword: Tracer FlexIO</p> <p>Reported in version: 2.2.1</p>   |
| 3592080   | <p>Description: When using UEK8 on the host in DPU mode, creating a VF on the host consumes about 100MB memory on BlueField</p> <p>Workaround: N/A</p> <p>Keyword: UEK; VF</p> <p>Reported in version: 2.2.1</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 data-bbox="235 1398 1463 1604">[ 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: 2.2.0</p> |

| Reference | Description   |
|-----------|---|
| 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: 2.2.0</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 BlueField, 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: 2.2.0</p>   |
| 3561723   | <p>Description: Running <code>mlxfwreset sync 1</code> on NVIDIA Converged Accelerators may be reported as supported although it is not. Executing the reset will fail.</p> <p>Workaround: N/A</p> <p>Keywords: mlxfwreset</p> <p>Reported in version: 2.2.0</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>Keywords: Longevity; mlxfwreset; DPU reboot</p> <p>Reported in version: 2.2.0</p> |
| 3534219   | <p>Description: On BlueField-3 devices, from DOCA 2.2.0 to 32.37.1306 (or lower), the host crashes when executing partial Arm reset (e.g., Arm reboot; BFB push; mlxfwreset).</p> <p>Workaround: Before downgrading the firmware:</p> <ol style="list-style-type: none"> <li>1. Run:</li> </ol>   |

| Reference           | Description  |
|---------------------|--|
|                     | <pre data-bbox="228 296 1469 499">echo 0 &gt; /sys/bus/platform/drivers/mlxbf- bootctl/large_icm</pre> <p data-bbox="228 499 1469 583">2. Reboot Arm.</p> <p data-bbox="228 583 1469 646">Keyword: BlueField-3; downgrade</p> <p data-bbox="228 646 1469 701">Reported in version: 2.2.0</p>   |
| 34<br>62<br>63<br>0 | <p data-bbox="228 701 1469 793">When trying to perform a PXE installation when UEFI Secure Boot is enabled, the following error messages may be observed:</p> <pre data-bbox="228 793 1469 1010">error: shim_lock protocol not found. error: you need to load the kernel first.</pre> <p data-bbox="228 1010 1469 1157">Workaround: Download a Grub EFI binary <a href="#">from the Ubuntu website</a>. For further information on Ubuntu UEFI Secure Boot PXE Boot, please visit Ubuntu's official website.</p> <p data-bbox="228 1157 1469 1220">Keyword: PXE; UEFI Secure Boot</p> <p data-bbox="228 1220 1469 1276">Reported in version: 2.0.2</p> |
| 34<br>48<br>84<br>1 | <p data-bbox="228 1276 1469 1381">Description: While running CentOS 8.2, switchdev Ethernet BlueField runs in "shared" RDMA net namespace mode instead of "exclusive".</p> <p data-bbox="228 1381 1469 1444">Workaround: Use <code>ib_core</code> module parameter <code>netns_mode=0</code>. For example:</p> <pre data-bbox="228 1444 1469 1604">echo "options ib_core netns_mode=0" &gt;&gt; /etc/modprobe.d/mlnx-bf.conf</pre> <p data-bbox="228 1604 1469 1667">Keyword: RDMA; isolation; Net NS</p> <p data-bbox="228 1667 1469 1730">Reported in version: 2.0.2</p>   |
| 27<br>06<br>80<br>3 | <p data-bbox="228 1730 1469 1835">Description: When an NVMe controller, SoC management controller, and DMA controller are configured, the maximum number of VFs is limited to 124.</p> <p data-bbox="228 1835 1469 1890">Workaround: N/A</p>   |

| Reference | Description  |
|-----------|--|
|           | Keyword: VF; limitation  |
|           | Reported in version: 2.0.2   |
| 3273435   | Description: Changing the mode of operation between NIC and DPU modes results in different capabilities for the host driver which might cause unexpected behavior.   |
|           | Workaround: Reload the host driver or reboot the host.   |
|           | Keyword: Modes of operation; driver  |
|           | Reported in version: 2.0.2   |
| 3264749   | Description: In Rocky and CentOS 8.2 inbox-kernel BFBs, RegEx requires the following extra huge page configuration for it to function properly: <pre> sudo hugeadm --pool-pages-min DEFAULT:2048M sudo systemctl start mlx-regex.service systemctl status mlx-regex.service </pre> |
|           | If these commands have executed successfully you should see <code>active (running)</code> in the last line of the output.  |
|           | Workaround: N/A  |
|           | Keyword: RegEx; hugepages  |
|           | Reported in version: 1.5.1   |
| 3240153   | Description: DOCA kernel support only works on a non-default kernel.   |
|           | Workaround: N/A  |
|           | Keyword: Kernel  |
|           | Reported in version: 1.5.0   |
| 3217627   | Description: The <code>doca_devinfo_rep_list_create</code> API returns success on the host instead of <code>Operation not supported</code> .   |
|           | Workaround: N/A  |
|           | Keyword: DOCA core; InfiniBand   |
|           | Reported in version: 1.5.0   |

## DOCA-Host and DOCA Drivers Known Issues

| Reference | Description  |
|-----------|--|
| 4155701   | Description: When offloading xfrm states to hardware, the offloading device is linked to the skb's secpath. If an skb is freed or deferred, an unregister netdevice operation may hang because the netdevice is still being reference-counted. |
|           | Workaround: Remove the netdevice from the xfrm states when the netdevice is unregistered.  |
|           | Keyword: IPSec Crypto Offload  |
|           | Reported in version: 2.10.0  |

## BlueField-3 Firmware Known Issues

| Reference | Description   |
|-----------|---|
| 427722    | <b>Description:</b> In the event of a shorted cable, no SEL (System Event Log) report is generated.   |
|           | <b>Workaround:</b> N/A  |
|           | <b>Keywords:</b> SEL  |
|           | <b>Discovered in Version:</b> 32.41.1000  |
| 3787618   | <b>Description:</b> NVIA register is not allowed for external host if any field of EXTERNAL_HOST_PRIV or EXTERNAL_HOST_PRIV_FAST TLVs is not set as the default.  |
|           | <b>Workaround:</b> N/A  |
|           | <b>Keywords:</b> Host privilege   |
|           | <b>Discovered in Version:</b> 32.41.1000  |
| 363631    | <b>Description:</b> When configuring BlueField-3 Arm cores as PCIe root-complex, all non-mlx5 devices must always set the BlueField-3's IOMMU to disabled or passthrough mode. Turning IOMMU "ON" requires special handling of interrupts in the driver or the use of polling. For further assistance, contact <a href="#">NVIDIA support</a> . |
|           | <b>Workaround:</b> N/A  |

| Reference | Description   |
|-----------|---|
|           | <b>Keywords:</b> IOMMU  |
|           | <b>Discovered in Version:</b> 32.39.2048  |
| 3614529   | <b>Description:</b> The supported DDR5 link speed in SKU B3220, is 5200 MT/s.   |
|           | <b>Workaround:</b> N/A  |
|           | <b>Keywords:</b> DDR5 link speed  |
|           | <b>Discovered in Version:</b> 32.39.2048  |
| 3728450   | <b>Description:</b> SW_RESET with a pending image is currently not supported.   |
|           | <b>Workaround:</b> N/A  |
|           | <b>Keywords:</b> SW_RESET   |
|           | <b>Discovered in Version:</b> 32.39.2048  |
| 3614288   | <b>Description:</b> Occasionally, the device may hang when there a hot plug is performed from a unknown direction.  |
|           | <b>Workaround:</b> N/A  |
|           | <b>Keywords:</b> Hot-plug operation   |
|           | <b>Discovered in Version:</b> 32.39.2048  |
| -         | <b>Description:</b> The I <sup>2</sup> C clock fall time is lower than the 12ns minimum defined in the I2C-bus specification.<br>For further information, refer to the I <sup>2</sup> C-bus Specification, Version 7.0, October 2021, <a href="https://www.i2c-bus.org/">https://www.i2c-bus.org/</a> . |
|           | <b>Workaround:</b> N/A  |
|           | <b>Keywords:</b> I <sup>2</sup> C clock   |
|           | <b>Discovered in Version:</b> 32.39.2048  |
| 3439438   | <b>Description:</b> When connecting to a High Speed Traffic Generator in 400G speed, the linkup time may takes up to 3 minutes.   |
|           | <b>Workaround:</b> N/A  |
|           | <b>Keywords:</b> 400G linkup time   |
|           | <b>Discovered in Version:</b> 32.38.1002  |

| Reference | Description   |
|-----------|---|
| 3534128   | <p><b>Description:</b> External flash access such as flash read using the MFT tools will fail if there is a pending image on the flash.</p> <p><b>Workaround:</b> N/A</p> <p><b>Keywords:</b> Flash access</p> <p><b>Discovered in Version:</b> 32.38.1002</p>  |
| 3534219   | <p><b>Description:</b> On BlueField-3 devices, from DOCA 2.2.0 to 32.37.1306 (or lower), the host crashes when executing partial Arm reset (e.g., Arm reboot; BFB push; mlxfwreset).</p> <p><b>Workaround:</b> Before downgrading the firmware, perform:</p> <ul style="list-style-type: none"> <li>• echo 0 &gt; /sys/bus/platform/drivers/mlxbf-bootctl/large_icm</li> <li>• Arm reboot</li> </ul> <p><b>Keywords:</b> BlueField-3; downgrade</p> <p><b>Discovered in Version:</b> 32.38.1002</p> |
| 3547022   | <p><b>Description:</b> When unloading the network drivers on an external host, sync1 reset may be still reported as 'supported' although it is not. Thus, initiating the reset flow may result in reset failure after a few minutes.</p> <p><b>Workaround:</b> N/A</p> <p><b>Keywords:</b> Sync 1 reset</p> <p><b>Discovered in Version:</b> 32.38.1002</p>   |
| 3439438   | <p><b>Description:</b> When connecting to a Spirent switch in 400G speed, the linkup time may takes up to 3 minutes.</p> <p><b>Workaround:</b> N/A</p> <p><b>Keywords:</b> Spirent, 400G, linkup time</p> <p><b>Discovered in Version:</b> 32.38.1002</p>   |
| 3178339   | <p><b>Description:</b> PCIe PML1 is disabled.</p> <p><b>Workaround:</b> N/A</p> <p><b>Keywords:</b> PCIe PML1</p> <p><b>Discovered in Version:</b> 32.38.1002</p>   |



| Reference | Description  |
|-----------|--|
| 3525865   | <b>Description:</b> Unexpected system behavior might be observed if the driver is loaded while reset is in progress.   |
|           | <b>Workaround:</b> N/A   |
|           | <b>Keywords:</b> Sync 1 reset, firmware reset  |
|           | <b>Discovered in Version:</b> 32.38.1002   |
| 3275394   | <b>Description:</b> When performing PCIe link secondary-bus-reset, disable/enable or mlxfwreset on AMD based Genoa systems, the device takes longer then expected to link up, due to a PCIe receiver termination misconfiguration.                 |
|           | <b>Workaround:</b> N/A   |
|           | <b>Keywords:</b> PCIe  |
|           | <b>Discovered in Version:</b> 32.37.1306   |
| 2878841   | Description: The firmware rollback fails for the signature retransmit flow if the QPN field is configured in the mkey (as it only allows the given QP to use this Mkey) as the firmware rollback flow relies on an internal QP that uses the mkey. |
|           | Workaround: N/A  |
|           | Keywords: Signature retransmit flow  |
|           | Discovered in Version: 32.37.1306  |
| 3412847   | Description: Socket-Direct is currently not supported.   |
|           | Workaround: N/A  |
|           | Keywords: Socket-Direct  |
|           | Discovered in Version: 32.37.1306  |

## BlueField-2 Firmware Known Issues

| Reference | Issue   |
|-----------|---|
| 2169950   | Description: When decapsulation on a packet occurs, the FCS indication is not calculated correctly. |

| Reference       | Issue  |
|-----------------|--|
|                 | Workaround: N/A  |
|                 | Keywords: FCS  |
|                 | Discovered in Version: 24.42.1000  |
| 375<br>491<br>3 | Description: PHYless Reset is currently not supported.   |
|                 | Workaround: N/A  |
|                 | Keywords: PHYless Reset  |
|                 | Discovered in Version: 24.40.1000  |
| 352<br>586<br>5 | Description: Unexpected system behavior might be observed if the driver is loaded while reset is in progress.  |
|                 | Workaround: N/A  |
|                 | Keywords: Sync 1 reset, firmware reset   |
|                 | Discovered in Version: 24.39.2048  |
| 354<br>702<br>2 | Description: When <code>tx_port_ts</code> is set to <code>"true"</code> , due to a compensation mechanism in the Tx timestamp available in some hardware Rx timestamp errors, a symmetrical error and no clock offset occur when using the timestamps to synchronize the device clock. |
|                 | This might also cause an error while using timestamps for delay measurements (e.g., delay measurements reported by a PTP daemon) and even negative delay measurements in some cases.   |
|                 | Workaround: N/A  |
|                 | Keywords: PTP path delay   |
|                 | Discovered in Version: 24.38.1002  |
| 354<br>702<br>2 | Description: When unloading the network drivers on an external host, sync1 reset may be still reported as 'supported' although it is not. Thus, initiating the reset flow may result in reset failure after a few minutes.   |
|                 | Workaround: N/A  |
|                 | Keywords: Sync 1 reset   |
|                 | Discovered in Version: 24.38.1002  |
| 345<br>747      | Description: Disabling the Relaxed Ordered (RO) capability ( <code>relaxed_ordering_read_pci_enabled=0</code> ) using the <code>vhca_resource_manager</code> is  |

| Reference | Issue  |
|-----------|--|
| 2         | currently not functional.  |
|           | Workaround: N/A  |
|           | Keywords: Relaxed Ordered  |
|           | Discovered in Version: 24.37.1300  |
| 3296463   | Description: fwreset is currently supported on PCI Gen 4 devices only.   |
|           | Workaround: N/A  |
|           | Keywords: fwreset, PCI Gen4  |
|           | Discovered in Version: 24.37.1300  |
| 2878841   | Description: The firmware rollback fails for the signature retransmit flow if the QPN field is configured in the mkey (as it only allows the given QP to use this Mkey) as the firmware rollback flow relies on an internal QP that uses the mkey. |
|           | Workaround: N/A  |
|           | Keywords: Signature retransmit flow  |
|           | Discovered in Version: 24.37.1300  |
| 3329109   | Description: MFS1S50-H003E cable supports only HDR rate when used as a split cable.  |
|           | Workaround: N/A  |
|           | Keywords: HDR, split cable, MFS1S50-H003E  |
|           | Discovered in Version: 24.37.1300  |
| 3267506   | Description: CRC is included in the traffic byte counters as a port byte counter.  |
|           | Workaround: N/A  |
|           | Keywords: Counters, CRC  |
|           | Discovered in Version: 24.35.2000  |
| 3141072   | Description: The "max_shaper_rate" configuration query via QEEC mlxreg returns a value translated to hardware granularity.   |
|           | Workaround: N/A  |
|           | Keywords: RX Rate-Limiter, Multi-host  |
|           | Discovered in Version: 24.34.1002  |

| Reference      | Issue   |
|----------------|---|
| 2870970        | Description: GTP encapsulation (flex parser profile 3) is limited to the NIC domain. Encapsulating in the FDB domain will render a 0-size length in GTP header. |
|                | Workaround: N/A   |
|                | Keywords: GTP encapsulation   |
|                | Discovered in Version: 24.34.1002   |
| 28990262853408 | Description: Some pre-OS environments may fail when sensing a hot plug operation during their boot stage.   |
|                | Workaround: N/A   |
|                | Keywords: BIOS; Hot plug; Virtio-net  |
|                | Discovered in Version: 24.33.1048   |
| 2870213        | Description: Servers do not recover after configuring <code>PCI_SWITCH_EMULATION_NUM_PORT</code> to 32 followed by power cycle.                                 |
|                | Workaround: N/A   |
|                | Keywords: VirtIO-net; power cycle   |
|                | Discovered in Version: 24.33.1048   |
| 2855592        | Description: When working with 3rd party device (e.g., Paragon) in 25GbE speed, the 25GbE speed must be configured in force mode.                               |
|                | Workaround: N/A   |
|                | Keywords: Force mode, 3rd party devices, 25GbE  |
|                | Discovered in Version: 24.33.1048   |
| 2850003        | Description: Occasionally, when rising a logical link, the link recovery counter is increase by 1.  |
|                | Workaround: N/A   |
|                | Keywords: Link recovery counter   |
|                | Discovered in Version: 24.33.1048   |
| 2616755        | Description: Forward action for IPoIB is not supported on RX RDMA Flow Table.   |
|                | Workaround: N/A   |
|                | Keywords: Steering, IPoIB   |

|           |                                   |
|-----------|-----------------------------------|
| Reference | Issue                             |
|           | Discovered in Version: 24.33.1048 |

**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.