



Known Issues

The following table lists the known issues and limitations for this release of DOCA SDK.

| Reference | Description |
|----------------|--|
| 4032924 | <p>Description: When upgrading to DOCA 2.8.0 on RPM-based OSES, a conflict between strongswan-bf or libreswan and strongSwan may occur.</p> <p>Workaround: Before upgrading, delete strongswan-bf and libreswan:</p> <pre data-bbox="235 594 1461 730">yum remove strongswan-bf strongswan-swanctl libreswan</pre> <p>Keyword: strongSwan; upgrade</p> <p>Reported in version: 2.8.0</p> |
| 4035553 | <p>Description: oper_sample_period does not always reflect the correct sample period. In some cases, it will reflect the admin_sample_period instead.</p> <p>Workaround: N/A</p> <p>Keyword: Core</p> <p>Reported in version: 2.8.0</p> |
| 4023257 | <p>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.</p> <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> |
| 40217524021748 | <p>Description: In all RDMA samples, if an error occurs in any of the following functions:</p> <ul data-bbox="277 1665 841 1787" style="list-style-type: none"> • Exporting RDMA/MMAP/Sync event • Connecting RDMA • Writing or reading the descriptors <p>An error is printed but the sample resumes and might:</p> |

| | |
|-----------------|---|
| | <p>1. Fail later, or be in busy-wait state indefinitely; and/or</p> <p>2. Result in access to an unknown address, causing an address sanitizer violation.</p> |
| | <p>Workaround for 1: Either:</p> <ul style="list-style-type: none"> • Follow the error logs to verify no errors occurred in the relevant function. And if it did, stop the sample. • Fix the issue locally. <p>Workaround for 2: The mentioned address sanitizer violation shall be ignored in case of an error in a relevant function.</p> |
| | Keyword: DOCA RDMA; samples |
| | Reported in version: 2.8.0 |
| 396 194 0 | <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> |
| 398 985 1 | <p>Description: A DOCA Flow pipe has multiple actions. When the action idx is not 0 and it has a shared endecap action, a crash occurs when attempting to create an entry.</p> <p>Workaround: N/A</p> <p>Keyword: DOCA Flow</p> <p>Reported in version: 2.8.0</p> |
| 398 890 4 | <p>Description: Failure to create a control entry with shared endecap action.</p> <p>Workaround: N/A</p> <p>Keyword: DOCA Flow</p> <p>Reported in version: 2.8.0</p> |
| 388 667 4 | <p>Description: Installing doca-all and other DOCA metapackages does not install the mlnx-nvme driver.</p> <p>Workaround: mlnx-nvme is only needed for NVMe-over-RDMA remote storage support. If you wish to install it, add the mlnx-nvme package to the install</p> |

| | |
|-----------------|---|
| | <p>command.</p> <ul style="list-style-type: none"> On RHEL: <pre>apt install doca-all mlnx-nvme-modules</pre> On Ubuntu: <pre>dnf install doca-all-kmod-mlnx-nvme</pre> |
| | Keyword: NVMe; DOCA profile |
| | Reported in version: 2.7.0 |
| 388 593 0 | <p>Description: When installing DOCA-Host on a system using NVMe storage (typically local NVMe disk), and the script doca-kernel-support is used to rebuild and install kernel modules, unloading the mlx5 drivers is only possible after also unmounting the NVMe storage, which would typically necessitate a reboot.</p> <p>Workaround: N/A</p> <p>Keyword: NVMe; doca-kernel-support; DOCA for host</p> <p>Reported in version: 2.7.0</p> |
| 383 725 5 | <p>Description: When running Arm shutdown from the host OS it is expected to get the message -E- Failed to send Register MRSI. 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 BlueField Arm cores from the BlueField BMC to verify that shutdown state has been reached. Run the following command:</p> <pre>ipmitool -C 17 -I lanplus -H <bmc_ip> -U root -P <password> 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> |

| | |
|-----------------|--|
| 384 470 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> |
| 387 772 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="233 642 1464 821">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 data-bbox="233 1014 1464 1234">[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> |
| 385 570 2 | <p>Description: Trying to jump from a steering level in the hardware to a lower level using software steering is not supported on rdma-core lower than 48.x.</p> <p>Workaround: N/A</p> <p>Keyword: RDMA; SWS</p> <p>Reported in version: 2.7.0</p> |
| 385 548 5 | <p>Description: When enabling the PCI_SWITCH_EMULATION_ENABLE NVconfig, the mlx devices, and potentially the RShim devices disappear. Also, looking at the kernel logs using dmesg shows the following messages:</p> <pre data-bbox="233 1822 1464 1944">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> |

| | |
|-----|---|
| | ... |
| | Workaround: N/A |
| | Keyword: NVconfig; RShim; dmsg |
| | Reported in version: 2.7.0 |
| 383 | 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. |
| 123 | Workaround: Upgrade to Linux Kernel version 5.10 or later. |
| 0 | Keyword: eMMC discard; trim functionality |
| | Reported in version: 2.7.0 |
| 374 | Description: mlxfwreset could timeout on servers where the RShim driver is running and INTx is not supported. The following error message is printed: BF reset flow encountered a failure due to a reset state error of negotiation timeout. |
| 387 | Workaround: Set PCIE_HAS_VFIO=0 and PCIE_HAS_UIO=0 in /etc/rshim.conf and restart the RShim driver. Then re-run the mlxfwreset command. If host Linux kernel lockdown is enabled, then manually unbind the RShim driver before mlxfwreset and bind it back after mlxfwreset: |
| 9 | <pre>echo "DROP_MODE 1" > /dev/rshim0/misc mlxfwreset <arguments> echo "DROP_MODE 0" > /dev/rshim0/misc</pre> |
| | Keyword: Timeout; mlxfwreset; INTx |
| | Reported in version: 2.7.0 |
| 366 | Description: Virtio-net controller fails to load if DPA_AUTHENTICATION is enabled. |
| 507 | Workaround: N/A |
| 0 | Keyword: Virtio-net; DPA |
| | Reported in version: 2.5.0 |

| | |
|-----|--|
| 367 | Description: If using BlueField with NVMe and mmcbld and configured to boot from mmcbld, users must create bf.cfg file with device=/dev/mmcbld0, then install the *.bfb as normal. |
| 806 | Workaround: N/A |
| 9 | Keyword: NVMe |
| | Reported in version: 2.5.0 |
| 368 | Description: When using strongSwan or OVS-IPsec as explained in the NVIDIA BlueField DPU BSP , 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. |
| 053 | Workaround: N/A |
| 8 | Keyword: IPsec |
| | Reported in version: 2.5.0 |
| N/A | Description: Execution unit partitions are still not implemented and would be added in a future release. |
| | Workaround: N/A |
| | Keyword: EU tool |
| | Reported in version: 2.5.0 |
| 366 | Description: Installing BFB using bfb-install when mlxconfig PF_TOTAL_SF>1700, triggers server reboot immediately. |
| 616 | Workaround: Change PF_TOTAL_SF to 0, perform a graceful shutdown , power cycle, then installing BFB. |
| 0 | Keyword: SF; PF_TOTAL_SF; BFB installation |
| | Reported in version: 2.2.1 |
| 359 | Description: When enabling Flex IO SDK tracer at high rates, a slow-down in processing may occur and/or some traces may be lost. |
| 483 | 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. |
| 6 | Keyword: Tracer FlexIO |
| | Reported in version: 2.2.1 |

| | |
|-----------------|---|
| 359 208 0 | <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> |
| 354 620 2 | <p>Description: After rebooting a BlueField-3 DPU running Rocky Linux 8.6 BFB, the kernel log shows the following error:</p> <div data-bbox="235 554 1463 688" style="border: 1px solid #ccc; padding: 5px; margin: 5px 0;"> <pre>[3.787135] mlxbf_gige MLNXBF17:00: Error getting PHY irq. Use polling instead</pre> </div> <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> |
| 356 604 2 | <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> |
| 354 647 4 | <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 /etc/bf.cfg file with the relevant PXE boot entries, then run the command bfcfg.</p> <p>Keyword: PXE; boot; MAC</p> <p>Reported in version: 2.2.0</p> |
| 356 172 3 | <p>Description: Running mlxfwreset sync 1 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> |

| | |
|-----------------|---|
| 330 648 9 | <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 intel_idle.max_cstate=1 entry to the kernel command line.</p> <p>Keywords: Longevity; mlxfwreset; DPU reboot</p> <p>Reported in version: 2.2.0</p> |
| 353 848 6 | <p>Description: When removing LAG configuration from BlueField, a kernel warning for uverbs_destroy_ufile_hw is observed if virtio-net-controller is still running.</p> <p>Workaround: Stop virtio-net-controller service before cleaning up bond configuration.</p> <p>Keywords: Virtio-net; LAG</p> <p>Reported in version: 2.2.0</p> |
| 353 421 9 | <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"> 1. Run: <div data-bbox="315 1171 1461 1310" style="border: 1px solid #ccc; padding: 5px; margin: 5px 0;"> <pre>echo 0 > /sys/bus/platform/drivers/mlxbf-bootctl/large_icm</pre> </div> 2. Reboot Arm. <p>Keyword: BlueField-3; downgrade</p> <p>Reported in version: 2.2.0</p> |
| 346 263 0 | <p>When trying to perform a PXE installation when UEFI Secure Boot is enabled, the following error messages may be observed:</p> <div data-bbox="237 1671 1461 1848" style="border: 1px solid #ccc; padding: 5px; margin: 5px 0;"> <pre>error: shim_lock protocol not found. error: you need to load the kernel first.</pre> </div> |

| | |
|-----|---|
| | <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: 2.0.2</p> |
| 344 | <p>Description: While running CentOS 8.2, switchdev Ethernet BlueField runs in "shared" RDMA net namespace mode instead of "exclusive".</p> |
| 884 | <p>Workaround: Use ib_core module parameter netns_mode=0. For example:</p> |
| 1 | <pre>echo "options ib_core netns_mode=0" >> /etc/modprobe.d/mlnx-bf.conf</pre> |
| | <p>Keyword: RDMA; isolation; Net NS</p> |
| | <p>Reported in version: 2.0.2</p> |
| 270 | <p>Description: When an NVMe controller, SoC management controller, and DMA controller are configured, the maximum number of VFs is limited to 124.</p> |
| 680 | <p>Workaround: N/A</p> |
| 3 | <p>Keyword: VF; limitation</p> |
| | <p>Reported in version: 2.0.2</p> |
| 327 | <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> |
| 343 | <p>Workaround: Reload the host driver or reboot the host.</p> |
| 5 | <p>Keyword: Modes of operation; driver</p> |
| | <p>Reported in version: 2.0.2</p> |
| 326 | <p>Description: In Rocky and CentOS 8.2 inbox-kernel BFBs, RegEx requires the following extra huge page configuration for it to function properly:</p> |
| 474 | <pre>sudo hugeadm --pool-pages-min DEFAULT:2048M sudo systemctl start mlx-regex.service systemctl status mlx-regex.service</pre> |
| 9 | |

| | |
|-----------------|---|
| | If these commands have executed successfully you should see active (running) in the last line of the output. |
| | Workaround: N/A |
| | Keyword: RegEx; hugepages |
| | Reported in version: 1.5.1 |
| 324 015 3 | Description: DOCA kernel support only works on a non-default kernel. |
| | Workaround: N/A |
| | Keyword: Kernel |
| | Reported in version: 1.5.0 |
| 321 762 7 | Description: The doca_devinfo_rep_list_create API returns success on the host instead of Operation not supported. |
| | Workaround: N/A |
| | Keyword: DOCA core; InfiniBand |
| | Reported in version: 1.5.0 |

Copyright 2024. PDF Generated on 08/15/2024