



HBN Service Release Notes

Table of contents

Changes and New Features

Supported Platforms and Interoperability

Supported BlueField Networking Platforms

Supported BlueField OS

Verified Scalability Limits

Known Issues

Bug Fixes

The following subsections provide information on HBN service new features, interoperability, known issues, and bug fixes.

Changes and New Features

HBN 2.5.0 offers the following new features and updates:

- Added support for Destination NAT (DNAT) for Overlay Gateways
- Added Configurable Control Plane Policing (CoPP) support
- Added ACL support for L3 sub-interfaces
- Added initial Bidirectional Forwarding Detection (BFD) support for BGP (alpha level)
- Bug fixes

HBN 2.5.0 does not include any user affecting changes if upgrading from the previous HBN version.

Supported Platforms and Interoperability

Supported BlueField Networking Platforms

HBN 2.5.0 has been validated on the following NVIDIA® BlueField® networking platforms:

- BlueField-2 DPUs:
 - BlueField-2 P-Series DPU 25GbE Dual-Port SFP56; PCIe Gen4 x8; Crypto Enabled; 16GB on-board DDR; 1GbE OOB management; HHHL
 - BlueField-2 P-Series DPU 25GbE Dual-Port SFP56; integrated BMC; PCIe Gen4 x8; Secure Boot Enabled; Crypto Enabled; 16GB on-board DDR; 1GbE OOB management; FHHL
 - BlueField-2 P-Series DPU 25GbE Dual-Port SFP56; integrated BMC; PCIe Gen4 x8; Secure Boot Enabled; Crypto Enabled; 32GB on-board DDR; 1GbE OOB management; FHHL
 - BlueField-2 P-Series DPU 100GbE Dual-Port QSFP56; integrated BMC; PCIe Gen4 x16; Secure Boot Enabled; Crypto Enabled; 32GB on-board DDR; 1GbE OOB management; FHHL

- BlueField-3 DPUs:
 - BlueField-3 B3210E E-Series FHHL DPU; 100GbE (default mode) / HDR100 IB; Dual-port QSFP112; PCIe Gen5.0 x16 with x16 PCIe extension option; 16 Arm cores; 32GB on-board DDR; integrated BMC; Crypto Enabled
 - BlueField-3 B3220 P-Series FHHL DPU; 200GbE (default mode)/NDR200 IB; Dual-port QSFP112; PCIe Gen5.0 x16 with x16 PCIe extension option; 16 Arm cores; 32GB on-board DDR; integrated BMC; Crypto Enabled
 - BlueField-3 B3240 P-Series Dual-slot FHHL DPU; 400GbE/NDR IB (default mode); Dual-port QSFP112; PCIe Gen5.0 x16 with x16 PCIe extension option; 16 Arm cores; 32GB on-board DDR; integrated BMC; Crypto Enabled

- BlueField-3 SuperNICs:
 - BlueField-3 B3210L E-series FHHL SuperNIC, 100GbE (default mode)/HDR100 IB, Dual port QSFP112, PCIe Gen4.0 x16, 8 Arm cores, 16GB on-board DDR, integrated BMC, Crypto Enabled
 - BlueField-3 B3220L E-Series FHHL SuperNIC, 200GbE (default mode)/NDR200 IB, Dual-port QSFP112, PCIe Gen5.0 x16, 8 Arm cores, 16GB on-board DDR, integrated BMC, Crypto Enabled
 - BlueField-3 B3140L E-Series FHHL SuperNIC, 400GbE/NDR IB (default mode), Single-port QSFP112, PCIe Gen5.0 x16, 8 Arm cores, 16GB on-board DDR, integrated BMC, Crypto Enabled
 - BlueField-3 B3140H E-series HHHL SuperNIC, 400GbE (default mode)/NDR IB, Single-port QSFP112, PCIe Gen5.0 x16, 8 Arm cores, 16GB on board DDR, integrated BMC, Crypto Enabled

 Note

BlueField platforms with 8GB on-board DDR memory are currently not supported with HBN.

Supported BlueField OS

HBN 2.5.0 supports DOCA 2.10.0 (BSP 4.10.0) on Ubuntu 22.04 OS.

Verified Scalability Limits

HBN 2.5.0 has been tested to sustain the following maximum scalability limits:

Limit	BlueField-2	BlueField-3	Comments
VTEP peers (BlueFields per control plane) in the fabric	8k ¹	8k ¹	Number of BlueFields (VTEPs) within a single overlay fabric (reachable in the underlay)
L2 VNIs/Overlay networks per BlueField	20	20	Total number of L2 VNIs in the fabric for L2 VXLAN use-case assuming every interface is associated with its own VLAN + L2 VNI
L3 VNIs/Overlay networks per BlueField	20 - for up to 4K VTEPs 10 - for up to 8K VTEPs	20 - for up to 4K VTEPs 10 - for up to 8K VTEPs	Total number of L3 VNIs in the fabric for L3 VXLAN use-case assuming every interface is associated with its own VLAN + L2 VNI + L3 VNI + VRF
BlueFields per a single L2 VNI network	8k	8k	Total number of DPUs, configured with the same L2 VNI (3 real DPUs, 2000 emulated VTEPs)
BlueFields per a single L3 VNI network	8k	8k	Total number of DPUs, configured with the same L3 VNI (3 real DPUs, 2000 emulated VTEPs)
Maximum number of local MAC/ARP entries per BlueField	20	20	Max total number of MAC/ARP entries learned from the host on the DPU
Maximum number of local BGP routes per BlueField	200	200	Max total number of BGP routes advertised by the host to the BlueField (BGP peering with the host): 100 IPv4 + 100 IPv6

Limit	BlueField-2	BlueField-3	Comments
Maximum number of remote L3 LPM routes (underlay)	8k	8k	IPv4 or IPv6 underlay LPM routes per BlueField (default + host routes + LPM)
Maximum number of EVPN type-2 entries	16K	16k	Remote overlay MAC/IP entries for compute peers stored on a single BlueField (L2 EVPN use case)
Maximum number of EVPN type-5 entries	32K	80K	Remote overlay L3 LPM entries for compute peers stored on a single BlueField (L3 EVPN use case)
Maximum number of Next-hops in ECMP Next-hop group	16	16	Max number of next-hops in ECMP Next-hop group (for overlay ECMP)
Maximum number of PFs on the Host side	2	2	Total number of PFs visible to the host
Maximum number of VFs on the Host side	16	16	Total number of VFs created on the host
Maximum number of SFs on BlueField side	2	2	Total number of SF devices created on BlueField Arm

1. Tested with 4 VNIs.

Known Issues

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

Reference ID	Description
4200335	<p>Description: DHCP issues may lead to incomplete resolve.conf on the DPU. The consequences can be DNS resolution failures and/or the hostname being set to 'localhost'.</p> <p>Workaround: Reboot.</p> <p>Keywords: DPU, DHCP, resolve.conf, hostname, localhost, DNS</p> <p>Reported in HBN Version: 2.5.0</p>
4196880	<p>Description: DHCP issues may lead to incomplete resolve.conf on the HBN container. The consequences can be DNS resolution failures and/or the hostname being set to 'localhost'.</p> <p>Workaround: Following is a list of possible workarounds.</p> <ol style="list-style-type: none"> 1. Stop container and restart kubelet@mgmt and containerd@mgmt, and then start the container. 2. Reboot. 3. systemctl restart dhclient@oob_net0 <p>Keywords: DHCP, resolve.conf, hostname, localhost, DNS</p> <p>Reported in HBN Version: 2.5.0</p>
4257285	<p>Description: ARP packets between DPU and Outside World that are forwarded using HBN are not HW offloaded.</p> <p>Workaround: N/A</p> <p>Keywords: ARP, Outside World, HW offload</p> <p>Reported in HBN Version: 2.5.0</p>
4279243	<p>Description: OVS does not punt the IPv6 Neighbour Advertisements with unicast Destination MAC address to CPU, hence endpoint MAC may not be learnt on the VTEP as long as the endpoint is silent (resulting in traffic towards endpoint to be software forwarded). This is applicable only for absolutely silent end hosts which do not initiate any IPv6 Neighbour Solicitation messages. Once the silent end host initiates the traffic, traffic will be hardware forwarded. This issue will persist only if the end points never initiate any traffic but only send IPv6 Neighbour Advertisements as a response to IPv6 Neighbour Solicitation (rare scenario)</p>

	<p>Workaround: User need to add following ovs rule on DPU: sudo ovs-ofctl add-flow br-hbn 'table=3,priority=100,icmp6,icmp_type=136 actions=resubmit(,98)' This rule will punt IPv6 NA packets to HBN To check if the rule is present: sudo ovs-ofctl dump-flows br-hbn --color --names table=3</p>
	Keywords: IPV6 ND (Neighbour Discovery), NA (Neighbour Advertisement), silent host
	Reported in HBN Version: 2.5.0
4 2 5	Description: When several ports are configured to be part of a bridge and later reconfigured to be L3 interfaces, only one port (the 1st port that was enslaved previously to bridge) gets correctly reprogrammed as an L3 interface in nl2doca. The remaining ports continue to appear as bridged ports in nl2doca.
5 7 0	Workaround: Restart HBN container after unconfiguring bridge ports and before reconfiguring those ports as L3 interfaces
8	Keywords: Bridge port, L3 Port
	Reported in HBN Version: 2.5.0
4 2 5	Description: ARP packets between DPU and Outside World that are forwarded using HBN are not HW offloaded.
7	Workaround: N/A
2 8	Keywords: ARP
5	Reported in HBN Version: 2.5.0
4 2 1 4	Description: Packets with destination port 4789 coming from host side will be dropped if HBN is configured as L3 evpn, leading to Customer / Tenant encapsulated VxLan traffic drop in L3EVPN scenario. This prevents running vxlan underlay over hbn vxlan overlay in L3 evpn scenarios.
6	Workaround: N/A
3 1	Keywords: 4789, vxlan overlay, vxlan underlay
	Reported in HBN Version: 2.5.0
4 1 9 3	Description: When LLDP is enabled on BlueField, it may not work on uplink ports when HBN service is running. This might happen if LLDP is running without any interface filter configuration.
0 4 6	Workaround: Configure LLDP to run only on interfaces where LLDP is required to be run, using a configuration file, <code>/etc/lldpd.d/ports.conf</code> , for the <code>lldpd</code> daemon. The interfaces can be specified using a regular expression pattern, if needed. For example:

- To run LLDP only on the uplinks (`p0` and `p1`), the configuration can be done as follows:

```
$ cat /etc/lldpd.d/ports.conf
```

Configure system interface pattern `p[01]`.

- To run LLDP on the uplinks plus some host-facing PFs or VFs, the configuration can be done as follows:

```
$ cat /etc/lldpd.d/ports.conf
```

Configure system interface pattern
`p[0-1], pf[0-1]hpf, pf[0-1]vf[0-12]`.

If this configuration file is changed while the LLDP service is running, it must be restarted using `systemctl restart lldpd`.

Keywords: LLDP

Reported in HBN version: 2.4.1

4 Description: Sometimes the DNS resolution may fail if `resolv.conf` is not
2 updated with the proper name server, leading to loss of OOB connectivity.
0

0 Workaround: N/A

3 Keywords: DNS; OOB connectivity
3

5 Reported in HBN version: 2.4.1

Description: The following critical error message is generated during HBN POD
reboot. It can be safely ignored.

4
0 Error message: "CRIT Server 'unix_http_server' running without any HTTP authentication
1 checking"
1

6
8 Workaround: N/A
8

Keywords: Log

Reported in HBN version: 2.4.0

4 Description: When using default BGP timers, an OVS restart may lead to extended
0 traffic loss due to BGP peering reset.
9

8 1	Workaround: N/A
5 8	Keywords: BGP; OVS
	Reported in HBN version: 2.4.0
3 7 4	Description: HBN container may hang in init-sfs during container restart when the HBN YAML file (i.e., <code>/etc/kubelet.d/doco_hbn.yaml</code>) is modified while container is running.
3 9	Workaround: If the container hangs in init-sfs for more than 1 minute, reload the DPU.
4 2	Keywords: Hang; container
	Reported in HBN version: 2.3.0
3 9 6 1 3	Description: The changing of the port number for NVUE REST API using nv CLI/API is not supported. The following command should not be used to change the port number: <code>nv set system api port <port-no></code>
8 7	Workaround: On HBN, NVUE is accessible through 8765 (i.e., default port number).
	Keywords: NVUE API; port number
	Reported in HBN version: 2.3.0
3 9 6 7 4	Description: The command <code>nv show system api connections</code> does not return any data.
7	Workaround: N/A
7 4	Keywords: REST API; nginx
8	Reported in HBN version: 2.3.0
3 8 6 5 6	Description: Packets with destination port 4789/8472 coming from host side will be dropped if HBN is configured as L3 evpn. Functional effect is that Customer / Tenant encapsulated VxLan traffic will be dropped in L3EVPN scenario.
5 6	Workaround: N/A
3 3	Keywords: 4789, 8472
	Reported in HBN version: 2.2.0

3 7 6	Description: A ping or other IP connectivity from a locally connected host in vrf-X to an interface IP address on the DPU/HBN itself in vrf-Y will not work, even if VRF route-leaking is enabled between these two VRFs.
9	Workaround: N/A
3 0	Keyword: IP
9	Reported in HBN version: 2.2.0
3 8 3	Description: Traffic entering HBN service on a host PF/VF main-interface and exiting on a sub-interface of the same PF/VF (and vice versa) is not hardware offloaded. Similarly, traffic entering HBN service on one sub-interface and exiting on another sub-interface of the same host PF/VF is also not hardware offloaded.
5 2	Workaround: N/A
9 5	Keyword: Hardware offload; interfaces
	Reported in HBN version: 2.2.0
3 7 7	Description: The DHCP relay gateway-interface IP address does not automatically pick up the IP address assigned to the associated VRF.
2	Workaround: The gateway-interface IP address must be explicitly configured.
5 5	Keyword: DHCP relay gateway; IP
2	Reported in HBN version: 2.2.0
3 8 9	Description: If NVUE-based routing policy (route map) configuration is used to associated route target extended communities with a EVPN route, only one route target can be specified.
1	Workaround: N/A
5 4	Keyword: NVUE; route target
2	Reported in HBN version: 2.2.0
3 7 5 7	Description: When the HBN container is coming up and applying a large configuration through the NVUE-startup service which includes entities used by DHCP relay (e.g., interfaces, SVIs and VRFs), the DHCP relay service may go into FATAL state. It can be observed using the following command:
6 8 6	<pre> supervisorctl status grep isc-dhcp-relay isc-dhcp-relay-vrf11 RUNNING pid 2069, uptime 0:11:31 isc-dhcp-relay-vrf12 RUNNING pid 2071, uptime 0:11:31 isc-dhcp-relay-vrf13 FATAL Exited too quickly (process log may have details) </pre>

	isc-dhcp-relay-vrf14 FATAL Exited too quickly (process log may have details)
	Workaround: Restart the DHCP relay service which is in FATAL state using the command: <pre>supervisorctl restart <relay-service-name></pre>
	Keyword: DHCP relay; fatal; container; restart
	Reported in HBN version: 2.1.0
	Description: When the DPU boots up after issuing a "reboot" command from the DPU itself, some host-side interfaces may remain down.
	Workaround: 1. Restart openibd: <pre>systemctl restart openibd</pre> 2. Recreate SR-IOV interfaces if they are needed. 3. Replay interface config. For example: <ul style="list-style-type: none"> ◦ If using ifupdown2: <pre>ifreload -a</pre> ◦ If using Netplan: <pre>netplan apply</pre>
	Keyword: Reboot
	Reported in HBN version: 1.5.0
3	Description: IPv6 stateless ACLs are not supported.
5	
4	Workaround: N/A
7	
1	Keyword: IPv6 ACL

0 3	Reported in HBN version: 1.5.0
3 3 3 9 3	Description: Statistics for hardware-offloaded traffic are not reflected on SFs inside an HBN container.
0 4	Workaround: Look up the stats using <code>ip -s link show</code> on PFs outside of the HBN container. PFs would show Tx/Rx stats for traffic that is hardware-accelerated in the HBN container.
	Keyword: Statistics; container
	Reported in HBN version: 1.4.0
3 3 5 2 0 0 3	Description: NVUE show, config, and apply commands malfunction if the <code>nvued</code> and <code>nvued-startup</code> services are not in the <code>RUNNING</code> and <code>EXITED</code> states respectively.
	Workaround: N/A
	Keyword: NVUE commands
	Reported in HBN version: 1.3.0
3 1 8 4 7 4 5	Description: The command <code>nv show interface <intf> acl</code> does not show correct information if there are multiple ACLs bound to the interface.
	Workaround: Use the command <code>nv show interface <intf></code> to view the ACLs bound to an interface.
	Keyword: ACLs
	Reported in HBN version: 1.2.0
3 1 5 8 9 3 4	Description: Deleting an NVUE user by removing their password file and restarting the <code>decrypt-user-add</code> service on the HBN container does not work.
	Workaround: Either respawn the container after deleting the file or delete the password file corresponding to the user by running <code>userdel -r username</code> .
	Keyword: User deletion
	Reported in HBN version: 1.2.0
3 1 8 5 0 0 3	Description: When a packet is encapsulated with a VXLAN header, it adds extra bytes which may cause the packet to exceed the MTU of link. Typically, the packet would be fragmented but its silently dropped and no fragmentation happens.
	Workaround: Make sure that the MTU on the uplink port is always 50 bytes more than host ports so that even after adding VXLAN headers, ingress packets do not exceed the MTU.

	Keyword: MTU; VXLAN
	Reported in HBN version: 1.2.0
3 1	Description: On VXLAN encapsulation, the DF flag is not propagated to the outer header. Such a packet may be truncated when forwarded in the kernel, and it may be dropped when hardware offloaded.
8 4 9 0	Workaround: Make sure that the MTU on the uplink port is always 50 bytes more than host ports so that even after adding VXLAN headers, ingress packets do not exceed the MTU.
5	Keyword: VXLAN
	Reported in HBN version: 1.2.0
3 1	Description: When stopping the container using the command <code>crictl stop</code> an error may be reported because the command uses a timeout of 0 which is not enough to stop all the processes in the HBN container.
8 8 6 8 8	Workaround: Pass a timeout value when stopping the HBN container by running: <pre>crictl stop --timeout 60 <hbn-container></pre>
	Keyword: Timeout
	Reported in HBN version: 1.2.0
3 1	Description: The same ACL rule cannot be applied in both the inbound and outbound direction on a port.
2 9	Workaround: N/A
7 4	Keyword: ACLs
9	Reported in HBN version: 1.2.0
3 1	Description: The system's time zone cannot be modified using NVUE in the HBN container.
2 6 5 6 0	Workaround: The time zone can be manually changed by symlinking the <code>/etc/localtime</code> file to a binary time zone's identifier in the <code>/usr/share/zoneinfo</code> directory. For example: <pre>sudo ln -sf /usr/share/zoneinfo/GMT /etc/localtime</pre>

	Keyword: Time zone; NVUE
	Reported in HBN version: 1.2.0
3 1 1	Description: Auto-BGP functionality (where the ASN does not need to be configured but is dynamically inferred by the system based on the system's role as a leaf or spine device) is not supported on HBN.
8 2	Workaround: If BGP is configured and used on HBN, the BGP ASN must be manually configured.
0 4	Keyword: BGP
	Reported in HBN version: 1.2.0
3 2 3 3	Description: Since checksum calculation is offloaded to the hardware (not done by the kernel), it is expected to see an incorrect checksum in the tcpdump for locally generated, outgoing packets. BGP keepalives and updates are some of the packets that show such incorrect checksum in tcpdump.
0	Workaround: N/A
8 8	Keyword: BGP
	Reported in HBN version: 1.2.0
2 8 2	Description: MAC addresses are not learned in the hardware but only in software. This may affect performance in pure L2 unicast traffic.
1	Workaround: N/A
7 8	Keyword: MAC; L2
5	Reported in HBN version: 1.3.0
3 0 1	Description: Due to disabled backend foundation units, some NVUE commands return <code>500 INTERNAL SERVER ERROR</code> / <code>404 NOT FOUND</code> . These commands are related to features or subsystems which are not supported on HBN.
7 2	Workaround: N/A
0	Keyword: Unsupported NVUE commands
2	Reported in HBN version: 1.3.0
2 8 2 8 8 3 8	Description: NetworkManager and other services not directly related to HBN may display the following message in syslog: <div style="border: 1px solid gray; padding: 10px; margin: 10px 0;"> <pre>"netlink: read: too many netlink events. Need to resynchronize platform cache"</pre> </div> The message has no functional impact and may be ignored.

Workaround: N/A
Keyword: Error
Reported in HBN version: 1.3.0

Bug Fixes

The following table lists the known issues which have been fixed for this release of HBN.

Reference	Description
4155959	<p>Description: With uplinks in the <code>br-sfc</code> bridge, IPv6 traffic in uplink-to-uplink direction results in OVS crash resulting in complete traffic drop.</p> <p>Fixed in HBN Version: 2.5.0</p>
4197067	<p>Description: The management VRF does not have an IPv6 address configured, resulting in the absence of a default IPv6 route in the management VRF. Consequently, IPv6 connectivity on the management port is unavailable, and only IPv4 connectivity is supported.</p> <p>Fixed in HBN Version: 2.5.0</p>
4093502	<p>Description: VRF interfaces have a loopback address, but these loopback addresses have scope global, not scope host which can break source IP address lookup for packets originating from the VRF.</p> <p>Fixed in HBN version: 2.4.0</p>
4029473	<p>Description: Rarely, after deletion then creation of an interface, BGP peering over that interface may announce IPv6 routes with an IPv4-mapped IPv6 address as the next hop, which the BGP peer device at the other end can reject.</p> <p>Fixed in HBN version: 2.4.0</p>
4125363	<p>Description: On newer BlueField-2 and BlueField-3 devices, <code>/sys/class/dmi/id/sys_vendor</code> shows <code>Nvidia</code> instead of <code>https://www.mellanox.com</code>, causing NVUE to fail to apply configurations if it is attempted.</p> <p>Fixed in HBN version: 2.4.0</p>
3965589	<p>Description: When SR-IOV VFs are created or deleted and recreated, some ports may stay in <code>ethX</code> naming format and not be properly renamed to <code>pfXvfY</code> format. This results in the port remaining in error state as when running the command <code>ovs-vsctl show</code> due to the SFC and HBN not recognizing it.</p> <p>Fixed in HBN version: 2.4.0</p>
4004191	<p>Description: Due to security fixes on BlueField-2, the number of context switches increased by 20% which may result in user applications (e.g., <code>nl2doca</code>) running slower.</p> <p>Fixed in HBN version: 2.4.0</p>

38 80 35 2	Description: Deleting and re-adding SR-IOV ports might result in some ports in br-hbn bridge going in error state. Fixed in HBN version: 2.4.0
39 60 82 5	Description: When either <code>ENABLE_SFC_HBN</code> or <code>ENABLE_BR_HBN</code> is set to <code>yes</code> in <code>bf.cfg</code> , the initial DHCP request from <code>oob_net0</code> during adapter boot does not contain the <code>NVIDIA/BF/00B</code> string in DHCP option 60 (vendor class identifier). Fixed in HBN version: 2.3.0
35 38 16 7	Description: An explicit restart of FRR service may be required if the BGP AS number is changed via NVUE. Fixed in HBN version: 2.3.0
33 60 69 9	Description: If it is required to decrease the default MTU on interfaces on which HBN operates, after the change is made on the BlueField as well as within HBN, the BlueField must be rebooted for the change to take effect properly. Fixed in HBN version: 2.3.0
38 64 08 0	Description: When an interface is toggled off and on, its sub-interfaces lose their IPv6 addresses and do not get them back. Fixed in HBN version: 2.3.0
36 32 34 4	Description: HBN interfaces on the BlueField side (outside the HBN container) may not get their proper MTU set from systemd-network. Fixed in HBN version: 2.2.0
37 60 86 9	Description: Datapath flow with very low PPS may be deleted before aging time (60 sec) in large scale of number of routes (16K+). Fixed in HBN version: 2.2.0
37 70 99 2	Description: It is not possible to configure an IPv6 default (<code>::/0</code>) static route using NVUE. Fixed in HBN version: 2.2.0
38 24 88 1	Description: When the number of unique ECMP groups used is more than 6, it results in failure of programming prefixes using ECMP-groups greater than 6. Uniqueness is based on ECMP content, so if multiple routes have same nexthop paths, they just use 1 ECMP group. Fixed in HBN version: 2.2.0

37 05 89 4	Description: In an EVPN Symmetric Routing scenario, IPv6 traffic is not hardware offloaded. Fixed in HBN version: 2.2.0
35 19 32 4	Description: The DOCA HBN container takes about 1 minute longer to spawn, as compared to previous HBN release (1.4.0) Fixed in HBN version: 2.1.0
32 19 53 9	Description: TC rules are programmed by OVS to map uplink and host representor ports to HBN service. These rules are ageable and can result in packets needing to get software forwarded periodically to refresh the rules. Fixed in HBN version: 2.1.0
36 10 97 1	Description: The output of the command <code>nv show interface</code> does not display information about VRFs, VXLAN, and bridge. Fixed in HBN version: 2.0.0
34 52 91 4	Description: IPv6 OOB connectivity from the HBN container stops working if the br-mgmt interface on the DPU goes down. When going down, the br-mgmt interface loses its IPv6 address, which is used as the gateway address for the HBN container. If the br-mgmt interface comes back up, its IPv6 address is not added back and IPv6 OOB connectivity from the HBN container will not work Fixed in HBN version: 1.5.0
31 91 43 3	Description: ECMP selection for the underlay path uses the ingress port and identifies uplink ports via round robin. This may not result in uniform spread of the traffic. Fixed in HBN version: 1.4.0
30 49 87 9	Description: When reloading (<code>ifreload</code>) an empty <code>/etc/network/interfaces</code> file, the previously created interfaces are not deleted. Fixed in HBN version: 1.4.0
32 84 60 7	Description: When an ACL is configured for IPv4 and L4 parameters (protocol tcp/udp, source, and destination ports) match, the ACL also matches IPv6 traffic with the specified L4 parameters. Fixed in HBN version: 1.4.0
32 82 11 3	Description: Some DPUs experience an issue with the clock settings after installing a BlueField OS in an HBN setting in which the date reverts back to "Thu Sep 8, 2022".

	Fixed in HBN version: 1.4.0
33 54 02 9	Description: If interfaces on which BGP unnumbered peering is configured are not defined in the <code>/etc/network/interfaces</code> configuration file, BGP peering does not get established on them.
	Fixed in HBN version: 1.4.0

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.