



Virtio-net Service Guide Release Notes

Table of contents

Changes and New Features in v24.10

Known Issues

Bug Fixes

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

Changes and New Features in v24.10

- High availability process is added to handle crashes and reduce downtime
- Added support for dynamic interrupt moderation (DIM)
- Support `VIRTIO_NET_F_CTRL_VLAN`
- Parallel admin VQ commands

Known Issues

The following are known limitations of this NVIDIA® BlueField® virtio-net software version.

Ref #	Issue
387 909 3	Description: When creating a large number of virtio-net VFs, the representor name of the SF may not be renamed.
	Workaround: Use the <code>ip</code> command to rename the representor manually.
	Keyword: Representor
	Reported in version: 24.10
394 390 5	Description: Host OS kernel <3.19 does not support 31 hotplug devices.
	Workaround: Avoid hotplugging more than 20 devices if host OS kernel is <3.19, or upgrade the kernel to ≥ 3.19 .
	Keyword: Host OS; kernel; hotplug
	Reported in version: 24.07
402 216 0	Description: Feature bit <code>VIRTIO_NET_F_CTRL_VLAN</code> is not supported. Enabling it from the hotplug device may results in anomalous behavior.
	Workaround: Disable <code>VIRTIO_NET_F_CTRL_VLAN</code> .
	Keyword: Feature bit
	Reported in version: 24.07

Ref #	Issue
400 126 1	Description: The <code>virtnet.conf</code> file does not check invalid values such as negative numbers or 0.
	Workaround: N/A
	Keyword: Virtnet; config; invalid value
	Reported in version: 24.07
396 559 8	Description: Admin-VQ-based transitional VF show a <code>vf_get</code> error when the controller is restarted. However, VF functionality is not affected.
	Workaround: N/A
	Keyword: Admin VQ; transitional device
	Reported in version: 24.07
396 195 1	Description: Out-of-memory call trace occurs when creating many (>300) VFs on a BlueField running OpenEuler or CentOS 7.6.
	Workaround: Update the kernel to support shared RQ.
	Keyword: OOM; OpenEuler; CentOS 7.6; virtual function
	Reported in version: 24.07
386 268 3	Description: Creating VFs and hotplug PFs in parallel can lead to controller crash.
	Workaround: Create VFs followed by hotplug PF or vice versa.
	Keyword: Virtio-net emulation
	Reported in version: 1.9.0
366 507 0	Description: Virtio-net controller fails to load if <code>DPA_AUTHENTICATION</code> is enabled.
	Workaround: N/A
	Keywords: Virtio-net; DPA
	Reported in version: DOCA 2.5.0
353 848 6	Description: When removing LAG configuration from BlueField, a kernel warning for <code>uverbs_destroy_ufile_hw</code> is observed if virtio-net-controller is still running.
	Workaround: Stop virtio-net-controller service before cleaning up bond configuration.
	Keywords: Virtio-net; LAG

Ref #	Issue
	Reported in version: DOCA 2.2.0
368	Description: Starting from kernel 5.14, the virtio-net TX path has a logic which may trigger infinite loop when <code>vq</code> is broken (e.g., device is removed) under heavy traffic.
380	Workaround: N/A
1	Keyword: Virtio-net
	Reported in version: DOCA 1.8.0
371	Description: When creating/destroying VFs back to back, make sure the virtio-net controller side does not see any alive VF before recreating them from the guest OS (i.e., <code>virtnet query</code>).
452	Workaround: N/A
2	Keyword: Virtio-net; VFs
	Reported in version: DOCA 1.8.0
369	Description: When restarting the virtio-net-controller from the DPU while the guest OS is booting, the guest OS may see kernel call trace while the controller is preparing the device. It recovers once the controller starts.
440	Workaround: N/A
2	Keyword: Virtio-net; hotplug; restart
	Reported in version: DOCA 1.8.0
363	Description: Jumbo MTU is only supported on a guest OS with kernel 4.11 and above.
345	Workaround: N/A
3	Keyword: Virtio-net; jumbo MTU
	Reported in version: DOCA 1.7.0
302	Description: When rebooting a DPU with a large number of VFs created on host, VF recovery may fail due to timeout.
196	Workaround: Restart the driver on the host after the DPU is up.
7	Keyword: Reboot; VFs
	Reported in version: DOCA 1.7.0

Ref #	Issue
323 244 4	Description: After live migration of virtio-net devices using the VFE driver, the <code>max_queues_size</code> output from the <code>virtnet list</code> may be wrong. This does not affect the actual value.
	Workaround: N/A
	Keywords: Virtio-net; live migration
	Reported in version: DOCA 1.4.0
280 178 0	Description: When running virtio-net-controller with host kernel older than 3.10.0-1160.el7, host virtio driver may get error (<code>Unexpected TXQ (13) queue failure: -28</code>) from dmesg in traffic stress test.
	Workaround: N/A
	Keywords: Virtio-net; error
	Reported in version: DOCA 1.2.0
287 021 3	Description: Servers do not recover after configuring <code>PCI_SWITCH_EMULATION_NUM_PORT</code> to 32 followed by power cycle.
	Workaround: Clear NVRAM and reset mlxconfig to default
	Keywords: Virtio-net; power cycle
	Reported in version: DOCA 1.2.0
268 519 1	Description: Once virtio-net is enabled, the mlx5 Windows VF becomes unavailable.
	Workaround: N/A
	Keywords: Virtio-net; virtual function; WinOF-2
	Reported in version: DOCA 1.2.0
270 239 5	Description: When a device is hot-plugged from the virtio-net controller, the host OS may hang when warm reboot is performed on the host and Arm at the same time.
	Workaround: Reboot the host OS first and only then reboot DPU.
	Keywords: Virtio-net controller; hot-plug; reboot
	Reported in version: DOCA 1.2.0

Bug Fixes

Ref #	Issue Description
397 489 3	Description: VLAN traffic does not work in virtio interface because <code>rq_attr.vlan_strip_disable</code> is set to 0 by default, stripping the VLAN tag a packet arrives at the virtio RQ.
	Keyword: VLAN
	Fixed in version: 24.10
393 643 5	Description: After changing uplink MTU to more than 1500, errors are printed from the virtio-net-controller side when using the vHost Acceleration Software Stack.
	Keyword: Virtio-net; vhost; live migration
	Fixed in version: 24.07
393 359 2	Description: When FLR times out, virtnet commands begin to hang and not return.
	Keyword: FLR; commands
	Fixed in version: 24.07

Notice
This document is provided for information purposes only and shall not be regarded as a warranty of a certain functionality, condition, or quality of a product. NVIDIA Corporation (“NVIDIA”) makes no representations or warranties, expressed or implied, as to the accuracy or completeness of the information contained in this document and assumes no responsibility for any errors contained herein. NVIDIA shall have no liability for the consequences or use of such information or for any infringement of patents or other rights of third parties that may result from its use. This document is not a commitment to develop, release, or deliver any Material (defined below), code, or functionality. NVIDIA reserves the right to make corrections, modifications, enhancements, improvements, and any other changes to this document, at any time without notice. Customer should obtain the latest relevant information before placing orders and should verify that such information is current and complete. NVIDIA products are sold subject to the NVIDIA standard terms and conditions of sale supplied at the time of order acknowledgement, unless otherwise agreed in an individual sales agreement signed by authorized representatives of NVIDIA and customer (“Terms of Sale”). NVIDIA hereby expressly objects to applying any customer general terms and conditions with regards to the purchase of the NVIDIA product referenced in this document. No contractual obligations are formed either directly or indirectly by this document. NVIDIA products are not designed, authorized, or warranted to be suitable for use in medical, military, aircraft, space, or life support equipment, nor in applications where failure or malfunction of the NVIDIA product can reasonably be expected to result in personal injury, death, or property or environmental damage. NVIDIA accepts no liability for inclusion and/or use of NVIDIA products in such equipment or applications and therefore such inclusion and/or use is at customer’s own risk. NVIDIA makes no representation or warranty that products based on this document will be suitable for any specified use. Testing of all parameters of each product is not necessarily performed by NVIDIA. It is customer’s sole responsibility to evaluate and determine the applicability of any information contained in this document, ensure the product is suitable and fit for the application planned by customer, and perform the necessary testing for the application in order to avoid a default of the application or the product. Weaknesses in customer’s product designs may affect the quality and reliability of the NVIDIA product and may result in additional or different conditions and/or requirements beyond those contained in this document. NVIDIA accepts no liability related to any default, damage, costs, or problem which may be based on or attributable to: (i) the use of the NVIDIA product in any manner that is contrary to this document or (ii) customer product designs. No license, either expressed or implied, is granted under any NVIDIA patent right, copyright, or other NVIDIA intellectual property right under this document. Information published by NVIDIA regarding third-party products or services does not constitute a license from NVIDIA to use such products or services or a warranty or endorsement thereof. Use of such information may require a license from a third party under the patents or other intellectual property rights of the third party, or a license from NVIDIA under the patents or other intellectual property

rights of NVIDIA.

Reproduction of information in this document is permissible only if approved in advance by NVIDIA in writing, reproduced without alteration and in full compliance with all applicable export laws and regulations, and accompanied by all associated conditions, limitations, and notices.

THIS DOCUMENT AND ALL NVIDIA DESIGN SPECIFICATIONS, REFERENCE BOARDS, FILES, DRAWINGS, DIAGNOSTICS, LISTS, AND OTHER DOCUMENTS (TOGETHER AND SEPARATELY, "MATERIALS") ARE BEING PROVIDED "AS IS." NVIDIA MAKES NO WARRANTIES, EXPRESSED, IMPLIED, STATUTORY, OR OTHERWISE WITH RESPECT TO THE MATERIALS, AND EXPRESSLY DISCLAIMS ALL IMPLIED WARRANTIES OF NONINFRINGEMENT, MERCHANTABILITY, AND FITNESS FOR A PARTICULAR PURPOSE. TO THE EXTENT NOT PROHIBITED BY LAW, IN NO EVENT WILL NVIDIA BE LIABLE FOR ANY DAMAGES, INCLUDING WITHOUT LIMITATION ANY DIRECT, INDIRECT, SPECIAL, INCIDENTAL, PUNITIVE, OR CONSEQUENTIAL DAMAGES, HOWEVER CAUSED AND REGARDLESS OF THE THEORY OF LIABILITY, ARISING OUT OF ANY USE OF THIS DOCUMENT, EVEN IF NVIDIA HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. Notwithstanding any damages that customer might incur for any reason whatsoever, NVIDIA's aggregate and cumulative liability towards customer for the products described herein shall be limited in accordance with the Terms of Sale for the product.

Trademarks

NVIDIA and the NVIDIA logo are trademarks and/or registered trademarks of NVIDIA Corporation in the U.S. and other countries. Other company and product names may be trademarks of the respective companies with which they are associated.

© Copyright 2025, NVIDIA. PDF Generated on 05/05/2025