



Intel® Ethernet Controller Products

27.7 Release Notes

NEX Cloud Networking Group (NCNG)

November 2022

Revision 1.0
743361-001

Revision History

Revision	Date	Comments
1.0	November 2022	Initial release.



1.0 Overview

This document provides an overview of the changes introduced in the latest Intel® Ethernet Controller/Adapter family of products. References to more detailed information are provided where necessary. The information contained in this document is intended as supplemental information only; it should be used in conjunction with the documentation provided for each component.

These release notes list the features supported in this software release, known issues, and issues that were resolved during release development.

1.1 New Features

1.1.1 Hardware Support

Release	New Hardware Support
27.7	<ul style="list-style-type: none">None for this release.

1.1.2 Software Features

Release	New Software Support
27.7	<ul style="list-style-type: none">Support for Microsoft* Windows* 11 22H2Microsoft* Azure Stack HCI, new feature releaseSupport for FreeBSD* 13.1 on most devices (not supported on Intel® Ethernet E820 Series devices)On Intel® Ethernet 700 Series devices:<ul style="list-style-type: none">i40e Linux driver allows speed adjustment with ethtool.FEC configuration through UEFI HII for 25G adaptersPruning of packets on bonded interfaces might be disabled.On Intel® Ethernet 810 Series devices:<ul style="list-style-type: none">Support for Debug Dump feature on Windows*, FreeBSD*, and ESXi*Support for Linux* Real Time Kernel 5.x and 4.xLinux support for RDMA RoCEv2 link aggregation in dual-port active-backup configurationICE driver support for GTP via switchdev and tc - OOTICE driver support for VXLAN/GRE in switchdev modeMiscellaneous:<ul style="list-style-type: none">The <i>auxiliary.ko</i> filename has been changed to <i>intel_auxiliary.ko</i>. This was done to avoid potential conflicts with the kernel and other vendors. This change affects the following drivers:<ul style="list-style-type: none">ice, version 1.9.7.1,i40e, version 2.20.10iavf, version 4.6.0rdma, version 1.10.5 <p>Please update these drivers to the latest version available.</p>

1.1.3 Removed Features

Release	Hardware/Feature Support
27.7	<ul style="list-style-type: none">None for this release.



1.1.4 Firmware Features

Release	New Firmware Support
27.7	<ul style="list-style-type: none"> None for this release



1.2 Supported Intel® Ethernet Controller Devices

Note: **Bold Text** indicates the main changes for this release.

For help identifying a specific network device as well as finding supported devices, click here:

<https://www.intel.com/content/www/us/en/support/articles/000005584/network-and-i-o/ethernet-products.html>

Note: Intel® ESXi drivers are available from VMware.

1.3 NVM

Table 1 shows the NVM versions supported in this release.

Table 1. Current NVM

Driver	NVM Version
810 Series	
E810	4.10
820 Series	
E822-C	2.28.1
E822-L	2.28.1
E823	2.28.1
700 Series	
700	9.10
500 Series	
X550	3.6
200 Series	
I210	2.0

Note: The NVM images for the E822-C, E822-L, and E823 are available to download from the Intel Resource & Documentation Center (RDC). An NDA is required.



1.4 Operating System Support

1.4.1 Levels of Support

The next sections use the following notations to indicate levels of support.

- Full Support = FS
- Not Supported = NS
- Inbox Support Only = ISO
- Supported Not Tested = SNT
- Supported by the Community = SBC

1.4.2 Linux

Table 2 shows the Linux distributions that are supported in this release and the accompanying driver names and versions.

Refer to [Section 1.4.1](#) for details on Levels of Support.

Table 2. Supported Operating Systems: Linux

Driver	Red Hat* Enterprise Linux* (RHEL) 9.0 & 8.6	RHEL 8.x (8.5 and previous)	SUSE* Linux Enterprise Server (SLES) 15 SP4	SLES 15 SP3 and previous	SLES 12 SP5	SLES 12 SP4 and previous	Canonical* Ubuntu* 22.04 LTS	Canonical Ubuntu 20.04 LTS	Debian* 11
Intel® Ethernet 810 Series									
ice	1.10.1.2	SNT	1.10.1.2	SNT	1.10.1.2	SNT	1.10.1.2	1.10.1.2	1.10.1.2
Intel® Ethernet 820 Series									
ice	1.10.1.2	SNT	1.10.1.2	SNT	1.10.1.2	SNT	1.10.1.2	SNT	SNT
Intel® Ethernet 700 Series									
i40e	2.21.12	SNT	2.21.12	SNT	2.21.12	SNT	2.21.12	2.21.12	SNT
Intel® Ethernet Adaptive Virtual Function									
iavf	4.6.1	SNT	4.6.1	SNT	4.6.1	SNT	4.6.1	4.6.1	SNT
Intel® Ethernet 10 Gigabit Adapters and Connections									
ixgbe	5.17.1	SNT	5.17.1	SNT	5.17.1	SNT	5.17.1	5.17.1	SNT
ixgbev	4.16.5	SNT	4.16.5	SNT	4.16.5	SNT	4.16.5	4.16.5	SNT
Intel® Ethernet Gigabit Adapters and Connections									
igb	5.12.3	SNT	5.12.3	SNT	5.12.3	SNT	5.12.3	5.12.3	SNT
Remote Direct Memory Access (RDMA)									
irdma	1.10.15	SNT	1.10.15	SNT	1.10.15	SNT	1.10.15	1.10.15	SNT

1.4.3 Windows Server

Table 3 shows the versions of Microsoft Windows Server that are supported in this release and the accompanying driver names and versions.

Refer to [Section 1.4.1](#) for details on Levels of Support.

Table 3. Supported Operating Systems: Windows Server

Driver	Microsoft Windows Server 2022	Microsoft Windows Server 2019	Microsoft Windows Server 2016	Microsoft Windows Server 2012 R2	Microsoft Windows Server 2012
Intel® Ethernet 810 Series					
icea	1.12.148.x	1.12.148.x	1.12.148.x	NS	NS
Intel® Ethernet 820 Series					
scea	1.1.285.x	1.1.285.x	NS	NS	NS
Intel® Ethernet 700 Series					
i40ea	1.16.206.x	1.16.202.x	1.16.202.x	1.16.202.x	1.16.62.x
i40eb	1.16.206.x	1.16.202.x	1.16.202.x	1.16.202.x	NS
Intel® Ethernet Adaptive Virtual Function					
iavf	1.13.8.x	1.13.8.x	1.13.8.x	1.13.8.x	NS
Intel® Ethernet 10 Gigabit Adapters and Connections					
ixe	NS	NS	NS	NS	2.4.36.x
ixn	NS	4.1.239.x	4.1.239.x	3.14.214.x	3.14.206.x
ixs	4.1.248.x	4.1.246.2	4.1.246.x	3.14.223.x	3.14.222.x
ixt	NS	4.1.228.x	4.1.229.x	3.14.214.x	3.14.206.x
sxa	4.1.248.x	4.1.243.x	4.1.243.x	3.14.222.x	3.14.222.x
sxb	4.1.248.x	4.1.239.x	4.1.239.x	3.14.214.x	3.14.206
vxn	NS	2.1.241.x	2.1.243.x	1.2.309.x	1.2.309.x
vxv	2.1.246.x	2.1.230.x	2.1.232.x	1.2.254.x	1.2.254.x
Intel® Ethernet 2.5 Gigabit Adapters and Connections					
e2f	1.1.3.28	1.1.3.28	NS	NS	NS

Table 3. Supported Operating Systems: Windows Server [continued]

Driver	Microsoft Windows Server 2022	Microsoft Windows Server 2019	Microsoft Windows Server 2016	Microsoft Windows Server 2012 R2	Microsoft Windows Server 2012
Intel® Ethernet Gigabit Adapters and Connections					
e1c	NS	NS	12.15.31.x	12.15.31.x	12.15.31.x
e1d	12.19.2.45	12.19.2.x	12.18.9.x	12.17.8.x	12.17.8.x
e1e	NS	NS	NS	NS	9.16.10.x
e1k	NS	NS	NS	NS	12.10.13.x
e1r	13.0.13.x	12.18.13.x	12.16.5.x	12.16.5.x	12.14.8.x
e1s	12.16.16.x	12.15.184.x	12.15.184.x	12.13.27.x	12.13.27.x
e1y	NS	NS	NS	NS	10.1.17.x
v1q	NS	1.4.7.x	1.4.7.x	1.4.5.x	1.4.5.x

1.4.4 Windows Client

Table 4 shows the versions of Microsoft Windows that are supported in this release and the accompanying driver names and versions.

Refer to [Section 1.4.1](#) for details on Levels of Support.

Table 4. Supported Operating Systems: Windows Client

Driver	Microsoft Windows 11	Microsoft Windows 10, version 1809 (and later)	Microsoft Windows 10	Microsoft Windows 8.1	Microsoft Windows 8
Intel® Ethernet 810 Series					
icea	1.12.148.x	NS	NS	NS	NS
Intel® Ethernet 700 Series					
i40ea	1.17.84.x	1.16.202.x	1.16.202.x	NS	NS
i40eb	1.17.80.x	1.16.202.x	NS	NS	NS
Intel® Ethernet Adaptive Virtual Function					
iavf	NS	NS	NS	NS	NS
Intel® Ethernet 10 Gigabit Adapters and Connections					
ixe	NS	NS	NS	NS	2.4.36.x
ixn	NS	4.1.239.x	4.1.239.x	3.14.214.x	3.14.206.x
ixs	4.1.248.x	4.1.246.x	4.1.246.x	3.14.223.x	3.14.222.x
ixt	NS	4.1.228.x	4.1.229.x	3.14.214.x	3.14.206.x
sxa	NS	4.1.243.x	4.1.243.x	3.14.222.x	3.14.222.x
sxb	NS	4.1.239.x	4.1.239.x	3.14.214.x	3.14.206.x
vxn	NS	NS	NS	NS	NS
vxv	NS	NS	NS	NS	NS
Intel® Ethernet 2.5 Gigabit Adapters and Connections					
e2f	2.1.1.x	1.1.3.x	NS	NS	NS

Table 4. Supported Operating Systems: Windows Client [continued]

Driver	Microsoft Windows 11	Microsoft Windows 10, version 1809 (and later)	Microsoft Windows 10	Microsoft Windows 8.1	Microsoft Windows 8
Intel® Ethernet Gigabit Adapters and Connections					
e1c	NS	NS	12.15.31.x	12.15.31.x	12.15.31.x
e1d	12.19.2.x	12.19.2.x	12.18.8.x	12.17.8.x	12.17.8.x
e1e	NS	NS	NS	NS	9.16.10.x
e1k	NS	NS	NS	NS	12.10.13.x
e1r	13.0.14.x	12.18.13.x	12.15.184.x	12.16.5.x	12.14.7.x
e1s	NS	12.15.184.x	12.15.184.x	12.13.27.x	12.13.27.x
e1y	NS	NS	NS	NS	10.1.17.x
v1q	NS	NS	1.4.7.x	NS	NS

1.4.5 FreeBSD

Table 5 shows the versions of FreeBSD that are supported in this release and the accompanying driver names and versions.

Refer to [Section 1.4.1](#) for details on Levels of Support.

Table 5. Supported Operating Systems: FreeBSD

Driver	FreeBSD 13	FreeBSD 12.3	FreeBSD 12.2 and previous
Intel® Ethernet 810 Series			
ice	1.37.7	1.37.7	SNT
Intel® Ethernet 820 Series			
ice	SNT	SNT	1.37.7
Intel® Ethernet 700 Series			
ixl	1.12.40	1.12.40	SNT
Intel® Ethernet Adaptive Virtual Function			
iavf	3.0.30	3.0.30	SNT
Intel® Ethernet 10 Gigabit Adapters and Connections			
ix	3.3.32	3.3.32	SNT
ixv	1.5.33	1.5.33	SNT
Intel® Ethernet Gigabit Adapters and Connections			
igb	2.5.24	2.5.24	SNT
Remote Direct Memory Access (RDMA)			
irdma	1.1.5	1.1.5	SNT
iw_ixl	0.1.30	0.1.30	SNT

2.0 Fixed Issues

2.1 Intel® Ethernet 800 Series Network Adapters

2.1.1 Intel® Ethernet 810 Series

2.1.1.1 General

- When RX/TX pause frames are enabled and the system is rebooted, or when the *icen* driver is unloaded, occasionally only one PF can be visible on ESX.
- Due to the previous bugs in PF-to-port mapping in both NVM and UEFI Driver, old NVMs are not compatible with the new UEFI driver. As it pertains to HII, NVMs can still be updated via FMP.

2.1.1.2 Linux Driver

- TC filters do not function correctly when the GPRS Tunnelling Protocol (GTP) encryption is enabled in the TC filter add command (**enc_key_id 1**) with *ice* driver version 1.10.X. This causes traffic to go to default TC 0.
- The sense of PFC is lost through a link bounce because the number of TCs is 1.
- Adding 16 MAC addresses to an untrusted VF that is attached to a VM, then changing the VF's trust status to on and off again, could cause the added MAC addresses to become non-functional.
- ADQ configuration commands that update settings on the interface (such as **tc qdisc add** and **ethtool -L**) cause the driver to close the associated RDMA interface and reopen it. This disrupts RDMA traffic on the interface. RDMA and ADQ configurations are not supported on the same interface simultaneously.
- Resetting VF device might lead to memory corruption. If VF is being reset with no intervals between resets, it is possible for the VF driver to free DMA mapping, while NIC still transmits packets. This will lead to memory corruption. To avoid this, limit rate resets of VF. Memory corruption will happen on the OS that VF is assigned to.
- Linux *ice* driver 1.2.1 cannot be compiled with E810 3.2 NVM images. The version on the kernel is 5.15.2.
- On RHEL8.5, **I2-fwd-offload** cannot be turned on.
- If the **rpmbuild** command of the new *iavf* version fails due to the existing auxiliary files installed, please use **--define "_unpackaged_files_terminate_build 0"** with the **rpmbuild** command.
- Usage/Workaround will look like **rpmbuild -tb iavf-4.4.0_rc53.tar.gz --define "_unpackaged_files_terminate_build 0" "**.
- When Double VLAN is created on a Virtual Machine, tx_tcp_cso [TX TCP Checksum Offload] and tx_udp_cso [TX UDP Checksum Offload] statistics might not increment correctly.
- On Intel® Ethernet Network Adapter E810-XXVDA4T, the driver might not link at 1000BASE-T and 1000BASE-X. The link might go down after advertising 1G.

2.1.1.3 Windows Driver

None for this release.

2.1.1.4 Linux RDMA Driver

- With Intel® Ethernet 800 Series Network Adapter iwarp, `ib_write_bw` fails with the a `RDMA_CM_EVENT_REJECTED` with at least 1000 QPs.
- When the ECN congestion control feature for rocev2 is enabled the code transmits the ECN codepoint as '00' in ACK packets instead of '10'. This marking causes fabric switches to drop RC ACK packets, drastically impacting performance when congestion is present and causing RC retries and aborted QP connections. This is the exact opposite of the intent of the feature and the subsequent retransmissions due to ACK drops makes congestion worse.

The issue is seen with multiple switch vendors.

2.1.1.5 NVM Update Tool

None for this release.

2.1.1.6 NVM

None for this release.

2.1.1.7 Firmware

- Per the specification, the **Get LLDP** command (0x28) response might contain only 2 TLVs (instead of 3).
- When software is requesting from firmware the port parameters on Port 0 (via AQ the connectivity type), the response is `BACKPLANE_CONNECTIVITY`, when it should be `CAGE_CONNECTIVITY`.
- Health status messages are not cleared with a PF reset, even after the reported issue is resolved.

2.1.1.8 Manageability

None for this release.

2.1.1.9 FreeBSD Driver

- When using 100G SR4 optics or active optical cables, it is possible for the adapter to fail to link if the `advertise_speed` sysctl is set to only link at 50G speeds.

2.1.1.10 VMware Driver

- On VMware ESXi vSphere 6.7 and Linux VM with Kernel version higher than 5.14, after performing warm reboot, traffic on SR-IOV VF adapters will be dropped.
 - **Workaround:** reload VF driver when issue occurs or instead of warm reboot perform power cycle from ESXi hypervisor.

2.1.1.11 Application Device Queues (ADQ)

None for this release.

2.1.1.11.1 ADQ Standard Issues

None for this release.

2.1.1.11.2 ADQ Configuration Script

None for this release



2.1.1.11.3 ADQ VF

- When ADQ is enabled on VFs, TC filters on the VF TCO (default TC) are not supported and will not pass traffic. It is not expected to add TC filters to TCO since it is reserved for non-filtered default traffic.
- VF Max Transmit Rate: TC qdisc add command on a VF interface does not verify that max_rate value(s) for the TCs are specified in increments of 500 Kbps. TC max_rate is expected to be a multiple of (or equal to) 500 Kbps.
- VF Max Transmit Rate: When ADQ is enabled on a VF interface, the tc qdisc add command causes the VF connection (ping) to drop when using *ice-1.8.X* and *iavf-4.4.X*.
 - **Workaround:** Use previous versions of the ice and iavf drivers for ADQ VF (*ice-1.7.16* with *iavf-4.3.19*).
- VF Max Transmit Rate: When a maximum TX transmit rate is specified in the **tc qdisc add** command on a VF interface, the maximum rate does not get applied correctly, causing an inconsistent TX rate limit for some applications.
- Configuring ADQ traffic classes with an odd number of hardware queues on a VF interface might result in a system hang in the *iavf* driver.
 - **Workaround:** To specify an even number of queues in the TC qdisc add the **dev** command for ADQ.

2.1.2 Intel® Ethernet 820 Series

2.1.2.1 General

- FW design issue caused the device to apply `LLDP enabled` property only to one port - usually port #0.

On some cards, enabling LLDP on non-0 ports might cause the device to disappear from the **NetworkDevices** collection after reboot.

All issues were fixed in Columbiaville 4.1 PC3 release.
- In addition to valid configurations, the HII Menu also lists invalid and extra configurations.

2.1.2.2 Firmware

- No link displayed on the 827 QSFP after setting to 50G.
- Asymmetrical link flow control is not functioning as expected with Tx=ON and Rx=OFF.

Works fine with Flow control set to TX=ON and RX=ON.
- When performing an NVM Update/inventory for a device running in recovery mode, the system returns **Exitcode 8** (No access to flash) instead of **Exitcode 0**. Even if the wrong Exitcode is observed, that device can still be initialized, perform updates, and exit from recovery mode.
- Current hardware and firmware limits the SFP cages to 1W (Power class 1). As a result, all optical modules that require power level 2 (1.5W) will not work with the ICX-D HCC Beta LEK images.

2.1.2.3 Linux Driver

- If a Virtual Function (VF) is not in trusted mode and eight or more VLANs are created on one VF, the VLAN that is last created might be non-functional and an error might be seen in dmesg.
- VXLAN traffic is not passing through on the expected queues.

- In DCB mode, when switching from IEEE to CEE mode, configurations are not applied.
 - **Workaround:** Either start in CEE mode or enter unwilling mode to perform configurations.
- When Double VLAN Mode is enabled on the host, disabling then re-enabling a Virtual Function attached to a Windows guest might cause error messages to be shown in **dmesg**. This behavior is not expected to affect functionality.
- Setting the MAC address of a VF to 00:00:00:00:00:00 from the host should reset the VF, not change the MAC address. We have observed the MAC address to appear as 00:00:00:00:00:00 from the host and the correct value from the VM.
- During RSS testing, incorrect values are displayed in OIDs for Maximum Number of RSS Queues and Processors.
- When all of the interfaces configured for VLAN traffic are connected to a vmSwitch port, and after reattaching one of the vmNICs to another port, it loses connectivity (i.e., cannot ping). The ability to ping is recovered after reattaching another vmNICs to the vmSwitch (where first vmNIC was attached).
- After a bond is created on two interfaces and the primary is disabled, the next "primary/secondary" will gain the MAC of the bond interface. The MAC of the bond always stays the same.

2.1.2.4 Windows RDMA Driver

- After enabling QoS for iWarp, no RDMA traffic is observed. This issue is not seen on RoCEv2.

2.1.2.5 NVM Update Tool

- If a device goes into recovery mode, a stripped image must be used to perform an update.

2.2 Intel® Ethernet 700 Series Network Adapters

2.2.1 General

None for this release.

2.2.2 VMware Driver

- When using a driver in ENS polling mode along with SR-IOV, packets might not come out from VF virtual machine. The issue appears after a PF reset or link down/up execution. [1](#)
 - **Workaround:** Reboot VM or reboot host
- When using driver in ENS polling mode along with SR-IOV, there might be no traffic between VF and VM adapters on single host. [1](#)
 - **Workaround:** None
- When using driver in ENS interrupt mode along with SR-IOV, there might be no traffic between VF and VM adapters with default MTU setting.
 - **Workaround:** Reboot VM or change MTU to any value, then change back to default.

2.2.3 Linux Driver:

- Intel® Ethernet Flow Director cloud filters: Outer MAC, Tenant ID, Inner MAC might not direct traffic to queues/VFs on the latest "Red Hat Enterprise Linux 9.0 (Plow)" 5.14.0-70.13.1.el9_0.x86_64.



2.2.4 Intel® PROSet:

None for this release.

2.2.5 EFI Driver

None for this release.

2.2.6 NVM

None for this release.

2.2.7 Firmware

None for this release

2.2.8 Windows Driver

None for this release.

2.2.9 Intel® Ethernet Flash Firmware Utility:

None for this release.

2.3 Intel® Ethernet 500 Series Network Adapters

None for this release.

2.4 Intel® Ethernet 300 Series Network Adapters

None for this release.

2.5 Intel® Ethernet 200 Series Network Adapters

None for this release.

3.0 Known Issues

3.1 Intel® Ethernet 800 Series Network Adapters

3.1.1 Intel® Ethernet 810 Series

3.1.1.1 General

- **Important:** For Intel® Ethernet Network Adapter E810-XXVDA4T users, **do not update** the NVM at this time.

If you must update the NVM, a power cycle reset to re-establish the link is required.

Note: Updating to the Release 27.7 drivers is supported and recommended.

- Intel's validation team found an issue in Windows Server 21H1. This OS version is unable to save a memory dump (crash dump) on disk. It is considered to be OS defect.
- When performing NVM Update/inventory for a device running into recovery mode, it returns with Exitcode 8 (No access to flash) instead of Exitcode 0. Even if the wrong exit code is observed, please keep in mind that the device can still be initialized, perform update, and exit from recovery mode.
- The Input-Output Memory Management Unit (IOMMU) feature of the processor prevents I/O devices from accessing memory outside the boundaries set by the OS. It also allows devices to be directly assigned to a Virtual Machine. However, IOMMU might affect performance, both in latency (each DMA access by the device must be translated by the IOMMU) and in CPU utilization (each buffer assigned to every device must be mapped in the IOMMU).

If you experience significant performance issues with IOMMU, try adding the following to the kernel boot command line:

intel_iommu=off

noiommu:

echo 1 > /sys/module/vfio/parameters/enable_unsafe_noiommu_mode

- Properties that can be modified through the manageability sideband interface **PLDM Type 6: RDE**, such as **EthernetInterface->AutoNeg** or **NetworkPort->FlowControlConfiguration** do not possess a permanent storage location on internal memory. Changes made through RDE are not preserved following a power cycle/PCI reset.
- Link issues (for example, false link, long time-to-link (TTL), excessive link flaps, no link) might occur when the Intel® Ethernet Connection C827 Series and Intel® Ethernet Connection XL827 Series (C827/XL827) retimer is interfaced with SX/LX, SR/LR, SR4/LR4, AOC limiting optics. This issue is isolated to C827/XL827 line side PMD RX susceptibility to noise.
- Intel® Ethernet 800 Series Network Adapters in 4x25GbE or 8x10GbE configurations are limited to a maximum total transmit bandwidth of roughly 28Gbps per port for 25GbE ports and 12Gbps per port on 10GbE ports.

This maximum is a total combination of any mix of network (leaving the port) and loopback (VF -> VF/VF -> PF/PF -> VF) TX traffic on a given port and is designed to allow each port to maintain port speed transmit bandwidth at the specific port speed when in 25GbE or 10GbE mode.

If the PF is transmitting traffic as well as the VF(s), under contention the PF has access to up to 50% TX bandwidth for the port and all VFs have access to 50% bandwidth for the port, which will also impact the total available bandwidth for forwarding.



Note: When calculating the maximum bandwidth under contention for bi-directional loopback traffic, the number of TX loopback actions are twice that of a similar unidirectional loopback case, since both sides are transmitting.

- The version of the Ethernet Port Configuration Tool available in Release 26.1 might not work as expected. This has been resolved in Release 26.4.
- The E810 currently supports a subset of 1000BASE-T SFP module types, which use SGMII to connect back to the E810. In order for the E810 to properly know the link status of the module's BASE-T external connection, the module must indicate the BASE-T side link status to the E810. An SGMII link between the E810 and the 1000BASE-T SFP module allows the module to indicate its link status to the E810 using SGMII Auto Negotiation. However 1000BASE-T SFP modules implement this in a wide variety of ways, and other methods which do not use SGMII are currently unsupported in the E810. Depending on the implementation, link might never be achieved. In other cases, if the module sends IDLEs to the E810 when there is no BASE-T link, the E810 might interpret this as a link partner sending valid data and might show link as being up even though it is only connected to the module and there is no link on the module's BASE-T external connection.
- If the PF has no link, then a Linux VM previously using a VF will not be able to pass traffic to other VMs without the patch found here.

<https://lore.kernel.org/netdev/BL0PR2101MB093051C80B1625AAE3728551CA4A0@BL0PR2101MB0930.namprd21.prod.outlook.com/T/#m63c0a1ab3c9cd28be724ac00665df6a82061097d>

This patch routes packets to the virtual interface.

Note: This is a permanent 3rd party issue. No expected action on the part of Intel.

- Some devices support auto-negotiation. Selecting this causes the device to advertise the value stored in its NVM (usually disabled).
- VXLAN switch creation on Windows Server 2019 Hyper V might fail.
- Intel does its best to find and address interoperability issues, however there might be connectivity issues with certain modules, cables or switches. Interoperating with devices that do not conform to the relevant standards and specifications increases the likelihood of connectivity issues.
- When priority or link flow control features are enabled, traffic at low packet rates might increment priority flow control and/or packet drop counters.
- For an Intel® Ethernet 800 Series Network Adapter-based adapter to reach its full potential, users must install it in a PCIe v4 x16 slot. Installing on fewer lanes (x8, x4, x2) and/or Gen3, Gen2 or Gen1, impedes the full throughput of the device.
- On certain platforms, the legacy PXE option ROM boot option menu entries from the same device are pre-pended with identical port number information (first part of the string that comes from BIOS).

This is not an option ROM issue. The first device option ROM initialized on a platform exposes all boot options for the device, which is misinterpreted by BIOS.

The second part of the string from the option ROM indicates the correct slot (port) numbers.

3.1.1.2 Firmware

- Promiscuous mode does not see all packets: it sees only those packets arriving over the wire (that is, not sent from the same physical function (PF) but a different virtual function (VF)).

3.1.1.3 Linux

- When attempting to reduce the combined Tx/Rx queues such as:

ethtool -L 8

a deadlock can happen between probing irdma and the ethtool command with hung task messages in the log.

— Workaround

Insure the queue reduction is complete without having irdma loaded.

- When the number of queues on the VF interface is changed using the **ethtool** command while traffic is flowing from client to SUT (using *ice* driver version 1.10.X and *iavf* version 4.6.X), the system reboots.
- Host system might hang when resuming from sleep or hibernate with active VFs.
Unexpected system hand or loss of network functionality can occur when waking from sleep or hibernation if VFs (Virtual Functions) are active. To restore normal operation, the system must be power cycled.
- LAN traffic is disabled between Windows VFs after changing TPID to 88a8.
- With the 810 Series 3.2 NVM in the Intel® Ethernet Network Adapter E810-CQDA2 card, if the 810 Series 2.2 *iavf* driver is installed, a fatal error is generated related to **pci-aspm.h**, and the installation fails.
- When Double VLAN is created on a Virtual Machine, **tx_tcp_cso [TX TCP Checksum Offload]** and **tx_udp_cso [TX UDP Checksum Offload]** statistics might not increment correctly.

3.1.1.4 Linux Driver

- Accelerated Receive Flow Steering (aRFS) might not work correctly with ice driver version 1.10.X, causing traffic to go to unexpected RX queues on the Intel® Ethernet Network Adapter E810 adapter.
- It might not be possible to create the maximum number of supported RDMA VFs. Attempting to create greater than 20 RDMA VFs will result in no RDMA devices being created for VFs.
- When attempting to reduce the combined Tx/Rx queues like:

```
ethtool -L 8
```

a deadlock can happen between probing irdma and the **ethtool** command, with hung task messages in the log.

— **Workaround:** Ensure the queue reduction is done without having *irdma* loaded.

- For Linux *ice* driver v1.7.16+ on E810 4-port SKUs, VF cannot be created on certain high-core count platform, because PF has exhausted all the MSIX interrupts.
- Double VLAN traffic might RSS into the first queue. If configuring VLAN interfaces on PF in a way that results in double VLAN tagging, received double VLAN packets will be concentrated on the first queue of interface.
- When using certain DDP package versions, 802.1ad type VLANs might not be correctly enabled on SIOV or SR-IOV interfaces.
- Repeatedly adding/deleting a VF from a namespace while also repeatedly changing its trust mode status can result in call trace after a significant number of iterations.



- When two VFs created from the same PF are assigned identical MAC addresses, they will not be able to pass traffic successfully unless VF spoof check is disabled on the VF interfaces.
- The Intel® Ethernet 800 Series Network Adapter in 8-port 10Gb configuration, the device might generate errors such as the example below on Linux PF or VF driver load due to RSS profile allocation. Ports that report this error will experience RSS failures resulting in some packet types not being properly distributed across cores.
 - **Workaround:** Disable RSS using the `--disable-rss` flag when starting DPDK. Afterwards, only enable the specific RSS profiles that are needed.

dmesg: VF add example

```
ice_add_rss_cfg failed for VSI:XX, error:ICE_ERR_AQ_ERROR
```

```
VF 3 failed opcode 45, retval: -5
```

DPDK v20.11 testpmd example:

```
Shutting down port 0...
```

```
Closing ports...
```

```
iavf_execute_vf_cmd(): No response or return failure (-5) for cmd 46
```

```
iavf_add_del_rss_cfg(): Failed to execute command of OP_DEL_RSS_INPUT_CFG
```

- When **spoofchk** is turn on, the VF device driver will have pending DMA allocations while it is released from the device.
- After changing link speed to 1G on the E810-XXVDA4, the PF driver cannot detect a link up on the adapter. As a workaround the user can force 1G on the second side.
- When using bonding mode 5 (i.e., balance-tlb or adaptive transmit load balancing), if you add multiple VFs to the bond, they are assigned duplicate MAC address. When the VFs are joined with the bond interface, the Linux bonding driver sets the MAC address for the VFs to the same value. The MAC address is based on the first active VF added to that bond. This results in balance-tlb mode not functioning as expected. PF interfaces behave as expected.
The presence of duplicate MAC addresses might cause further issues, depending on your switch configuration.
- Changing the FEC value from BaseR to RS results in an error message in **dmesg**, and might result in link issues.
- UEFI PXE installation of Red Hat Enterprise Linux 8.4 on a local disk results with the system failing to boot.
- If single VLAN traffic is active on a PF interface and a CORER or GLOBR reset is triggered manually, PF traffic will resume after the reset whereas VLAN traffic might not resume as expected. For a workaround, issue the ethtool command: **ethtool -K PF_devname rx-vlan-filter off** followed by **ethtool -K PF_devname rx-vlan-filter on** and VLAN traffic will resume.
- Adding a physical port to the Linux bridge might fail and result in Device or Resource Busy message if SR-IOV is already enabled on a given port. To avoid this condition, create SR-IOV VFs after assigning a physical port to a Linux bridge. Refer to *Link Aggregation is Mutually Exclusive with SR-IOV and RDMA* in the *ice* driver README.
- When using a Windows Server 2019 RS5 Virtual Machine on a RHEL host, a VLAN configured on the VF using **iproute2** might not pass traffic correctly when an *ice* driver older than version 1.3.1 is used in combination with a *iavf* driver version.
- It has been observed that when using iSCSI, the iSCSI initiator intermittently fails to connect to the iSCSI target.



- When the Double VLAN Mode is enabled on the host, disabling and re-enabling a Virtual Function attached to a Windows guest might cause error messages to be displayed in **dmesg**. These messages do not affect functionality.
- With the current *ice* PF driver, there might not be a way for a trusted DPDK VF to enable unicast promiscuous without turning on "ethtool --priv-flags" with vf-true-promisc-support."
- If a VLAN with an Ethertype of 0x9100 is configured to be inserted into the packet on transmit, and the packet, prior to insertion, contains a VLAN header with an Ethertype of 0x8100, the 0x9100 VLAN header is inserted by the device after the 0x8100 VLAN header. The packet is transmitted by the device with the 0x8100 VLAN header closest to the Ethernet header.
- A PCI reset performed on the host might result in traffic failure on VFs for certain guest operating systems.
- On RHEL 7.x and 8.x operating systems, it has been observed that the `rx_gro_dropped` statistic might increment rapidly when Rx traffic is high. This appears to be an issue with the RHEL kernels.
- When *ice* interfaces are part of a bond with `arp_validate=1`, the backup port link status flaps between up and down.
 - **Workaround:** It is recommended to not enable `arp_validate` when bonding *ice* interfaces.
- Changing a Virtual Function (VF) MAC address when a VF driver is loaded on the host side might result in packet loss or a failure to pass traffic. As a result, the VF driver might need to be restarted.
- Current limitations of minimum Tx rate limiting on SR-IOV VFs:
 - If DCB or ADQ are enabled on a PF then configuring minimum Tx rate limiting on SR-IOV VFs on that PF is rejected.
 - If both DCB and ADQ are disabled on a PF then configuring minimum Tx rate limiting on SR-IOV VFs on that PF is allowed.
 - If minimum Tx rate limiting on a PF is already configured for SR-IOV VFs and a DCB or ADQ configuration is applied, then the PF can no longer guarantee the minimum Tx rate limits set for SR-IOV VFs.
 - If minimum Tx rate limiting is configured on SR-IOV VFs across multiple ports that have an aggregate bandwidth over 100Gbps, then the PFs cannot guarantee the minimum Tx rate limits set for SR-IOV VFs.
- Some distros might contain an older version of **iproute2/devlink** which might result in errors.
 - **Workaround:** Please update to the latest **devlink** version.

3.1.1.5 FreeBSD Driver

- At the end of an **NVMUpdate** or **NURA** tool execution, after the NVM has finished updating, the tool might indicate that an EMP reset failed—while the adapter continues to work normally and **NVMUpdate** indicates that the update was successful.

This is a known issue that is planned to be addressed in a future tools release.

- Due to changes in the reset path, several errors are seen in **dmesg** when trying to re-establish link after an **nvmupdate**. This does not effect functionality, as **nvmupdate** requires a reboot to be complete.



- The driver can be configured with both link flow control and priority flow control enabled even though the adapter only supports one mode at a time. In this case, the adapter will prioritize the priority flow control configuration. Verify that link flow control is active or not by checking the **active:** line in ifconfig.
- IAVF virtual interfaces in FreeBSD-13.0 guests might experience poor receive-performance during stress.
- Unable to ping after removing the primary NIC teaming adapter. The connection can be restored after restarting the VM adapters. This issue is not observed after the secondary adapter is removed, and is not OS specific.
- The visibility of the iSCSI LUN is dependent upon being able to establish a network connection to the LUN. In order to establish this connection, factors such as the initialization of the network controller, establishing link at the physical layer (which can take on the order of seconds) must be considered. Because of these variables, the LUN might not initially be visible at the selection screen.
- Intel® Ethernet Controller E810 devices are in the DCBX CEE/IEEE willing mode by default. In CEE mode, if an Intel® Ethernet Controller E810 device is set to non-willing and the connected switch is in non-willing mode as well, this is considered an undefined behavior.
 - **Workaround:** Configure Intel® Ethernet Controller E810 devices for the DCBX willing mode (default).
- In order to use guest processor numbers greater than 16 inside a VM, you might need to remove the *RssMaxProcNumber (if present) from the guest registry.

3.1.1.6 Windows RDMA Driver

- In heavy RDMA read traffic, some packets can be dropped and cause errors. To avoid that PFC needs to be configured with no-drop policy for RDMA traffic. The Intel® Ethernet Network Adapter E810 might experience an adapter-wide reset on all ports. When in firmware managed mode, a DCBx willing mode configuration change that is propagated from the switch removes a TC that was enabled by RDMA. This typically occurs when removing a TC associated with UP0 because it is the default UP on which RDMA based its configuration. The reset results in a temporary loss in connectivity as the adapter re-initializes.
- With a S2D storage cluster configuration running Windows Server 2019, high storage bandwidth tests might result in a crash for a BSOD bug check code 1E (KMODE_EXCEPTION_NOT_HANDLED) with `smbdirect` as the failed module. Customers should contact Microsoft via the appropriate support channel for a solution.

3.1.1.7 Linux RDMA Driver

- A usermode user verbs thread, if run concurrently with a probe of irdma, might panic as follows:

```
[11620.922765] RIP: 0010:irdma_dbg_save_ucontext+0x4e/0x60 [irdma]
```

There is linked list not properly initialized prior to the kernel device registration.

— Workaround

The issue happens most frequently with the **ibacm** service which is the IB path record caching service.

The service is not needed for irdma. If the panic shows:

```
Comm: ibacm
```

A detour is possible by either disabling the ibacm service or removing the package associated with ibacm.

- Previous to this release, the ToS-to-user priority mappings were hard coded in Linux to use only 0, 2, 4, 6 VLAN PFC priority values with RDMA traffic. Now in order to use other priorities with RDMA traffic (i.e. 1, 3, 5, 7), a VLAN is required to be setup using the `egress-qos-map` option. For example to map all priority 0 as priority 3, one can use: `ip link add link <ifname> name <vlan-ifname> type vlan id <vlan-id> egress-qos-map 0:3 1:0`
- The `ib_read_bw` micro benchmark can fail when two separate read sessions are sourced from a single port to two ports on the BSC FPGA device.

The commands are as follows:

— Server

- `Ib_read_bw -x <dev1-gid> -d <dev1> -s 4096 -n 5 -q 1 -t 128 -p 17516 &`
- `Ib_read_bw -x <dev2-gid> -d <dev2> -s 4096 -n 5 -q 1 -t 128 -p 17517`

— Client

- `Ib_read_bw -x <cvl1-gid> -d <cvl1> -s 4096 -n 5 -q 1 -t 128 -p 17516 <dev1-ip-addr> &`
- `Ib_read_bw -x <cvl1-gid> -d <cvl2> -s 4096 -n 5 -q 1 -t 128 -p 17516 <dev2-ip-addr> &`

Traces show the all RDMA read response only packets seem to be ignored by the CVL, resulting in repeated retries until the retry count is exhausted and terminates the two connections.

- **Workaround:** Avoid having the two sessions where the RDMA read response packets are coming from two separate ports.

- On RHEL 7.9, installing `rdma-core v35.0 debuginfo` rpms can prevent the installation of `debuginfo` rpms from indistro products like `libfabric`.
- When using Intel® MPI Library in Linux, Intel recommends to enable only one interface on the networking device to avoid MPI application connectivity issues or hangs. This issue affects all Intel® MPI Library transports, including TCP and RDMA. To avoid the issue, use `ifdown <interface>` or `ip link set down <interface>` to disable all network interfaces on the adapter except for the one used for MPI. OpenMPI does not have this limitation.
- A `rpings` can fail after a system reboot.

The issue is caused when the `modprobe` of `ib_core.ko` is the target of a signal at the point the `initramfs` switch root is done.

A signature of the issue is this log message:

```
[ 6.869539] infiniband irdma0: Couldn't open port 1
```

- **Workaround:** Adding a kernel boot line argument: `printk.devkmsg=on` might delay the `irdma` load until after the switch root.

3.1.1.8 NVM Update Tool

- On some adapters, after making the NVMUpdate we can observe that adapter will be in pending reboot state. POR is needed to exit this state.
- Updating using an external OROM (FLB file) and opting for delayed reboot in the configuration file is not supported.



3.1.1.9 VMware Driver

- When instantiating the maximum number of VFs in NSX-T, adding a Transport Node afterwards might fail due to timeout.
- Configuring the NSX-T Virtual Distributed Switch uplink port might fail when SR-IOV is enabled in the PF.
- When configured for NSX-T, ENS, and VXLAN, heavy traffic might cause pings to fail. Setting LFC for PF might fail.
- Excessive VF reset might cause multicast ping from VF to fail.
- TCP traffic on VMs might be interrupted if PF undergoes a reset.
- The output of the PTP signal might have an incorrect period than the one configured.
- Received packets with incorrect length can generate alarms in VMware ESXi. These alarms can be ignored. Please following article for more details: <https://kb.vmware.com/s/article/83627>
- When configuring a switch to use IEEE LLDP version for DCB, the PF host driver is unable to change the CEE LLDP version for DCB, even if the switch is configured for CEE LLDP.
- When running Release 2.2 NVM drivers on Release 3.2 NVMs, users might encounter warning messages regarding Null pointer errors. These are expected warnings when running older drivers on newer NVMs.
- When entering the Pause Parameter via the CLI, related configurations in quick or rapid succession could cause a configuration failure or unexpected results. In NSX-T 3.1.0, a Guest Virtual Machine associated with ENS NSX-T Virtual Distributed Switch (NVDS) might experience guest operating system kernel panic when receiving TCP traffic with VXLAN overlay. VMware ESX 7.0 operating system with NSX-T 3.1.0 might experience a kernel panic (also known as PSOD) when changing NUMA node in NSX-T Virtual distributed switch. Rebooting a Red Hat 8.2 Linux VF VM multiple times might cause traffic to stop on that VF. A VLAN tag is not inserted automatically when DCB PFC is enabled on an interface. This might cause RDMA issues if no VLAN is configured.
- **Workaround:** Since PFC for icen is VLAN-based, create a VLAN tag for DCB to be fully operational. After a PF Reset, Windows VF traffic might fail.

3.1.1.10 VMware RDMA Driver

- Running Unreliable Datagram (UD) RDMA mixed traffic with more than two QPs might lead to a receiver side UD application hang. To recover, restart the RDMA UD application. This is not expected to impact storage (NVMeoF, iSER, VSAN) applications because they do not rely on UD communication.
- Before loading a non-default DDP package, all RDMA traffic must be stopped and the *irdma* driver must be unloaded on the system.
If this recommendation is not followed, the following failures might occur:
 - Loading a DDP package during RDMA traffic might lead to system hang that requires a server reset to recover.
 - Loading a DDP package with RDMA enabled (without RDMA traffic running) might fail, and the device can become unusable for RDMA traffic until recovered by reboot.
- For RDMA, VMware recommends lossless traffic to be configured on priority 3, but due to a known issue in E810 adapters it requires priority '2' to be configured. The *irdma* driver will automatically translate this to priority '3' in the ESXi OS. If other priority settings are used, the priority seen in the packet might be different than DCB setting on the NIC and RDMA traffic will not be properly configured as lossless.

3.1.1.11 Application Device Queues (ADQ)

The code contains the following known issues:

ADQ Standard Issues

- After clearing the ADQ VF configuration, not all queues are available. As a result, the ADQ VF reconfiguration fails.
 - **Workaround:** Add at least 2 seconds of sleep time to restore all queues before reconfiguring the ADQ.
 - Creating more than 10k TC filters on an interface can result in errors talking to the kernel and the filters fail to get created (maximum number of supported tc filters is 32k).
 - The latest RHEL and SLES distros have kernels with back-ported support for ADQ. For all other OS distros, you must use the LTS Linux kernel v4.19.58 or higher to use ADQ. The latest out-of-tree driver is required for ADQ on all Operating Systems.
 - ADQ configuration must be cleared following the steps outlined in the *Intel® Ethernet Controller E810 Application Device Queues (ADQ) Configuration Guide*. The following issues can result if steps are not executed in the correct order:
 - Removing a TC qdisc prior to deleting a TC filter will cause the qdisc to be deleted from hardware and leave an unusable TC filter in software.
 - Deleting a ntuple rule after deleting the TC qdisc, then re-enabling ntuple, might leave the system in an unusable state which requires a forced reboot to clear.
 - Mitigation — Follow the steps documented in the *Intel® Ethernet Controller E810 Application Device Queues (ADQ) Configuration Guide* section "Clear the ADQ Configuration."
 - ADQ configuration is not supported on a bonded or teamed Intel® Ethernet Network Adapter E810 interface. Issuing the **ethtool** or **tc** commands to a bonded E810 interface will result in error messages from the *ice* driver to indicate the operation is not supported.
 - If the application stalls for some reason, this can cause a queue stall for application-specific queues for up to two seconds.
 - **Workaround** - Recommend configuration of only one application per Traffic Class (TC) channel.
 - DCB and ADQ are mutually exclusive and cannot coexist. A switch with DCB enabled might remove the ADQ configuration from the device.
 - **Workaround** - Do not enable DCB on the switch ports being used for ADQ. Disable LLDP on the interface by turning off firmware LLDP agent using:

```
ethtool --set-priv-flags $iface fw-lldp-agent off
```
- Note:** (unrelated to Intel drivers): The 5.8.0 Linux kernel introduced a bug that broke the interrupt affinity setting mechanism.
- **Workaround** - Use an earlier or later version of the kernel to avoid this error.
- Commands such as **tc qdisc add** and **ethtool -I** cause the driver to close the associated RDMA interface and reopen it. This disrupts RDMA traffic for 3-5 seconds until the RDMA interface is available again for traffic.
 - When the number of queues is increased using **ethtool -L**, the new queues will have the same interrupt moderation settings as queue 0 (i.e., Tx queue 0 for new Tx queues and Rx queue 0 for new Rx queues). This can be changed using the **ethtool** per-queue coalesce commands.
 - To fully release hardware resources and have all supported filter type combinations available, the *ice* driver must be unloaded and re-loaded.



- If a reset occurs on a PF interface containing TC filter(s), traffic does not resume to the TC filter(s) after the PF interface is restored.
- TC filters can unexpectedly match packets that use IP protocols other than what is specified as the **ip_proto** argument in the **tc filter add** command. For example, UDP packets might be matched on a TCP TC filter created with **ip_proto tcp** without any L4 port matches.
- ADQ does not work as expected with NVMe/TCP using Linux kernel v5.16.1 and later. When **NVMe connect** is issued on an initiator with kernel v5.16.1 (or later), a system hang might be observed on the host system. This issue is not specific to Intel® Ethernet drivers, it is related to NVMe changes in the 5.16 kernel. Issue can also be observed with older versions of the *ice* driver using a 5.16+ kernel.

ADQ Configuration Script

Refer to the *Intel® Ethernet 800 Series Network Adapters Application Device Queues (ADQ) Configuration Guide* “Known Issues” section for ADQ Configuration Script updates.

ADQ Virtual Function Issues

- The *iavf* driver must use Trusted mode with ADQ: Trusted mode must be enabled for ADQ inside a VF. If TC filters are created on a VF interface with trusted mode off, the filters are added to the software table but are not offloaded to the hardware.
- VF supports Max Transmit Rate only: the *iavf* driver only supports setting maximum transmit rate (**max_rate**) for Tx traffic. Minimum transmit rate (**min_rate**) setting is not supported with a VF.
- A core-level reset of an ADQ-configured VF port (rare events usually triggered by other failures in the NIC/*iavf* driver) results in loss of ADQ configuration. To recover, reapply ADQ configuration to the VF interface.
- VF errors occur when deleting TCs or unloading the *iavf* driver in a VF: *ice* and *iavf* driver error messages might get triggered in a VF when TCs are configured, and TCs are either manually deleted or the *iavf* driver is unloaded. Reloading the *ice* driver recovers the driver states.

3.1.1.12 Manageability

- Intel updated the E810 FW to align the sensor ID design as defined by DMTF DSP2054 starting from Release 26.4. Previous versions of the E810 FW were based on draft version of the specification. As a result updating to the newer NVM with this FW will result in updating numbering for the thermal sensorsIDs and PDR handlers. Anyone using hard coded values for these will see changes. A proper description of the system through PLDM type 2 PDRs shall give a BMC enough information to understand what sensors are available, what they are monitoring and what their ID is.

3.1.2 Intel® Ethernet 820 Series

3.1.2.1 General

- During validation, an issue was discovered in Windows Server 21H1. This OS version is unable to save a memory dump (crash dump) to a disk. It is considered to be OS defect.
- With a Windows host and Linux virtual machine (VM), the last transmit (Tx) queue might not increment when there are multiple Tx queues.
- **Insufficient PCI-Express bandwidth available for device** might be logged for Intel® Ethernet E820 Series Network Adapters. The E820 Series does not use a PCI-Express interface and this appears to only be a logging issue.
- The 4x25G NVM lists 100GbE and 50GbE link speeds in the device advanced tab on Windows Server.

- There is a lack of output from the `Get-NetQoSPolicy` command, even though iSCSI is working.
- On 82X platforms using **lanconf** in the EFI shell and EFI networking enabled, under the **EDKII Menu --> Platform Configuration --> Network Configuration**, the **EFI Network** option is disabled by default.

If this option is enabled, then **lanconf** in the EFI shell hangs and is unusable.

— **Workaround:** Disable the EFI Network option.

3.1.2.2 Firmware

- The DCB-MAP Configuration is not displaying on the SUT Interface from the Extreme Switch after enabling the firmware Mode in the SUT.
- After changing port options in EPCT and the system returns blank mode, perform a second reboot.
- The **Get LLDP** command (0x28) response contains only 2 TLV types. There should be a third TLV type.

3.1.2.3 Linux Driver

- Virtual Functions (VF) do not run on one of CPK Physical Functions (PF). The Single Root I/O Virtualization (SR-IOV) cannot be used on one of PFs.
 - In a double VLAN setup with set to promiscuous mode, packets are not seen in Wireshark on the expected ports.
 - After assigning a Locally Administered Address (LAA), the system can still wake from S5 by using the Burned In Address (BIA) but does not wake up if the LAA is used.
- **Workaround:** use BIA for waking the system from S5.
- On RHEL 7.9 VMs, VF traffic does not resume after the VF's MAC address is changed on the host side. This appears to be a limitation with RHEL 7.9.

There are two workarounds options to resume VF traffic. Only one has to be applied.

1. Manually set the MAC of the VF interface in the guest OS to match the one set on the host

— `$ ip link set <eth> mac <mac_set_on_the_host_side>`

2. Bring the link administratively down/up on the guest OS

— `$ ip link set <eth> down && ip link set <eth> up`

- DCB-MAP Configuration is not reflected from switch on Intel® Ethernet Connection C827 Series Port with CEE and SW Mode on SLES15 SP3 OS.
- **Module is not present** error message is displayed after loading the *ice* driver with cages filled.

3.1.2.4 FreeBSD Driver

- During successive driver unload/load cycles, single-port adapters might experience initialization failure when 5-layer topology is enabled.
- When a driver is loaded with an empty cage, an Admin Queue (AQ) error is recorded instead of the expected AHS link messages.
- The available memory decreases slightly when reloading driver. This should have minimal impact under normal use.
- Using FreeBSD, while receiving packets from client, the connection between the client and the system under test (SUT) fails after the reboot of the SUT.



3.1.2.5 Windows RDMA Driver

None for this release.

3.1.2.6 Linux RDMA Driver

None for this release.

3.1.2.7 NVM Update Tool

None for this release.

3.1.2.8 VMware Driver

None for this release.

3.1.2.9 Application Device Queues (ADQ)

None for this release.

3.1.2.10 Manageability

None for this release.

3.1.2.11 NVM

- The Single Root I/O Virtualization interface (SR-IOV) might fail when running with a 50G image on a 100G Si.
- The 100M option, is visible in Windows* Device Manager. However, when it is selected, a link cannot be established.
- Using the EPCT tool, after switching the QSFP PKVL VMC to 4x25, the driver can no longer attach interfaces. The only way to restore the device is to re-flash the NVM using LanConf.
- LAN traffic is disabled between peer VMs after disabling SR-IOV.

3.2 Intel® Ethernet 700 Series Network Adapters

3.2.1 General

- Devices based on the Intel® Ethernet Controller XL710 (4x10 GbE, 1x40 GbE, 2x40 GbE) have an expected total throughput for the entire device of 40 Gb/s in each direction.
- The first port of Intel® Ethernet Controller 700 Series Network Adapter display the correct branding string. All other ports on the same device display a generic branding string.
- In order for an Intel® Ethernet Controller 700 Series Network Adapter to reach its full potential, users must install it in a PCIe Gen3 x8 slot. Installing on fewer lanes (x4, x2) and/or Gen2 or Gen1, impedes the full throughput of the device.

3.2.2 Intel® Ethernet Controller V710-AT2/X710-AT2/TM4

- Incorrect *DeviceProviderName* is returned when using RDE *NegotiateRedfishParameters*. This issue has been root caused and the fix should be integrated in the next firmware release.



3.2.3 Windows Driver

None for this release.

3.2.4 Linux Driver

- **ptp4l** since v1.8 has supported monitoring link state, and will not reset after detecting a fault when the link status of the port is down. However, if `fault_reset_interval` is configured to ASAP (or its numeric equivalent of -128), **ptp4l** will attempt to reset immediately without checking the link status. This causes **ptp4l** to continuously fault and reset as the link for the port is still down. This results in a clear and unexpected behavioral difference when using ASAP vs when using another `fault_reset_interval`. A fix was proposed to the PTP4l Opensource project, but has not yet been accepted by the LinuxPTP community

3.2.5 Intel® PROSet

None for this release.

3.2.6 EFI Driver

None for this release.

3.2.7 VMware Driver

- On some hosts with an AMD CPU, there is no Tx traffic with multiple Intel® Ethernet Network Adapter I710 ports connected to one vSwitch.
 - **Workaround:** Do not connect multiple Intel® Ethernet Network Adapter I710 ports to one vSwitch.
- When using driver in ENS polling mode along with SR-IOV, packets might not come out from VF virtual machine. Issue appears after PF reset or link down/up execution.
 - **Workaround:** Reboot VM or reboot host
- When using driver in ENS polling mode along with SR-IOV there might be no traffic between VF and VM adapters on single host.
 - **Workaround:** None
- When using driver in ENS interrupt mode along with SR-IOV there might be no traffic between VF and VM adapters with default MTU setting.
 - **Workaround:** Reboot VM or change MTU to any value, then change back to default.

3.2.8 NVM

None for this release.

3.3 Intel® Ethernet 500 Series Network Adapters

3.3.1 General

None for this release.

3.3.2 EFI Driver

None for this release.

3.3.3 Windows Driver

None for this release.

3.4 Intel® Ethernet 300 Series Network Adapters

3.4.1 EFI Driver

None for this release.

3.4.2 VMware Driver

None for this release.

3.5 Intel® Ethernet 200 Series Network Adapters

None for this release.

3.6 Legacy Devices

Some older Intel® Ethernet adapters do not have full software support for the most recent versions of Microsoft Windows*. Many older Intel® Ethernet adapters have base drivers supplied by Microsoft Windows. Lists of supported devices per operating system are available [here](#).

4.0 NVM Upgrade/Downgrade 800 Series/700 Series and X550

Refer to the Feature Support Matrix (FSM) links listed in [Related Documents](#) for more detail. FSMs list the exact feature support provided by the NVM and software device drivers for a given release.

5.0 Languages Supported

Note: This only applies to Microsoft Windows and Windows Server Operating Systems.

This release supports the following languages:

Languages	
English French German Italian Japanese	Spanish Simplified Chinese Traditional Chinese Korean Portuguese

6.0 Related Documents

Contact your Intel representative for technical support about Intel® Ethernet Series devices/adapters.

6.1 Feature Support Matrix

These documents contain additional details of features supported, operating system support, cable/modules, etc.

Device Series	Support Link
Intel® Ethernet 800 Series: — E810 — E820 Intel® Ethernet Controller E810 and Intel® Ethernet Connection E82X Feature Comparison Matrix	https://cdrdv2.intel.com/v1/dl/getContent/630155 https://cdrdv2.intel.com/v1/dl/getContent/739764 https://cdrdv2.intel.com/v1/dl/getContent/751546
Intel® Ethernet 700 Series: — X710/XXV710/XL710 — X722 — X710-TM4/AT2 and V710-AT2	https://cdrdv2.intel.com/v1/dl/getContent/332191 https://cdrdv2.intel.com/v1/dl/getContent/336882 https://cdrdv2.intel.com/v1/dl/getContent/619407
Intel® Ethernet 500 Series	https://cdrdv2.intel.com/v1/dl/getContent/335253
Intel® Ethernet 300 Series	N/A
Intel® Ethernet 200 Series	N/A

6.2 Specification Updates

These documents provide the latest information on hardware errata as well as device marking information, SKU information, etc.

Device Series	Support Link
Intel® Ethernet 800 Series	https://cdrdv2.intel.com/v1/dl/getContent/616943
Intel® Ethernet 700 Series: — X710/XXV710/XL710 — X710-TM4/AT2 and V710-AT2	https://cdrdv2.intel.com/v1/dl/getContent/331430 https://cdrdv2.intel.com/v1/dl/getContent/615119
Intel® Ethernet 500 Series — X550 — X540	https://cdrdv2.intel.com/v1/dl/getContent/333717 https://cdrdv2.intel.com/v1/dl/getContent/334566
Intel® Ethernet 300 Series	https://cdrdv2.intel.com/v1/dl/getContent/333066
Intel® Ethernet 200 Series — I210 — I211	https://cdrdv2.intel.com/v1/dl/getContent/332763 https://cdrdv2.intel.com/v1/dl/getContent/333015

6.3 Software Download Package

The release software download package can be found [here](#).

6.4 SourceForge Ethernet Drivers and Utilities

For additional information regarding Linux kernel drivers, please refer to the [Intel® Ethernet Drivers and Utilities](#) SourceForge project page.

6.5 Intel Product Security Center Advisories

Intel product security center advisories can be found at:

<https://www.intel.com/content/www/us/en/security-center/default.html>

NOTE: ***This page intentionally left blank.***



LEGAL

No license (express or implied, by estoppel or otherwise) to any intellectual property rights is granted by this document.

This document (and any related software) is Intel copyrighted material, and your use is governed by the express license under which it is provided to you. Unless the license provides otherwise, you may not use, modify, copy, publish, distribute, disclose or transmit this document (and related materials) without Intel's prior written permission. This document (and related materials) is provided as is, with no express or implied warranties, other than those that are expressly stated in the license.

Intel disclaims all express and implied warranties, including without limitation, the implied warranties of merchantability, fitness for a particular purpose, and non-infringement, as well as any warranty arising from course of performance, course of dealing, or usage in trade.

This document contains information on products, services and/or processes in development. All information provided here is subject to change without notice. Contact your Intel representative to obtain the latest forecast, schedule, specifications and roadmaps.

The products and services described may contain defects or errors which may cause deviations from published specifications.

Intel and the Intel logo are trademarks of Intel Corporation in the U.S. and/or other countries.

Other names and brands may be claimed as the property of others.

© 2022 Intel Corporation.

