# Intel® Ethernet Controller Products

## 27.5 Release Notes

**Ethernet Products Group**

*July 2022*

Revision 1.0
733990-001

# Revision History

| Revision | Date | Comments |
|----------|------|----------|
| 1.0 | July 2022 | Initial release. |

# 1.0     Overview

This document provides an overview of the changes introduced in the latest Intel® Ethernet Network 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 |
|---|---|
| 27.5 | • None for this release. |

### 1.1.2     Software Features

| Release | New Software Support |
|---|---|
| 27.5 | **Support for the following Linux\* operating system distributions:**<br>• SUSE\* Linux Enterprise Server (SLES) 15 SP4<br>• Red Hat\* Enterprise Linux\* (RHEL) 9.0.<br>• Red Hat Enterprise Linux 8.6.<br>• Canonical\* Ubuntu\* 22.04 LTS.<br>  Drivers are no longer tested on RHEL 8.5, RHEL 7.9, and Ubuntu 18.04.<br>**On Intel® Ethernet 810 Series devices:**<br>• Support for Intel® Scalable I/O Virtualization on Linux operating systems.<br>• Support for transmit balancing feature.<br>• Support for QinQ in VMware ESXi\*.<br>• Support for updating DDP packages in VMware ESXi using esxcli. The minimum driver version for this functionality is icen-1.9.1.x or higher.<br>• Support for ADQ with ADQ configuration script.<br>• Support for ADQ with Independent Pollers. |

### 1.1.3     Removed Features

| Release | Hardware/Feature Support |
|---|---|
| 27.5 | • None for this release. |

### 1.1.4     Firmware Features

| Release | New Firmware Support |
|---|---|
| 27.5 | • Improved PTP time stamp SW polling latency reduction.<br>• LLDP / RDE read / Write support is now compliant with the DMTF standard.<br>• Added support for No-FEC option to support some optical modules within 25G speeds with the switches.<br>• Added support for PLDM PCIe reset avoidance at NVMUpdate. |

## 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** | **4.0** |
| **820 Series** | |
| **E822-C** | 2.28.1 |
| **E822-L** | 2.28.1 |
| **E823** | 2.28.1 |
| **700 Series** | |
| **700** | 8.7 |
| **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

### 1.4.1.1 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* Enter-prise Linux* (RHEL) 9.0 | Red Hat* Enter-prise Linux* (RHEL) 8.6 | RHEL 8.x (8.5 and previous) | RHEL 7.x | SUSE* Linux Enter-prise 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 | 1.9.11 | 1.9.11 | 1.9.11 | SNT | 1.9.11 | 1.9.11 | 1.9.11 | SNT | 1.9.11 | 1.9.11 | 1.9.11 |
| **Intel® Ethernet 820 Series** | | | | | | | | | | | |
| ice | 1.9.11 | SNT | **SNT** | SNT | 1.9.11 | SNT | 1.9.11 | SNT | 1.9.11 | SNT | SNT |
| **Intel® Ethernet 700 Series** | | | | | | | | | | | |
| i40e | 2.20.12 | 2.20.12 | SNT | SNT | 2.20.12 | SNT | 2.20.12 | SNT | 2.20.12 | 2.20.12 | SNT |
| **Intel® Ethernet Adaptive Virtual Function** | | | | | | | | | | | |
| iavf | 4.5.3 | 4.5.3 | SNT | SNT | 4.5.3 | SNT | 4.5.3 | SNT | 4.5.3 | 4.5.3 | SNT |
| **Intel® Ethernet 10 Gigabit Adapters and Connections** | | | | | | | | | | | |
| ixgbe | 5.16.5 | 5.16.5 | SNT | SNT | 5.16.5 | SNT | 5.16.5 | SNT | 5.16.5 | 5.16.5 | SNT |
| ixgbevf | 4.16.5 | 4.16.5 | SNT | SNT | 4.16.5 | SNT | 4.16.5 | SNT | 4.16.5 | 4.16.5 | SNT |
| **Intel® Ethernet Gigabit Adapters and Connections** | | | | | | | | | | | |
| igb | 5.11.4 | 5.11.4 | SNT | SNT | 5.11.4 | SNT | 5.11.4 | SNT | 5.11.4 | 5.11.4 | SNT |
| **Remote Direct Memory Access (RDMA)** | | | | | | | | | | | |
| irdma | 1.9.30 | 1.9.30 | SNT | SNT | 1.9.30 | SNT | 1.9.30 | SNT | 1.9.30 | 1.9.30 | SNT |

## 1.4.2        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** | **1.12.144.0** | **1.12.144.0** | **1.12.144.0** | NS | NS |
| **Intel® Ethernet 820 Series** | | | | | |
| **scea** | 1.1.285.0 | 1.1.285.0 | NS | NS | NS |
| **Intel® Ethernet 700 Series** | | | | | |
| **i40ea** | 1.16.202.x | 1.16.202.x | 1.16.202.x | 1.16.202.x | 1.16.62.x |
| **i40eb** | 1.16.202.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.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** | 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 |
| **vxs** | 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 |

intel.

**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 | 12.19.2.45 | 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 |
| e1q | NS | NS | NS | NS | 12.7.28.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 |

## 1.4.3 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 | Microsoft Windows 10 | Microsoft Windows 8.1 | Microsoft Windows 8 |
|---|---|---|---|---|---|
| **Intel® Ethernet 810 Series** | | | | | |
| **icea** | NS | NS | NS | NS | NS |
| **Intel® Ethernet 700 Series** | | | | | |
| **i40ea** | 1.17.80.0 | 1.16.202.0 | NS | NS | NS |
| **i40eb** | 1.17.80.0 | 1.16.202.0 | 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 |
| **vxs** | NS | NS | NS | NS | NS |
| **Intel® Ethernet 2.5 Gigabit Adapters and Connections** | | | | | |
| **e2f** | 2.1.1.7 | 1.1.3.28 | NS | NS | NS |

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

| Driver | Microsoft Windows 11 | Microsoft Windows 10, version 1809 | 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 | 12.19.2.45 | 12.19.2.45 | 12.18.8.4 | 12.17.8.7 | 12.17.8.7 |
| e1e | NS | NS | NS | NS | 9.16.10.x |
| e1k | NS | NS | NS | NS | 12.10.13.x |
| e1q | NS | NS | NS | NS | 12.7.28.x |
| e1r | 13.0.14.0 | 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 |

## 1.4.4 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** | **1.35.5** | **1.35.5** | SNT |
| **Intel® Ethernet 820 Series** | | | |
| **ice** | SNT | SNT | **1.35.5** |
| **Intel® Ethernet 700 Series** | | | |
| **ixl** | 1.12.35 | 1.12.35 | SNT |
| **Intel® Ethernet Adaptive Virtual Function** | | | |
| **iavf** | **3.0.30** | **3.0.30** | SNT |
| **Intel® Ethernet 10 Gigabit Adapters and Connections** | | | |
| **ix** | 3.3.31 | 3.3.31 | SNT |
| **ixv** | 1.5.32 | 1.5.32 | SNT |
| **Intel® Ethernet Gigabit Adapters and Connections** | | | |
| **igb** | 2.5.24 | 2.5.24 | SNT |
| **Remote Direct Memory Access (RDMA)** | | | |
| **irdma** | 1.**0.9** | **1.0.9** | 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

- 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.

- When a VF interface is set as '**up**' and assigned to a namespace, and the namespace is then deleted, the dmesg log may show the error **Failed to set LAN Tx queue context, error: ICE_ERR_PARAM** followed by error codes from the *ice* and *iavf* drivers.

#### 2.1.1.2    Linux Driver

- Repeatedly assigning a VF interface to a network namespace then deleting that namespace might result in an unexpected error message and might possibly result in a call trace on the host system.

- When the queue settings of a port are modified using the **ethtool -L ethx combined XX** command, the Interrupt Moderation settings reset to default.

- When the maximum allowed number of VLAN filters are created on a trusted VF, and the VF is then set to untrusted and the VM is rebooted, the *iavf* driver may not load correctly in the VM and may show errors in the VM dmesg log.

- If trusted mode is enabled for a VF while promiscuous mode is disabled and multicast promiscuous mode is enabled, unicast packets may be visible on the VF and multicast packets may not be visible on the VF. Alternatively, if promiscuous mode is enabled and multicast promiscuous mode is disabled, then both unicast and multicast packets may not be visible on the VF interface.

- A VF may incorrectly receive additional packets when trusted mode is disabled but promiscuous mode is enabled.

- Receive hashing might not be enabled by default on Virtual Functions when using an older iavf driver in combination with a newer PF driver version.

- Linux **sysctl** commands, or any automated scripting that alerts or sets **/proc/sys/** attributes using **sysctl**, might encounter a system crash that includes **irdma_net_event** in the **dmesg** stack trace.

  **Workaround**: With OOT irdma-1.8.X installed on the system, avoid running **sysctl** while drivers are being loaded or unloaded.

- irdma stops working if the number of ice driver queues are changed (ethtool -L) while the irdma driver is loaded. As a workaround, remove (if previously loaded) and reload irdma after changing the number of queues.

#### 2.1.1.3    Windows Driver

None for this release.

#### 2.1.1.4    Linux RDMA Driver

None for this release.

### 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

- When sending an AQC to start, the LLDP agent does not change the property **NetworkAdapter.LldpEnabled** to **True**.
- Flow control settings have no effect on traffic, and counters do not increment with flow control set to TX=ON and Rx=OFF. However, flow control works fine with values set to TX=On RX=ON.

### 2.1.1.8 Manageability

None for this release.

### 2.1.1.9 FreeBSD Driver

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

### 2.1.1.10 VMware Driver

None for this release.

### 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

## 2.2 Intel® Ethernet 700 Series

### 2.2.1 General

None for this release.

### 2.2.2 VMware Driver

None for this release.

### 2.2.3 Linux driver:

- When large 1G aligned buffers are used for memory registration there is a possibility that the data is placed at an incorrect location in the buffer for Intel® Ethernet X722 devices resulting in data corruption. The irdma driver incorrectly uses a 1G host page size which is not supported on Intel® Ethernet X722 devices. The fix is in the irdma driver 1.9.27 or later to stops advertising 1G host page size for Intel® Ethernet X722 devices.

### 2.2.4 Intel® PROSet:

None for this release.

### 2.2.5 EFI Driver

None for this release.

### 2.2.6 NVM

- If the error message **OS layer initialization failed** is displayed, please update the Windows QV driver to the version included in this release.

**Note:** If you are using Proset, an update of the QV driver may also require updating the Proset.

### 2.2.7 Windows driver:

None for this release.

### 2.2.8 Intel® Ethernet Flash Firmware Utility:

None for this release.

## 2.3 Intel® Ethernet 500 Series

None for this release.

## 2.4 Intel® Ethernet 300 Series

None for this release.

## 2.5 Intel® Ethernet 200 Series

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

- Sometimes, when RX/TX pause frames is enabled and the system is rebooted, or when the icen driver is unloaded, only one PF can be visible on ESX.

- 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 may 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) may 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 will be 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 may not be working as expected. This has been resolved in Release 26.4.

- 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 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 E810. Depending on the implementation, link may never be achieved. In other cases, if the module sends IDLEs to the E810 when there is no BASE-T link, the E810 may interpret this as a link partner sending valid data and may 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.c
  om/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.

- In order 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.

- When having link issues (including no link) at link speeds faster than 10 Gb/s, check the switch configuration and/or specifications. Many optical connections and direct attach cables require RS-FEC for connection speeds faster than 10 Gb/s. One of the following might resolve the issue:

  Configure the switch to use RS-FEC mode.

  — Specify a 10 Gb/s, or slower, link speed connection.

  — If attempting to connect at 25 Gb/s, try using an SFP28 CA-S or CS-N direct attach cable. These cables do not require RS-FEC.

  — If the switch does not support RS-FEC mode, check with the switch vendor for the availability of a software or firmware upgrade.

### 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).

- Per the specification, the **Get LLDP** command (0x28) response may 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.

### 3.1.1.3      Linux

- 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 may not increment correctly.

### 3.1.1.4      Linux Driver

- 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.

- 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.

- It may 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 using certain DDP package versions, 802.1ad type VLANs may not be correctly enabled on SIOV or SR-IOV interfaces.

- 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.

- Repeatedly adding/deleting a VF from a namespace while also repeatedly changing its trust mode status may 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 device may 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.

  **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**VXLAN stateless offloads (checksum, TSO), as well as TC filters directing traffic to a VXLAN interface are not supported with Linux v5.9 or later.

- 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, **l2-fwd-offload** cannot be turned on.

- 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.

- 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 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 may cause further issues, depending on your switch configuration.

- Changing the FEC value from BaseR to RS results in an error message in dmesg, and may 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 may 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.

- 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.

- 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 will not affect functionality.

- With the current ice PF driver, there may not be a way for a trusted DPDK VF to enable unicast promiscuous without turning on "ethtool --priv-flags" with vf-true-promic-support.".

- 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.

- 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 may contain an older version of **iproute2/devink** which may result in errors. **Workaround**: Please update to the latest **devlink** version.

- On Intel® Ethernet Network Adapter E810-XXVDA4T, the driver may not link at 1000baseT and 1000baseX. The link may go down after advertising 1G.

### 3.1.1.5 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.

- 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 may 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

- On RHEL 7.9, installing rdma-core v35.0 debuginfo rpms can prevent the installation of debuginfo rpms from indistro products like libfabric.

- With Intel® Ethernet 800 Series Network Adapter iwarp, ib_write_bw fails with the a RDMA_CM_EVENT_REJECTED with at least 1000 QPs.

- Simultaneously running four mpi_stress mpi jobs with IEFS on Intel® Ethernet 800 Series Network Adapters and rocev2 across four nodes with eight ranks per node can fail to connect.

  Eventually, the only recourse may be to unload and reload the irdma driver.

- When using Intel MPI 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 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.

### 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.

- After downgrading to Release 25.6 (and previous), a loss of traffic may result. Workaround: Unload and reload the driver to resume traffic. Rebooting the system would also help.

### 3.1.1.9 VMware Driver

- Configuring the NSX-T Virtual Distributed Switch uplink port may 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 may fail.

- Excessive VF reset may cause multicast ping from VF to fail.

- TCP traffic on VMs maybe be interrupted if PF undergoes a reset.

- The output of the PTP signal may 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

- On VMware ESXi vSphere 6.7 and Linux VM with Kernel version higher than 5.14, after performing warm reboot, traffic on SRIOV VF adapters will be dropped.

- Workaround: reload VF driver when issue occurs or instead of warm reboot perform power cycle from ESXi hypervisor.

- 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 may 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 may 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 may fail.

### 3.1.1.9.1    VMware RDMA Driver

- Running Unreliable Datagram (UD) RDMA mixed traffic with more than two QPs may 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 irdman driver must be unloaded on the system.

  If this recommendation is not followed, the following failures may occur:

  — Loading a DDP package during RDMA traffic may lead to system hang that requires a server reset to recover.

  — Loading a DDP package with RDMA enabled (without RDMA traffic running) may fail, and the device may become unusable for RDMA traffic until recovered by reboot.

- For RDMA, VMware recommends lossless traffic to configured on priority 3, but due to a known issue in E810 adapters it requires priority '2' to be configured. The irdman driver will automatically translate this to priority '3' in the ESXi OS. If other priority settings are used, the priority seen in the packet may be different than DCB setting on the NIC and RDMA traffic will not be properly configured as lossless.

### 3.1.1.10 Application Device Queues (ADQ)

The code contains the following known issues:

**ADQ Standard Issues**

- 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 ADQ Configuration Guide. The following issues may 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, may leave the system in an unusable state which requires a forced reboot to clear.

  — Mitigation — Follow the steps documented in the ADQ Configuration Guide to "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 -L` 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 may 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 may 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**

- When ADQ is enabled on VFs, TC filters on the VF TC0 (default TC) are not supported and will not pass traffic. It is not expected to add TC filters to TC0 since it is reserved for non-filtered default traffic.

- 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.

- 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*.

- 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.

- 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.

- Configuring ADQ traffic classes with an odd number of hardware queues on a VF interface may 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.

## 3.1.1.11    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

- In addition to valid configurations, the HII Menu also lists invalid and extra configurations.

- With a Windows host and Linux virtual machine (VM), the last transmit (Tx) queue may not increment when there are multiple Tx queues.

- **Insufficient PCI-Express bandwidth available for device** may 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.

- 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

- After changing port options in EPCT and the system returns blank mode, perform a second reboot.

- 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.

- 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.

- No link displayed on the 827 QSFP after setting to 50G.

- The **Get LLDP** command (0x28) response contains only 2 TLV types. There should be a third TLV type.

### 3.1.2.3 Linux Driver

- 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.

- 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.

- 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).

- During RSS testing, incorrect values are displayed in OIDs for Maximum Number of RSS Queues and Processors.

- 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.

- When Double VLAN Mode is enabled on the host, disabling then re-enabling a Virtual Function attached to a Windows guest may cause error messages to be shown in **dmesg**. This behavior is not expected to affect functionality.

- In DCB mode, when switching from IEEE to CEE mode, configurations are not applied.

  Workarounds are to either start in CEE mode or enter unwilling mode to perform configurations.

- On RHEL 7.9 VMs, VF traffic doesn't 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

- VXLAN traffic is not passing through on the expected queues.

- 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.

- At 100G, packets are not counted correctly between two virtual machines (VM) and the associated virtual functions (VF).

### 3.1.2.4 FreeBSD Driver

- During successive driver unload/load cycles, single-port adapters may 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

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

### 3.1.2.6 Linux RDMA Driver

None for this release.

### 3.1.2.7 NVM Update Tool

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

### 3.1.2.8 VMware Driver

- LAN traffic is disabled between Windows VFs after changing TPID to 88a8.

### 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) may 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

### 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-based adapters 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-based 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

None for this release.

### 3.2.5　　　Intel® PROSet

None for this release.

### 3.2.6　　　EFI Driver

- In the BIOS Controller Name as part of the Controller Handle section, a device path appears instead of an Intel adapter branding name.

### 3.2.7　　　VMware Driver

None for this release.

### 3.2.8　　　NVM

None for this release.

## 3.3　　　Intel® Ethernet 500 Series

### 3.3.1　　　General

None for this release.

### 3.3.2　　　EFI Driver

- In the BIOS Controller Name as part of the Controller Handle section, a device path appears instead of an Intel adapter branding name.

### 3.3.3　　　Windows Driver

None for this release.

## 3.4　　　Intel® Ethernet 300 Series

### 3.4.1　　　EFI Driver

- In the BIOS Controller Name as part of the Controller Handle section, a device path appears instead of an Intel adapter branding name.

### 3.4.2　　　VMware Driver

None for this release.

## 3.5　　　Intel® Ethernet 200 Series

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 | Spanish |
| French | Simplified Chinese |
| German | Traditional Chinese |
| Italian | Korean |
| Japanese | Portuguese |

# 6.0    Related Documents

Contact your Intel representative for technical support about Intel® Ethernet Series Controllers/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 Network Adapter:<br>— E810 | https://cdrdv2.intel.com/v1/dl/getContent/630155 |
| Intel® Ethernet 700 Series:<br>— X710/XXV710/XL710<br>— X722<br>— X710-TM4/AT2 and V710-AT2 | https://cdrdv2.intel.com/v1/dl/getContent/332191<br>https://cdrdv2.intel.com/v1/dl/getContent/336882<br>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 Network Adapter | https://cdrdv2.intel.com/v1/dl/getContent/616943 |
| Intel® Ethernet 700 Series:<br>— X710/XXV710/XL710<br>— X710-TM4/AT2 and V710-AT2 | https://cdrdv2.intel.com/v1/dl/getContent/331430<br>https://cdrdv2.intel.com/v1/dl/getContent/615119 |
| Intel® Ethernet 500 Series<br>— X550<br>— X540 | https://cdrdv2.intel.com/v1/dl/getContent/333717<br>https://cdrdv2.intel.com/v1/dl/getContent/334566 |
| Intel® Ethernet 300 Series | https://cdrdv2.intel.com/v1/dl/getContent/333066 |
| Intel® Ethernet 200 Series<br>— I210<br>— I211 | https://cdrdv2.intel.com/v1/dl/getContent/332763<br>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.

Copies of documents which have an order number and are referenced in this document may be obtained by calling 1-800-548-4725 or by visiting www.intel.com/design/literature.htm.

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.