Intel® Ethernet Controller Products

27.2.1 Release Notes

Ethernet Products Group

April 2022

Revision 1.1
727561-001

Revision History

Revision	Date	Comments
1.1	April 2022	Updated for Release 27.2.1: • NDIS driver(s) • FreeBSD driver(s)
1.0	April 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
27.2.1	None for this release.

1.1.2 Software Features

Release	New Software Support
27.2.1	• Update to the 800 Series Network Driver Interface Specification (NDIS) drivers. • Support for Microsoft* Windows* 11. • On Microsoft Windows 11 (and later), devices based on the Intel® Ethernet I225 controller use a NetAdapter miniport driver. • **NOTE**: The **e2fn** driver, which is built around the NetAdapter framework (not NDIS), is incompatible with the NVMUpdate tool.

1.1.3 Removed Features

Release	Hardware/Feature Support
27.2.1	• None for this release.

1.1.4 Firmware Features

Release	New Firmware Support
27.2.1	• 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

1.3 NVM

Table 1 shows the NVM versions supported in this release.

Table 1. Current NVM

Driver	NVM Version
800 Series	
E810	3.20
700 Series	
700	8.6
500 Series	
X550	3.6
200 Series	
I210	2.0

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* Enter-prise Linux* (RHEL) 8.5	RHEL 8.x (8.4 and previous)	RHEL 7.9	RHEL 7.x (7.8 and previous)	SUSE* Linux Enter-prise Server (SLES) 15 SP3	SLES 15 SP2 and previous	SLES 12 SP5	SLES 12 SP4 and previous	Canonical * Ubuntu* 20.04 LTS	Canonical Ubuntu 18.04 LTS	Debian* 11
Intel® Ethernet 800 Series											
ice	**1.8.8**	SNT	**1.8.8**	SNT	**1.8.8**	SNT	**1.8.8**	SNT	**1.8.8**	**1.8.8**	**1.8.8**
Intel® Ethernet 700 Series											
i40e	2.18.9	SNT	2.18.9	SNT	2.18.9	SNT	2.18.9	SNT	2.18.9	2.18.9	SNT
Intel® Ethernet Adaptive Virtual Function											
iavf	**4.4.2.1**	SNT	**4.4.2.1**	SNT	**4.4.2.1**	SNT	**4.4.2.1**	SNT	**4.4.2.1**	**4.4.2.1**	SNT
Intel® Ethernet 10 Gigabit Adapters and Connections											
ixgbe	5.14.6	SNT	5.14.6	SNT	5.14.6	SNT	5.14.6	SNT	5.14.6	5.14.6	SNT
ixgbevf	4.14.5	SNT	4.14.5	SNT	4.14.5	SNT	4.14.5	SNT	4.14.5	4.14.5	SNT
Intel® Ethernet Gigabit Adapters and Connections											
igb	5.9.3	SNT	5.9.3	SNT	5.9.3	SNT	5.9.3	SNT	5.9.3	5.9.3	SNT
Remote Direct Memory Access (RDMA)											
irdma	1.8.45	SNT	1.8.45	SNT	1.8.45	SNT	1.8.45	SNT	1.8.45	1.8.45	SNT

1.4.3 Windows Server

Table 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 800 Series					
icea	1.11.44.0	1.11.44.0	1.11.44.0	NS	NS
Intel® Ethernet 700 Series					
i40ea	1.16.139.x	1.16.62.x	1.16.62.x	1.16.62.x	1.16.62.x
i40eb	1.16.141.x	1.16.62.x	1.16.62.x	1.16.62.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	NS	**1.0.2.22**	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	NS	12.19.1.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
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

1.4.4 Windows Client

Table 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 800 Series					
icea	NS	NS	NS	NS	NS
Intel® Ethernet 700 Series					
i40ea	**1.17.80.0**	NS	NS	NS	NS
i40eb	**1.17.80.0**	NS	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	NS
ixn	NS	4.1.239.x	4.1.239.x	3.14.214.x	NS
ixs	**4.1.248.0**	4.1.246.x	4.1.246.x	3.14.223.x	NS
ixt	NS	4.1.228.x	4.1.229.x	3.14.214.x	NS
sxa	NS	4.1.243.x	4.1.243.x	3.14.222.x	NS
sxb	NS	4.1.239.x	4.1.239.x	3.14.214.x	NS
vxn	NS	2.1.249.x	2.1.243.x	1.2.309.x	NS
vxs	NS	2.1.230.x	2.1.232.x	1.2.254.x	NS
Intel® Ethernet 2.5 Gigabit Adapters and Connections					
e2f	**2.1.1.7**	NS	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**	12.19.1.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
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	1.4.7.x	1.4.7.x	1.4.5.x	1.4.5.x

1.4.5 FreeBSD

Table 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 12.2	FreeBSD 12.1 and previous
Intel® Ethernet 800 Series			
ice	**1.34.6**	**1.34.6**	SNT
Intel® Ethernet 700 Series			
ixl	1.12.32	1.12.32	SNT
Intel® Ethernet Adaptive Virtual Function			
iavf	**3.0.29**	**3.0.29**	SNT
Intel® Ethernet 10 Gigabit Adapters and Connections			
ix	3.3.30	3.3.30	SNT
ixv	1.5.31	1.5.31	SNT
Intel® Ethernet Gigabit Adapters and Connections			
igb	2.5.21	2.5.21	SNT

Table 5. Supported Operating Systems: FreeBSD [continued]

Driver	FreeBSD 13	FreeBSD 12.2	FreeBSD 12.1 and previous
Remote Direct Memory Access (RDMA)			
irdma	1.0.0	1.0.0	SNT
iw_ixl	0.1.30	0.1.30	SNT

2.0 Fixed Issues

2.1 Intel® Ethernet 800 Series

2.1.1 General

2.1.2 Linux Driver

- Prior to irdma version 1.8.45, installing the OOT irdma driver on a system with RDMA-capable Intel® Ethernet Connection X722/Intel® Ethernet Network Adapter X722 ports and using an OS or kernel with an in-tree irdma driver could cause a system crash. To prevent a system crash when using OOT irdma drivers, either use irdma 1.8.45, or update i40e version (2.18.9 or greater) and load it before this new irdma is loaded.
- AF_XDP based applications may cause system crash on packet receive with RHEL based 4.18 kernels.
- During a long reboot cycle test (about 250-500 reboots) of the Intel Ethernet 800 Series adapters, the Intel ICE and iavf driver may experience kernel panics leading to an abnormal reboot of the server.
- The commands **ethtool -C [rx|tx]-frames** are not supported by the iavf driver and will be ignored.

  Setting [tx|rx]-frames-irq using ethtool -C may not correctly save the intended setting and may reset the value back to the default value of 0.
- Interrupt Moderation settings reset to default when the queue settings of a port are modified using the **ethtool -L ethx combined XX** command.
- When a VM is running heavy traffic loads and is attached to a Virtual Switch with either SR-IOV enabled or VMQ offload enabled, repeatedly enabling and disabling the SR-IOV/VMQ setting on the vNIC in the VM may result in a BSOD. Linux RDMA Driver
- In order to send or receive RDMA traffic, the network interface associated with the RDMA device must be up. If the network interface experiences a link down event (for example, a disconnected cable or ip link set <interface> down), the associated RDMA device is removed and no longer available to RDMA applications. When the network interface link is restored, the RDMA device is automatically re-added.
- RHEL 8.5 only: Any usermode test that uses **ibv_create_ah** (For example, a RoCEv2 usermode test such as udaddy) will fail.
- Due to a nondeterministic race condition, if the irdma driver is loaded in Linux by an Intel® Ethernet 800 series device with non-standard MTU (i.e., non-1500B MTU), the system's network interfaces may fail to load after reboot. After failing to load, interactions with the networking stack may hang on the system. Multiple reboots may be required to avoid the condition.
- The Devlink command **devlink dev param show** (DEVLINK_CMD_PARAM_GET) does not report MinSREV values for firmware (fw.mgmt.srev) and OROM (fw.undi.srev). This defect was also seen on the NVMUpdate tool, which caused an inventory error.

2.1.3 Windows Driver

- When a VM is running heavy traffic loads and is attached to a Virtual Switch with either SR-IOV enabled or VMQ offload enabled, repeatedly enabling and disabling the SR-IOV/VMQ setting on the vNIC in the VM, may result a VM freeze/hang.

2.1.4 Linux RDMA Driver

- iWARP mode requires a VLAN to be configured to fully enable PFC.

2.1.5 NVM Update Tool

None for this release.

2.1.6 NVM

None for this release.

2.1.7 Firmware

- Following a firmware update and reboot/power cycle on the Intel Ethernet CQDA2 Adapter, Port 1 is displaying NO-CARRIER and is not functional.
- Added a state machine to the thermal threshold activity so that when the switch page fails, it tries again from the same state.
- FW not allow link if module not supported in lenient mode.
- RDE Device is reporting a **RevisionID** property of **PCIeFunctions** schema as 0x00, instead 0x02.
- The RDE device reports its status as **Starting** (with low power), even though it is in standby mode.
- Wake On LAN flow is unexpectedly triggered by the E810 CQDA2 for OCP 3.0 adapter. The server unexpectedly wakes up automatically from S5 power state in few seconds after shut down from the OS, and it is impossible to shut down the server.
- Fixed an issue where the FW was reporting a module power value of module from an incorrect location.

2.1.8 Manageability

None for this release.

2.1.9 FreeBSD Driver

None for this release.

2.1.10 Application Device Queues (ADQ)

None for this release.

2.2 Intel® Ethernet 700 Series

2.2.1 General

None for this release.

2.2.2 Linux driver:

None for this release.

2.2.3 Intel® PROSet:

None for this release.

2.2.4 EFI Driver

None for this release.

2.2.5 NVM

None for this release.

2.2.6 Windows driver:

None for this release.

2.2.7 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

3.1.1 General

- 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 Parkvale (C827/XL827) retimer is interfaced with SX/LX, SR/LR, SR4/LR4, AOC limiting optics. This issue is isolated to Parkvale line side PMD RX susceptibility to noise.
- Intel Ethernet 800 Series 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.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.
- In order for an Intel® Ethernet 800 Series-based adapter to reach its full potential, users must install it in a PCIe Gen4 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.

3.1.2 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.
- 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.

3.1.3 Linux Driver

- 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.
- 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" "**.
- 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.
- 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 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.
- 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.
- 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.
- 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.
- 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.

- 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 newer AVF 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 might not be a way for a trusted VF to enable unicast promiscuous and multicast promiscuous mode without turning on ethtool --priv-flags with vf-true-promisc-support. As such, the expectation is to not use vf-true-promisc-support to gate VF's request for unicast/multicast promiscuous mode.
- 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.
- 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.
- 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 Adapter XXVDA4T, the driver may not link at 1000baseT and 1000baseX. The link may go down after advertising 1G.

3.1.4 FreeBSD Driver

- 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.
- During stress, the FreeBSD-13.0 virtual guest interfaces may experience poor receive performance. Windows Driver
- 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.

3.1.5 Windows RDMA Driver

- 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.

3.1.6 Linux RDMA 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.

3.1.7 NVM Update Tool

- 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.

3.1.8 Application Device Queues (ADQ)

The code contains the following known issues:

- 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.
- 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 ice-1.8.2, 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.
- 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® E810 Network adapter 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.
- 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 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.
- 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.
- 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.
- 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.

3.1.9 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.

3.2 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-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-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

- The i40e driver does not compile under SLES SLE15 SP3 with kernel 5.3.18-150300.59.43.1 or newer (or any kernel that has **150300** in its version string).

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 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.

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.

3.5 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 languages listed in the table that follows:

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 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	https://cdrdv2.intel.com/v1/dl/getContent/630155
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 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