



Intel[®] Ethernet Controller Products

28.2 Release Notes

NEX Cloud Networking Group (NCNG)

October 2023

Revision 1.3
778690-001

Revision History

Revision	Date	Comments
1.3	October 2023	Added an issue in section 3.2.2 about Windows Event ID 411.
1.2	September 2023	Corrected an issue in section 3.2.1 about DMARemappingCompatible.
1.1	August 2023	Added section 1.6 DDP Versions Supported.
1.0	August 2023	Initial release.

Contents

1.0 Overview	5
1.1 New Features	5
1.1.1 Hardware Support	5
1.1.2 Software Features	5
1.1.3 Firmware Features	5
1.2 Removed Features	6
1.3 Operating Systems Supported	8
1.3.1 Linux	8
1.4 Windows Server	8
1.4.1 Windows Client	10
1.4.2 FreeBSD	11
1.4.3 ESXi Drivers	11
1.5 NVM Versions Supported	12
1.6 DDP Versions Supported	12
2.0 Fixed Issues	13
2.1 Intel® Ethernet 800 Series Network Adapters	13
2.1.1 Intel® Ethernet 810 Series	13
2.1.2 Intel® Ethernet 820 Series	15
2.2 Intel® Ethernet 700 Series Network Adapters	16
2.2.1 VMware Driver	16
2.2.2 Firmware/NVM/NVM Update	16
2.2.3 Windows Driver	17
3.0 Known Issues	18
3.1 Intel® Ethernet 800 Series Network Adapters	18
3.1.1 Intel® Ethernet 810 Series	18
3.1.2 Intel® Ethernet 820 Series	24
3.2 Intel® Ethernet 700 Series Network Adapters	28
3.2.1 General	28
3.2.2 Windows Driver	28
3.2.3 Intel® Ethernet Controller V710-AT2/X710-AT2/TM4	28
3.2.4 Linux Driver	28
3.2.5 VMware Driver	28
3.2.6 Firmware/NVM/NVM Update	29
3.3 Intel® Ethernet 500 Series Network Adapters	29
3.4 Legacy Devices	29
4.0 NVM Upgrade/Downgrade 800 Series/700 Series and X550	30
5.0 Languages Supported	30
6.0 Related Documents	30
6.1 Feature Support Matrix	30
6.2 Specification Updates	31
6.3 Software Download Package	31
6.4 SourceForge Ethernet Drivers and Utilities	31
6.5 Intel Product Security Center Advisories	31



NOTE: *This page intentionally left blank.*

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
28.2	<ul style="list-style-type: none"> • Support for the Intel® Ethernet Connection (20) I219-LM • Support for the Intel® Ethernet Connection (24) I219-LM • Support for the Intel® Ethernet Connection (25) I219-LM • Support for the Intel® Ethernet Connection (26) I219-LM • Support for the Intel® Ethernet Connection (27) I219-LM • Support for the Intel® Ethernet Connection (20) I219-V • Support for the Intel® Ethernet Connection (24) I219-V • Support for the Intel® Ethernet Connection (25) I219-V • Support for the Intel® Ethernet Connection (26) I219-V • Support for the Intel® Ethernet Connection (27) I219-V • Support for the Intel® Ethernet Network Adapter I226-T1

1.1.2 Software Features

Release	New Software Support
28.2	<ul style="list-style-type: none"> • Support for Microsoft* Windows* 10 on Intel Intel(R) Ethernet 800 Series 10Gbps devices • Support for Microsoft Windows 10 on Intel Intel(R) Ethernet 800 Series 25Gbps devices • Support for SUSE* Linux Enterprise Server (SLES) 15 SP5 • The following Microsoft* PowerShell* cmdlets now support Intel(R) Ethernet 700 Series devices <ul style="list-style-type: none"> – Get-IntelEthernetLogConfig – Set-IntelEthernetLogConfig – Start-IntelEthernetLog – Stop-IntelEthernetLog – Reset-IntelEthernetLogConfig – Disable-IntelEthernetLogConfig • The following Microsoft PowerShell cmdlets now support "Driver" as a valid value for the "Module" parameter. The new parameter "SubModule" was added to the following cmdlets: <ul style="list-style-type: none"> – Get-IntelEthernetLogConfig – Set-IntelEthernetLogConfig • The Linux iavf driver now supports allocating up to 256 transmit/receive queue pairs

1.1.3 Firmware Features

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

1.2 Removed Features

Release	Hardware/Feature Support
28.2	<p>End of support for RSS on Microsoft Windows operating systems for the following devices:</p> <ul style="list-style-type: none"> • Intel® Ethernet Connection I217-LM • Intel® Ethernet Connection I218-LM • Intel® Ethernet Connection (2) I218-LM • Intel® Ethernet Connection (3) I218-LM • Intel® Ethernet Connection I218-V • Intel® Ethernet Connection (2) I218-V • Intel® Ethernet Connection (3) I218-V • Intel® Ethernet Connection I217-V • Intel® Ethernet Network Adapter I226-T1 • Intel® Ethernet Network Adapter I225-T1 • Intel® Ethernet Controller (2) I225-IT • Intel® Killer(TM) E3100X 2.5 Gigabit Ethernet Controller • Intel® Ethernet Controller (2) I225-LM • Intel® Ethernet Controller (2) I225-LMvP • Intel® Ethernet Controller (2) I225-V • Intel® Ethernet Controller (3) I225-IT • Intel® Killer(TM) E3100X 2.5 Gigabit Ethernet Controller • Intel® Ethernet Controller (3) I225-LM • Intel® Ethernet Controller (3) I225-LMvP • Intel® Ethernet Controller (3) I225-V • Intel® Ethernet Controller I226-IT • Intel® Killer™ E3100X 2.5 Gigabit Ethernet Controller (3) • Intel® Ethernet Controller I226-LM • Intel® Ethernet Controller I226-LMvP • Intel® Ethernet Controller I226-V • Intel® Ethernet Controller I225-LM • Intel® Ethernet Controller I225-V • Intel® Ethernet Connection (2) I219-LM • Intel® Ethernet Connection (2) I219-V • Intel® Ethernet Connection (3) I219-LM • Intel® Ethernet Connection (4) I219-LM • Intel® Ethernet Connection (4) I219-V • Intel® Ethernet Connection (5) I219-LM • Intel® Ethernet Connection (5) I219-V • Intel® Ethernet Connection I219-LM • Intel® Ethernet Connection (10) I219-LM • Intel® Ethernet Connection (11) I219-LM • Intel® Ethernet Connection (12) I219-LM • Intel® Ethernet Connection (13) I219-LM • Intel® Ethernet Connection (14) I219-LM • Intel® Ethernet Connection (15) I219-LM • Intel® Ethernet Connection (16) I219-LM

Release	Hardware/Feature Support
28.2	<ul style="list-style-type: none"> • Intel® Ethernet Connection (17) I219-LM • Intel® Ethernet Connection (18) I219-LM • Intel® Ethernet Connection (19) I219-LM • Intel® Ethernet Connection (20) I219-LM • Intel® Ethernet Connection (22) I219-LM • Intel® Ethernet Connection (23) I219-LM • Intel® Ethernet Connection (24) I219-LM • Intel® Ethernet Connection (25) I219-LM • Intel® Ethernet Connection (26) I219-LM • Intel® Ethernet Connection (27) I219-LM • Intel® Ethernet Connection (6) I219-LM • Intel® Ethernet Connection (7) I219-LM • Intel® Ethernet Connection (8) I219-LM • Intel® Ethernet Connection (9) I219-LM • Intel® Ethernet Connection I219-V • Intel® Ethernet Connection (10) I219-V • Intel® Ethernet Connection (11) I219-V • Intel® Ethernet Connection (12) I219-V • Intel® Ethernet Connection (13) I219-V • Intel® Ethernet Connection (14) I219-V • Intel® Ethernet Connection (16) I219-V • Intel® Ethernet Connection (17) I219-V • Intel® Ethernet Connection (18) I219-V • Intel® Ethernet Connection (19) I219-V • Intel® Ethernet Connection (20) I219-V • Intel® Ethernet Connection (22) I219-V • Intel® Ethernet Connection (23) I219-V • Intel® Ethernet Connection (24) I219-V • Intel® Ethernet Connection (25) I219-V • Intel® Ethernet Connection (26) I219-V • Intel® Ethernet Connection (27) I219-V • Intel® Ethernet Connection (6) I219-V • Intel® Ethernet Connection (7) I219-V • Intel® Ethernet Connection (8) I219-V • Intel® Ethernet Connection (9) I219-V • Intel® Ethernet Connection (3) I218-LM • Intel® Ethernet Connection (3) I218-V

1.3 Operating Systems Supported

1.3.1 Linux

Operating Systems supported:

- Linux Real Time Kernel 5.x and 4.x (only on Intel Ethernet E810 Series)
- Linux, v2.4 kernel or higher
- Red Hat* Enterprise Linux* (RHEL) 9.2
- Red Hat Enterprise Linux 8.8
- SUSE* Linux Enterprise Server (SLES) 15 SP5
- SUSE Linux Enterprise Server 12 SP5
- Canonical* Ubuntu* 22.04 LTS
- Canonical Ubuntu 20.04 LTS
- Debian* 11

Table 1. Supported Operating Systems: Linux

Product	PF Driver	VF Driver	RDMA Driver
Intel® Ethernet 810/820 Series	1.12.6	4.9.1	1.12.55
Intel® Ethernet 700 Series	2.23.17	4.9.1	1.12.55
Intel® Ethernet 10 Gigabit Adapters	5.19.6	4.18.7	Not Supported
Intel® Ethernet Gigabit Adapters	5.14.16	Not Supported	Not Supported

1.4 Windows Server

Operating Systems supported:

- Microsoft Windows Server 2022
- Microsoft Windows Server 2019, Version 1903
- Microsoft Windows Server 2016
- Microsoft Windows Server 2012 R2
- Microsoft Windows Server 2012
- Microsoft Azure Stack HCI

Table 2. Supported Operating Systems: Windows Server

Driver	Windows Server 2022	Windows Server 2019	Windows Server 2016	Windows Server 2012 R2	Windows Server 2012
Intel® Ethernet 800 Series					
icea	1.13.242.0	1.13.242.0	1.13.242.0	Not Supported	Not Supported
scea	1.12.174.0	1.12.174.0	Not Supported	Not Supported	Not Supported
Intel® Ethernet 700 Series					
i40ea	1.18.367.0	1.18.367.0	1.18.367.0	1.16.202.0	1.16.62.0
i40eb	1.18.367.0	1.18.367.0	1.18.367.0	1.16.202.0	Not Supported
Intel® Ethernet Adaptive Virtual Function					
iavf	1.14.201.0	1.14.201.0	1.14.201.0	1.13.8.0	Not Supported
v40e	Not Supported	Not Supported	Not Supported	Not Supported	1.5.86.1
Intel® Ethernet 10 Gigabit Adapters and Connections					
ixs	4.1.252.0	4.1.251.0	4.1.246.0	3.14.223.0	3.14.222.0
sxa	4.1.252.0	4.1.251.0	4.1.249.0	3.14.224.0	3.14.224.0
sxb	4.1.252.0	4.1.251.0	4.1.249.0	3.14.223.0	3.14.223.0
ixt	Not Supported	4.1.228.0	4.1.229.0	3.14.214.0	3.14.206.0
ixn	Not Supported	4.1.251.0	4.1.249.0	3.14.223.0	3.14.223.0
vxs	2.1.246.0	2.1.230.0	2.1.232.0	1.2.254.0	1.2.254.0
vxn	Not Supported	2.1.249.0	2.1.249.0	1.2.317.0	1.2.317.0
Intel® Ethernet 2.5 Gigabit Adapters and Connections					
e2f	2.1.3.15	1.1.4.42	Not Supported	Not Supported	Not Supported
Intel® Ethernet Gigabit Adapters and Connections					
e1r	13.0.16.0	12.18.16.0	12.16.7.0	12.16.5.0	12.14.8.5
e1s	12.16.16.0	12.15.184.4	12.15.184.4	12.13.27.2	12.13.27.2
v1q	Not Supported	1.4.7.3	1.4.7.3	1.4.5.4	1.4.5.0

1.4.1 Windows Client

Operating Systems Supported:

- Microsoft Windows 11 22H2
- Microsoft Windows 11 21H2
- Microsoft Windows 10 21H2
- Microsoft Windows 10, Version 1809

Table 3. Supported Operating Systems: Windows Client

Driver	Windows 11	Windows 10 21H2 / Windows 10 RS5	Windows 10 RS1
Intel® Ethernet 800 Series			
icea	1.13.242.0	1.13.242.0	Not Supported
Intel® Ethernet 700 Series			
i40ea	1.18.367.0	1.18.367.0	Not Supported
Intel® Ethernet Adaptive Virtual Function			
iavf	1.14.201.0	1.14.201.0	1.14.201.0
Intel® Ethernet 10 Gigabit Adapters and Connections			
ixs	4.1.252.0	4.1.251.0	4.1.246.0
ixt	Not Supported	4.1.228.0	4.1.229.0
ixn	Not Supported	4.1.251.0	4.1.249.0
vxs	2.1.246.0	2.1.230.0	2.1.232.0
vxn	Not Supported	2.1.249.0	2.1.249.0
Intel® Ethernet 2.5 Gigabit Adapters and Connections			
e2f	2.1.3.3	1.1.4.39	Not Supported
Intel® Ethernet Gigabit Adapters and Connections			
e1r	13.0.16.0	12.18.16.0	12.16.7.0
e1s	Not Supported	12.15.184.4	12.15.184.4
e1d	12.19.2.50	21H2: 12.19.2.50 RS5: 12.18.9.10	12.18.9.10
e1c	Not Supported	Not Supported	12.15.31.4
v1q	Not Supported	1.4.7.3	1.4.7.3

1.4.2 FreeBSD

Operating Systems supported:

- FreeBSD 13.1
- FreeBSD 12.4

Table 4. Supported Operating Systems: FreeBSD

Driver	PF Driver	VF Driver	RDMA Driver
Intel® Ethernet 810/820 Series	1.38.16	3.0.32	1.2.17
Intel® Ethernet 700 Series	1.13.4	3.0.32	1.2.17
Intel® Ethernet 10 Gigabit Adapters	3.3.35	1.5.34	Not Supported
Intel® Ethernet Gigabit Adapters	2.5.28	Not Supported	Not Supported

1.4.3 ESXi Drivers

Note: Intel® ESXi drivers are available from VMware.

- VMWare ESXi 8.0
- VMware ESXi 7.0

Refer to VMWare's download site for the latest ESXi drivers for Intel® Ethernet® devices.

1.5 NVM Versions Supported

The following table shows the NVM versions supported in this release.

Table 5. Current NVM

Product	NVM Version
810 Series	
E810	4.30
820 Series	
E822	3.26
E823-C	3.26
E823-L	3.26
700 Series	
X710	9.30
X722	6.20
500 Series	
X550	3.60
X552NS	2.10
X552DE	2.10
X553	2.10
200 Series	
I210	2.00

1.6 DDP Versions Supported

The following table shows the versions supported in this release.

Table 6. Current DDP

Package	DDP Version
OS Package	1.3.35.0
Comms Package	1.3.45.0
Wireless Edge Package	1.3.13.0

2.0 Fixed Issues

2.1 Intel® Ethernet 800 Series Network Adapters

2.1.1 Intel® Ethernet 810 Series

2.1.1.1 Linux Driver

- Resolved an issue where when using certain DDP package versions, 802.1ad type VLANs might not be correctly enabled on SIOV or SR-IOV interfaces.
- Resolved an issue where running traffic via Open vSwitch (when an SR-IOV VF have specific ipv6 address assigned), works for 20min (until the neighbor table is clean on the switch/router) then it stops until IPv6 from VF is deleted.
- Previously, the Switchdev feature's VF to VF communication was not functional in the ice driver versions *1.10.1.2*, *1.10.1.2.2*, and *1.11.14*.
- Resolved an issue where the NVM update process would need Power On Reset on PTP devices unless the user manually turned off PTP activity before the NVM update process was started.
- Resolved an issue where a performance issue caused the GNSS write commands to be delayed and not delivered until the next commands appeared.
- Resolved an issue where modifying inner or outer VLAN offload caps while the link is down might break communications when the link is later brought back up.
- Resolved an issue for Linux *ice* driver v1.7.16+ on E810 four-port SKUs where VF could not be created on certain high-core count platform because PF had exhausted all the MSIX interrupts.
- Resolved an issue where repeatedly adding/deleting a VF from a namespace while also repeatedly changing its trust mode status could result in call trace after a significant number of iterations.
- Previously an issue existed in the IAVF driver where the VF couldn't communicate on inner or outer VLANs if the number of rx queues was modified after switching off VLAN offloads.
- Resolved an issue when **spoofchk** was turned on, the VF device driver would have pending DMA allocations while it was released from the device.
- Previously a UEFI PXE installation of Red Hat Enterprise Linux 8.4 on a local disk resulted with the system failing to boot.
- Resolved an issue where sometimes the EMP reset via devlink didn't work and consequently the E810 NVM image update with the `nvupdate` tool reported an error.
- Resolved an issue where auxiliary module conflicts would be reported between **i40e** and **ice** during the generation and installation of Linux rpm packages.
- Resolved an issue where "ethtool -L/G" made E810 link down in Linux Anolis 8.6.
- Previously, during a power cycle, in low probability, E810 25 G adapters showed the wrong media type as other by command "ethtool port-name."
- Resolved an issue where unloading a driver without a DDP profile would disable PF interfaces.
- Limiting the Maximum Bitrate for a Transmit Queue can be enabled by using following sysfs

```
/sys/class/net/<ethx>/queues/tx-0/tx_maxrate
```

Example:

Enable Tx Limit: `echo 300 > /sys/class/net/<ethx>/queues/tx-0/tx_maxrate`

Removing the limit: `# echo 0 > /sys/class/net/<ethx>/queues/tx-0/tx_maxrate`

- Previously, the **set_irq_affinity** script reported errors when more than 64 queues were configured on a server with >64 local cores. This also applies when trying to access any TX queues >=64.

2.1.1.2 Windows Driver

- Windows 11 base driver **icea.sys** was being placed in system32 instead of the driver store.

2.1.1.3 Firmware/NVM/NVM Update

- Resolved an issue where an incorrectly interpreted BDF during MCTP Endpoint discovery led to sending PLDM events to a non-existing device.
- Resolved an issue where a failed RDE operation was improperly terminated, and blocked subsequent RDE operations.
- Resolved an issue where the id field of Event BEJ was changed from the originating resource Id as the resource number to a deferred binding resource Id.
- A "Cannot initialize port" failure occurs when running `nvmupdate 1.39.32.6` in SLES15SP3, because the tool doesn't support `ioctl`.
- A SyncE-related transaction between driver and device may cause a PHY loading failure, during NVM upgrade. Subsequently, PHY errors are reported to UEFI driver during reboot, resulting in a BIOS hang.
- Resolved an issue where E810-XXV-4T didn't expose one of the On Chip thermal sensors via PLDM. The sensor was not visible via BMC.
- Resolved an issue where CVL4.1 and CVL4.2 had a problem with Recovery Clock configuration on port 4, in the case of fully reversed PF<>MAC mapping. Other configurations were not impacted.
- Resolved an issue where the NC-SI "Reset Channel" (0x05) command was not reverting the port link state change made by for Real Time Comprehensive Embedded Management "Network Port Control" (0x34) command.
- Resolved an issue where, during communication via PLDM Type 6, FW did not move the operation to the failed state and did not return `INVALID_DATA`. Instead, it returned `UNSUPPORTED` in the case of an invalid payload.
- Reverted an incorrectly added strict handling of the payload-less action.
- Resolved an issue where the Network Device and Port Metrics are disabled due to an BMC issue and missing support on BMC side.
- Resolved an issue where in order to change the Forward Error Correction setting in the UEFI HII the user needs to change the **Media detection** setting, save the changes, and then change the FEC setting.
- Previously, promiscuous mode did not see all packets; it only saw those packets arriving over the wire (that is, not sent from the same physical function (PF) but a different virtual function (VF)).
- Resolved an issue where during the second time of a POST operation, `NetworkAdapter.ResetSettingsToDefault` failed with the BMC reporting "Internal Server Error."
- Resolved an issue when both RMI-based transport (RBT) and Management Component Transport Protocol (MCTP) were enabled, and PT was enabled on RBT, then OEM command 0x26 (config LLDP) was being rejected from NC-SI over MCTP with reason code 0x5081.

- Corrected an issue where updating an NVM image of E810 could result in a "PHY NVM is not supported in devlink" warning message.
- Previously, E8102CQDA2 adapter showed Linkcap width as x16.
- Support for `ioctl` in interactive mode was added in version 1.39.45.2
- Previously during MCTP Endpoint discovery, an incorrectly interpreted BDF led to sending PLDM events to non-existing devices.
- Resolved an issue where on some systems, after making the NVMUpdate with delayed reboot, FW reset (EMPR) was required to complete the update and might fail.

2.1.1.4 VMware Driver

- Resolved an issue where excessive VF reset might cause multicast ping from VF to fail.
- Resolve the issue where the output of the PTP signal might have a period different than the one configured.
- Previously, encapsulated packets with inner packet padding were reported as packets with incorrect checksum. Driver collected these statistics and passed to the networking stack, which resulted in high pNIC error alarm raised by the OS.

2.1.1.5 FreeBSD Driver

- Resolved an issue where during successive driver unload/load cycles, single-port adapters might experience initialization failure when 5-layer topology is enabled.
- Resolved an issue where at the end of an **NVMUpdate** or **NURA** tool execution, after the NVM has finished updating, the tool might indicate that an EMP reset failed while the adapter continues to work normally and **NVMupdate** indicates that the update was successful.
- Resolved several errors seen in **dmesg** when trying to re-establish link after an **nvmupdate**.

2.1.1.6 QV Tools

- Corrected an issue where **eeupdate** could not update the serial mac address of E810 when port 0 mac address was changed.

2.1.2 Intel® Ethernet 820 Series

2.1.2.1 Firmware/NVM/NVM Update

- Resolved an issue where a PHY DOWNGRADE configuration file entry was incorrectly handled in some scenarios.
- Resolved an issue where cage power limit was not properly reported when NVM had a lower limit than the netlist, which could potentially could result in a power over-budget.
- Resolved issue where the AQ commands Get Input Frequency List (0x0C6C) and Get Output Frequency List (0x0C6D) returned an empty list of allowed frequencies. The commands now return the correct list of allowed frequencies.
- Resolved an issue where Real Time Comprehensive Embedded Management NC-SI command "Get Temperature" (0x13) was always responding with Response Code: Command Completed (0x3), Reason Code: Unknown/Unsupported Command Type (0x7fff), without providing values.
- Resolved an issue where a power over limit was not reported in the 'Invalid node max power limit' bit (Byte1, bit7) of the Get Link Status AQ output.

- Resolve an issue where the Subvendor ID and Subsystem ID in pcie space was not overridden with the value configured in IDEEPROM after upgrade MezzCard EEPROM. Now SVID and SSID application flow is implemented in FW.
- Resolved an issue where it was possible to use out-of-range clock_ref values in commands 0x0C62/0x0C63/0x0C64/0x0C65. Now input is allowed only within range of 0x00-0x1F and everything outside this range will cause EINVAL.
- Resolved an issue where using the EPCT tool, after switching the QSFP PKVL VMC to 4x25, resulting in the driver no longer being able to attach to interfaces. The only way to restore the device was to re-flash the NVM using LanConf.
- Resolved an issue where you might encounter the following situation:
 - Update is done on a device with an external programmable PHY FW.
 - Update is for a new MAC FW, but PHY FW included in the update is identical to that already on the device.
- Resolved an issue where changing the FEC value from BaseR to RS resulted in an error message in **dmesg**, and might result in link issues.
- Resolved an issue where all POR CFGs with the Coppervale (CPVL) Phy like CFG2.1, CFG5.x and CFG7.0-CPVL wouldn't link after forcing link mode with ethtool.
- It is no longer expect that the Single Root I/O Virtualization interface (SR-IOV) might fail when running with a 50 G image on 100 G Si.

2.1.2.2 Linux Driver

- Resolved an issue where after a power reset, or update, FW phy/link commands decline requests during this time resulting in PHY NVM update failing due to incomplete link initialization procedure.
- Resolved an issue where DefaultQueueVmmqQueuePairs for the virtual switch could not be set to more than 2 for default/main vport. Now, the number of queue pair for the default vport provided during creating a virtual switch is honored and setting the queue count accordingly.
- Resolved an issue where a **ptp4l** run failure with multi-master was causing a limitation, and the user needed to avoid running multiple **ptp4l** instances in one quad.
- Resolved an issue with **ptp4l** run timeout in non-pf0 ports with the ice 1.11.3 driver.

2.2 Intel® Ethernet 700 Series Network Adapters

2.2.1 VMware Driver

- Previously, in ESXi7U3 the X520 can not setup 1 G connection by default setting.

2.2.2 Firmware/NVM/NVM Update

- Resolved an issue where a rare PCI speed degradation with Gen3 to Gen2 was seen during extensive link tests after a reset is caused by a core reset on an ICX-D platform.
- Resolved an issue where in RDE's LLDPEnabled implementation, the function responsible for the NVM write does not check the validity bit of the settings saved in NVM. This caused a skip of write to NVM in some cases.
- Resolved an issue where during a write, with http PATCH method triggered, the LLDP Disable Valid bit was ignored, which resulted in a skip of write in all scenarios instead of just for the same LLDP Enabled value.

- Resolved an issue where if BMC bonding was disabled, NCSI mode was set to a Manual Switch Mode (depending on the platform), and BMC share port has changed before doing a power cycle, recovery mode would occur. Other prerequisites include disabled WOL, disabled LLDP, enabled PXE no drop, and disabled auto power-on after power loss (in BIOS).
- Fixed an issue where the processing port id was always the same as the mac address.
- Previously, the addition of PCIe re-timers added to the total channel latency causing replay timer timeout correctable errors.
- Resolved an issue where PLDM type MCTP packets were not accepted due to missing configuration for PCIe VDM.

2.2.3 Windows Driver

- Resolved an issue where the following message was appearing in the system log:
For Intel I350 NIC, "Source:NDIS, ID:10400" of warning log will be registered on system log on Windows Server 2019.
- Previously, DIRID **DefaultDestDir** values were showing 11 on Windows 11 instead of 13.

3.0 Known Issues

3.1 Intel® Ethernet 800 Series Network Adapters

3.1.1 Intel® Ethernet 810 Series

3.1.1.1 General

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

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

```
intel_iommu=off
```

```
noiommu:
```

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

- Properties that can be modified through the manageability sideband interface **PLDM Type 6: RDE**, such as **EthernetInterface->AutoNeg** or **NetworkPort->FlowControlConfiguration** do not possess a permanent storage location on internal memory. Changes made through RDE are not preserved following a power cycle/PCI reset.
- Link issues (for example, false link, long time-to-link (TTL), excessive link flaps, no link) might occur when the Intel® Ethernet Connection C827 Series and Intel® Ethernet Connection XL827 Series (C827/XL827) retimer is interfaced with SX/LX, SR/LR, SR4/LR4, AOC limiting optics. This issue is isolated to C827/XL827 line side PMD RX susceptibility to noise.
- Intel® Ethernet 800 Series Network Adapters in 4x25 GbE or 8x10 GbE configurations are limited to a maximum total transmit bandwidth of roughly 28 Gbps per port for 25 GbE ports and 12 Gbps per port on 10 GbE 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 25 GbE or 10 GbE 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.

- 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 third 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.

3.1.1.2 Firmware/NVM/NVM Update

- The PLDM Type 5 Command "Activate Firmware" can potentially cause a subsequent PLDM Type 2 event.
- During an ActiveFirmware request coming from Update Agent (BMC) PLDM Type2, a PlatformEventMessage is generated by the Firmware Device (Network Adapter). As a result, the response does not happen, and Update Agent (BMC) treats it as Update Failed.
- The second part of the string from the option ROM indicates the correct slot (port) numbers.

3.1.1.3 Linux

- When using devlink-rate and custom TX scheduler tree topology in switch mode, setting a node's/leaf's tx_priority ('devlink port function rate set pci/<PF_PCI>/<node_name> tx_priority 1') along with a BW limitation (for example, tx_max - 'devlink port function rate set pci/<PF_PCI>/<node_name> tx_max 1Gbit' - on the parent node) can result in starvation of the node/leaf (0 BW on Tx), with a lower priority set, and trigger a Tx hang detection. This can cause a reset, disrupting the custom topology (for example, reverting the affected node/leaf to the default parent node).

Workaround: To prevent this, a tx_max with a value smaller than the parent's tx_max can be configured for a higher priority node/leaf, leaving bandwidth available for lower priority ones.

- DPDK traffic is stopped after FLR reset. This issue has been documented in the **rte_eth_dev_reset** API.

Workaround: **testpmd** can be used to recover a VF after a reset.

- When a VF reset happens, **testpmd** will print out "port reset" event to the console.
- Use the "port reset" command to call **rte_eth_dev_reset**, and everything will go back to normal.

- There is a compatibility issue between Real-Time Linux kernel-3.10.0 and ice-1.10.x and later versions. It is suggested to use 4.x/5.x rt kernel.
- The Celo process may not be ended or killed while exiting the application. As the result, the current console is non-responsive. The stability of the system is not endangered, and the user can start the next console session.
- Changing the inner or outer VLAN tag protocols after setting the private flag `vf-true-promisc-support` disables the promiscuity on the VF's VLAN interfaces.

- When DCB configuration is changed to use the firmware agent in willing mode on the first port of a two port card, **lldpad** core dumps, and the second port of the reports the following messages:

```
[3784.370726] ice 0000:af:00.1: Commit DCB Configuration to the hardware
[3784.393109] INFO: Flow control is disabled for this traffic class (0) on this vsi.
[3784.458480] 8021q: Adding VLAN 0 to HW filter on device eth3
```

The messages repeat forever.

Workaround: Insure the DCB configuration to use the firmware agent in willing mode is on the second port as well as the first.

- When the number of queues on the VF interface is changed using the **ethtool** command while traffic is flowing from client to SUT (using *ice* driver version 1.10.X and *iavf* version 4.6.X), the system reboots.
- Host system might hang when resuming from sleep or hibernate with active VFs.
Unexpected system hand or loss of network functionality can occur when waking from sleep or hibernation if VFs (Virtual Functions) are active. To restore normal operation, the system must be power cycled.
- 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.
- It might not be possible to create the maximum number of supported RDMA VFs. Attempting to create greater than 20 RDMA VFs will result in no RDMA devices being created for VFs.
- Double VLAN traffic might RSS into the first queue. If configuring VLAN interfaces on PF in a way that results in double VLAN tagging, received double VLAN packets will be concentrated on the first queue of interface.
- When two VFs created from the same PF are assigned identical MAC addresses, they will not be able to pass traffic successfully unless the VF spoof check is disabled on the VF interfaces.
- The Intel® Ethernet 800 Series Network Adapter in eight-port 10 Gb configuration, the device might generate errors such as the example below on Linux PF or VF driver load due to RSS profile allocation. Ports that report this error will experience RSS failures resulting in some packet types not being properly distributed across cores.
 - Workaround:** Disable RSS using the `--disable-rss` flag when starting DPDK. Afterwards, only enable the specific RSS profiles that are needed.

dmesg: VF add example

```
ice_add_rss_cfg failed for VSI:XX, error:ICE_ERR_AQ_ERROR
VF 3 failed opcode 45, retval: -5
```

DPDK v20.11 testpmd example:

```
Shutting down port 0...
Closing ports...
iavf_execute_vf_cmd(): No response or return failure (-5) for cmd 46
iavf_add_del_rss_cfg(): Failed to execute command of OP_DEL_RSS_INPUT_CFG
```

- After changing link speed to 1 G on the E810-XXVDA4, the PF driver cannot detect a link up on the adapter. As a workaround the user can force 1 G on the second side.

- When using bonding mode 5 (i.e., balance-tlb or adaptive transmit load balancing), if you add multiple VFs to the bond, they are assigned duplicate MAC address. When the VFs are joined with the bond interface, the Linux bonding driver sets the MAC address for the VFs to the same value. The MAC address is based on the first active VF added to that bond. This results in balance-tlb mode not functioning as expected. PF interfaces behave as expected.

The presence of duplicate MAC addresses might cause further issues, depending on your switch configuration.

- If single VLAN traffic is active on a PF interface and a CORER or GLOBR reset is triggered manually, PF traffic will resume after the reset whereas VLAN traffic might not resume as expected. For a
 - **Workaround:** Issue the ethtool command: **ethtool -K PF_devname rx-vlan-filter off** followed by **ethtool -K PF_devname rx-vlan-filter on** and VLAN traffic will resume.
- Adding a physical port to the Linux bridge might fail and result in Device or Resource Busy message if SR-IOV is already enabled on a given port. To avoid this condition, create SR-IOV VFs after assigning a physical port to a Linux bridge. Refer to *Link Aggregation is Mutually Exclusive with SR-IOV and RDMA* in the *ice* driver README.
- When using a Windows Server 2019 RS5 Virtual Machine on a RHEL host, a VLAN configured on the VF using **iproute2** might not pass traffic correctly when an *ice* driver older than version 1.3.1 is used in combination with a *ivf* driver version.
- It has been observed that when using iSCSI, the iSCSI initiator intermittently fails to connect to the iSCSI target.
- With the current *ice* PF driver, there might not be a way for a trusted DPDK VF to enable unicast promiscuous without turning on "ethtool --priv-flags" with "vf-true-promisc-support".
- If a VLAN with an Ethertype of 0x9100 is configured to be inserted into the packet on transmit, and the packet, prior to insertion, contains a VLAN header with an Ethertype of 0x8100, the 0x9100 VLAN header is inserted by the device after the 0x8100 VLAN header. The packet is transmitted by the device with the 0x8100 VLAN header closest to the Ethernet header.
- A PCI reset performed on the host might result in traffic failure on VFs for certain guest operating systems.
- On RHEL 7.x and 8.x operating systems, it has been observed that the `rx_gro_dropped` statistic might increment rapidly when Rx traffic is high. This appears to be an issue with the RHEL kernels.
- 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 100 Gbps, then the PFs cannot guarantee the minimum Tx rate limits set for SR-IOV VFs.
- Some distros might contain an older version of **iproute2/devlink** which might result in errors.
 - **Workaround:** Update to the latest **devlink** version.

- When DCB configuration is changed to use the firmware agent in willing mode on the first port of a two port card, the **Ildpad** core dumps, and the second port reports the following messages:
 - [3784.370726] ice 0000:af:00.1: Commit DCB Configuration to the hardware
 - [3784.393109] INFO: Flow control is disabled for this traffic class (0) on this vsi.
 - [3784.458480] 8021q: adding VLAN 0 to HW filter on device eth3

The messages repeat forever.

Workaround: Ensure the DCB configuration is set to use the firmware agent in willing mode on the second port as well as the first.

- When Large Send Offload (LSO) V2 is enabled, the network adapter is unable to transmit frames larger than the MTU, which can impact network performance. Additionally, the incorrect incrementing of checksums `OID_INTEL_OFFLOAD_LARGE_SEND_VXLAN_COUNT` may lead to inaccurate network statistics.

Workaround: Users can temporarily disable Large Send Offload V2 on their network adapters to allow the transmission of frames larger than the MTU. However, note that this workaround may impact other aspects of network performance. We recommend using this workaround only if absolutely necessary and awaiting the software update for a comprehensive solution.

- During TC configuration, using the "ethtool -S <vf_interface>" command results in a crash due to invalid memory access during reconfiguration of queues.

3.1.1.4 FreeBSD Driver

- During traffic in RoCEv2 mode, if using a large number of QPs (>64), a PE Critical Error may occur. In such circumstances the card may become nonoperational, and reboot is required to restore RDMA capability.
- The driver can be configured with both link flow control and priority flow control enabled even though the adapter only supports one mode at a time. In this case, the adapter will prioritize the priority flow control configuration. Verify that link flow control is active or not by checking the **active:** line in `ifconfig`.
- IAVF virtual interfaces in FreeBSD-13.0 guests might experience poor receive-performance during stress.
- Unable to ping after removing the primary NIC teaming adapter. The connection can be restored after restarting the VM adapters. This issue is not observed after the secondary adapter is removed, and is not OS specific.
- The visibility of the iSCSI LUN is dependent upon being able to establish a network connection to the LUN. In order to establish this connection, factors such as the initialization of the network controller, establishing link at the physical layer (which can take on the order of seconds) must be considered. Because of these variables, the LUN might not initially be visible at the selection screen.
- Intel® Ethernet Controller E810 devices are in the DCBX CEE/IEEE willing mode by default. In CEE mode, if an Intel® Ethernet Controller E810 device is set to non-willing and the connected switch is in non-willing mode as well, this is considered an undefined behavior.
 - **Workaround:** Configure Intel® Ethernet Controller E810 devices for the DCBX willing mode (default).
- In order to use guest processor numbers greater than 16 inside a VM, you might need to remove the `*RssMaxProcNumber` (if present) from the guest registry.

3.1.1.5 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.
- Running Unreliable Datagram (UD) RDMA mixed traffic with more than two QPs might lead to a receiver side UD application hang. To recover, restart the RDMA UD application. This is not expected to impact storage (NVMeoF, iSER, VSAN) applications because they do not rely on UD communication.
- 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.
- On RHEL 7.9, installing `rdma-core v35.0 debuginfo` rpms can prevent the installation of `debuginfo` rpms from indistro products like `libfabric`.
- When using Intel® MPI Library in Linux, Intel recommends to enable only one interface on the networking device to avoid MPI application connectivity issues or hangs. This issue affects all Intel® MPI Library transports, including TCP and RDMA. To avoid the issue, use `ifdown <interface>` or `ip link set down <interface>` to disable all network interfaces on the adapter except for the one used for MPI. OpenMPI does not have this limitation.
- VLAN priority for Unreliable Datagram (UD) traffic is incorrect if supplied ToS is not set to Priority 0.

3.1.1.6 VMware Driver

- Using Native Mode and ENS Mode ICEN driver with the latest DDP can cause queue configuration issues.
- When instantiating the maximum number of VFs in NSX-T, adding a Transport Node afterwards might fail due to timeout.
- Configuring the NSX-T Virtual Distributed Switch uplink port might fail when SR-IOV is enabled in the PF.
- Setting LFC for PF might fail.
- Received packets with incorrect length can generate alarms in VMware ESXi. These alarms can be ignored. See the following article for more details: <https://kb.vmware.com/s/article/83627>
- When configuring a switch to use IEEE LLDP version for DCB, the PF host driver is unable to change the CEE LLDP version for DCB, even if the switch is configured for CEE LLDP.
- When running Release 2.2 NVM drivers on Release 3.2 NVMs, users might encounter warning messages regarding Null pointer errors. These are expected warnings when running older drivers on newer NVMs.
- When entering the Pause Parameter via the CLI, related configurations in quick or rapid succession could cause a configuration failure or unexpected results. In NSX-T 3.1.0, a Guest Virtual Machine associated with ENS NSX-T Virtual Distributed Switch (NVDS) might experience guest operating system kernel panic when receiving TCP traffic with VXLAN overlay. VMware ESX 7.0 operating system with NSX-T 3.1.0 might experience a kernel panic (also known as PSOD) when changing NUMA node in NSX-T Virtual distributed switch. Rebooting a Red Hat 8.2 Linux VF VM multiple times

might cause traffic to stop on that VF.A VLAN tag is not inserted automatically when DCB PFC is enabled on an interface. This might cause RDMA issues if no VLAN is configured.

- **Workaround:** Since PFC for *icen* is VLAN-based, create a VLAN tag for DCB to be fully operational. After a PF Reset, Windows VF traffic might fail.

3.1.1.7 Application Device Queues (ADQ)

The code contains the following known issues:

ADQ Standard Issues

- ADQ dynamically registers and unregisters queue group-specific devlink parameters. Newer kernels (6.2.x and later) expect all devlink parameters to be registered during probe time and unregistered upon device removal. As a result, warning messages are to be expected when ADQ queue groups are created or destroyed for the first time.
- 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).
 - **Workaround** - Recommend configuration of only one application per Traffic Class (TC) channel.
- ADQ does not work as expected with NVMe/TCP using Linux kernel v5.16.1 and later. When **NVMe connect** is issued on an initiator with kernel v5.16.1 (or later), a system hang might be observed on the host system. This issue is not specific to Intel® Ethernet drivers, it is related to NVMe changes in the 5.16 kernel. Issue can also be observed with older versions of the *ice* driver using a 5.16+ kernel.
- When using E810 4.1 NVM along with E810 2.4 *iavf* drivers, the ADQ VF traffic does not hit either the default TC or the ADQ TC queue set when using TCP protocol. However, traffic utilizing the UDP protocol hits the default TC queue set, not the ADQ TC queue set.

3.1.2 Intel® Ethernet 820 Series

3.1.2.1 General

- During validation, an issue was discovered in Windows Server 21H1. This OS version is unable to save a memory dump (crash dump) to a disk. It is considered to be OS defect.
- We have noted certain native VF drivers in combination with this release's PF driver will cause it to not bond.
- With a Windows host and Linux virtual machine (VM), the last transmit (Tx) queue might not increment when there are multiple Tx queues.
- **Insufficient PCI-Express bandwidth available for device** might be logged for Intel® Ethernet E820 Series Network Adapters. The E820 Series does not use a PCI-Express interface and this appears to only be a logging issue.
- The 4x25 G NVM lists 100 GbE and 50 GbE 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/NVM/NVM Update

- 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).
- The DCB-MAP Configuration is not displaying on the SUT Interface from the Extreme Switch after enabling the firmware Mode in the SUT.
- Using the EPCT tool to change port configuration requires 2 reboots to complete the programming process. Blank mode or PTP initialization failures may be observed after a single reboot and will be resolved after performing a second reboot.
- The **Get LLDP** command (0x28) response use to contain only 2 TLV types. The new implementation requires a third TLV type, which contains a copy of the whole LLDP frame payload.
- The 100 MB option, is visible in Windows* Device Manager. However, when it is selected, a link cannot be established.

3.1.2.3 Linux Driver

- Driver load may fail in SLES12SP5 and cause call traces in dmesg. If using SLES12S95, this issue may be seen when upgrading the ICE driver version, as this is a limitation of SLES12SP5.
- With RDMA enabled, the number of MSIXs available in SLES12SP5 is less than the ones requested. Workaround: Continue with earlier releases of ICE Driver.

ICE Driver Version	SLE Version	Working	Not Working	Comments
1.10.5	SLE12SP5	X		Working because RDMA is disabled.
1.9.12	SLE12SP5	X		Working because RDMA is disabled.
1.11.7	SLES12SP5		X	Not working because RDMA is enabled and it requests additional MSIXs.
1.11.RC72	SLE15SP4	X		Working since the Distro supports more availability of MSIX interrupts.

- **Ethtool** module output is printing offset with values instead of human-readable module information and details.
- Current netlist support on SyncE platforms will show support for 2x25 and when selected will only enable 1x25. This is in error as the 2x25 configuration is not supported on this platform.
- You may see an error message after re-seating the network that the module is not present. This is a logging issue. The port is connected and functioning correctly.
- It has been observed in certain scenarios where inner VLAN traffic will concentrate on a single queue creating packet drops.
- Inventory output doesn't report **MinSrev** update. The issue is observed only when `-optin` is requested together with the NVM update. If the NVM is updated first, and after this request `-optin` it should be successful. Additionally we can still use the **iiostl** interface to update the NVM image and update **MinSrev** together.
- DCB-MAP configuration is not reflecting on the SUT Interface from Extreme Switch after enabling the Firmware Mode in the SUT.
 - **Note:** Software Mode can be a possible workaround for this issue.
- Virtual Functions (VF) do not run on one of CPK Physical Functions (PF). The Single Root I/O Virtualization (SR-IOV) cannot be used on one of PFs.
- In a double VLAN setup with set to promiscuous mode, packets are not seen in Wireshark on the expected ports.
- After assigning a Locally Administered Address (LAA), the system can still wake from S5 by using the Burned In Address (BIA) but does not wake up if the LAA is used.
 - **Workaround:** use BIA for waking the system from S5.

- On RHEL 7.9 VMs, VF traffic does not resume after the VF's MAC address is changed on the host side. This appears to be a limitation with RHEL 7.9.

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

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

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

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

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

- DCB-MAP Configuration is not reflected from switch on Intel® Ethernet Connection C827 Series Port with CEE and SW Mode on SLES15 SP3 OS.
- **Module is not present** error message is displayed after loading the *ice* driver with cages filled.
- All ports link is not coming up after updating driver with WRCP 22.12 OS. Also an error is encountered when trying to make link up RTNETLINK answers: Input/output error.

Workaround: Resolve the issue using the following workaround.

- Copy to `/lib/firmware/updates/intel/ice/ddp`
- In **-sf ice-1.3.26.0.pkg ice.pkg**
- **rmmod ice**
- **modprobe ice**

3.1.2.4 FreeBSD Driver

- 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 Driver

- Some adapters are disabled after changing MTU on Switch Embedded Teaming vSwitch
 - The issue takes place on LCC platform
 - The issue does not take place on HCC platform
 - The issue takes place on Switch Embedded Teaming switch with Single Root I/O Virtualization:
New `-Vmswitch -EnableEmbeddedTeaming $true -EnableIov $true`
 - The issue does not take place on the switch without Embedded Teaming: by default -
`EnableEmbeddedTeaming` is None
 - The issue takes place after MTU size was changed 5-6 times with 1 minute intervals between changes (one test iteration contains 3 MTU size changes)
 - The issue takes place after MTU size was changed 50-60 times with 10 minute intervals between changes

Based on our research we recommend to avoid Switch Embedded Teaming if possible, or increase timeout between MTU size changes.

3.1.2.6 LANConf Tool

- There is a limitation of the LANConf tool that only allows the user to program ID EEPROM on Quad 0. This also means that each attempt to program ID EEPROM on Quad 1 will cause an update on Quad 0.

3.2 Intel® Ethernet 700 Series Network Adapters

3.2.1 General

- A rare PCI speed degradation with Gen3 to Gen2 has been seen during extensive link tests after a reset is caused by a core reset on an ICX-D platform.
- Setting PHY debug mode with disabling link during initialization causes link to go down.
- The *DMARemappingCompatible* registry is missing in the previous driver. The *DMARemappingCompatible* set as '1' res is applicable to all client drivers (1G/10G).

3.2.2 Windows Driver

- **i40ea** version 1.16.130 and later provides low throughput numbers in Windows 2022 Hyper-V VMs when vSwitch is set to Switch independent teaming interface without SRIOV enabled.
- When drivers are installed in Windows 2019 Server, Windows Event ID 411 may be reported in the Windows Event Log.

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

- Do not connect Intel® Ethernet Controller V710 based cards to switch Edge-core 5812-54T-O-AC-B because this switch does not work well with them.

3.2.4 Linux Driver

- **ptp4l** since v1.8 has supported monitoring link state, and will not reset after detecting a fault when the link status of the port is down. However, if `fault_reset_interval` is configured as ASAP (or its numeric equivalent of -128), **ptp4l** will attempt to reset immediately without checking the link status. This causes **ptp4l** to continuously fault and reset as the link for the port is still down. This results in a clear and unexpected behavioral difference when using ASAP vs when using another `fault_reset_interval`. A fix was proposed to the PTP4l Opensource project, but has not yet been accepted by the LinuxPTP community
- Flow director outer MAC L2 filter is not able to direct traffic to queues/VFs.

3.2.5 VMware Driver

- On some hosts with an AMD CPU, there is no Tx traffic with multiple Intel® Ethernet Network Adapter X710 ports connected to one vSwitch.
 - **Workaround:** Do not connect multiple Intel® Ethernet Network Adapter X710 ports to one vSwitch.
- Multicast packets sent from VF are not visible on **vmkernel** interface connected to the first NIC port in ENS mode. The problem results in packet drop only on **vmkernel** interface (packets are not seen). This does not affect vf-vf traffic, and this is only an issue when sending multicast packets from VM to **vmkernel**.

3.2.6 Firmware/NVM/NVM Update

- A core reset clears the CSRs during the block read of MAC CSR, resulting in stalling of other CSR read operations.
- 40 G QSFP modules from Intel can not support NC-SI OEM command 0x4b02 to query temperature with reason code 0x5089.

3.3 Intel® Ethernet 500 Series Network Adapters

None for this release.

3.4 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 [Feature Support Matrix](#) for more detail. FSMs list the exact feature support provided by the NVM and software device drivers for a given release.

5.0 Languages Supported

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

This release supports the following languages:

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

6.0 Related Documents

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

6.1 Feature Support Matrix

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

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

6.2 Specification Updates

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

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

6.3 Software Download Package

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

6.4 SourceForge Ethernet Drivers and Utilities

For additional information regarding Linux kernel drivers, 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>

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 that are referenced in this document can be obtained by visiting the [Intel Resource and Documentation Center](#).

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.

© 2023 Intel Corporation.