Intel® Ethernet Controller Products

Release Notes for Software Release 25.3

Intel Corporation

Revision 1.0
September 2020
631171-001

Revision History

Revision	Date	Comments
1.0	September 2020	• Initial release (Intel Public).

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 Supported Intel® Ethernet Controller Devices

Note: **Bold Text** indicates the main changes for Software Release 25.3.

Controller Family	Windows* Driver Version	Linux* Driver Version	FreeBSD Driver Version	NVM Version
Intel® Ethernet 800 Series	**1.4.38.x**	**1.1.x**	**0.26.x**	**2.1x**
Intel® Ethernet 700 Series Intel® Ethernet Network Connection X722	1.12.177.x	2.12.x	1.12.x	8.0 N/A
Intel® Ethernet 500 Series Intel® Ethernet Controller X550	4.1.197.x	5.8.x	3.3.x	N/A 3.0
Intel® Ethernet 300 Series	12.18.11.x	5.3.x	2.5.x	N/A
Intel® Ethernet 200 Series	12.19.0.x	5.3.x	2.5.x	N/A
Intel® Ethernet 225 Series	**1.01.x**	**N/A**	**N/A**	**N/A**
Intel® Ethernet Controller 82598	2.4.36.x	5.8.x	3.3.x	N/A
Intel® Ethernet Controller 8257x	9.16.10.x	5.3.x	2.5.x	N/A

For help identifying your network device and finding supported devices, click here:

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

1.2 New Features

1.2.1 Hardware Support

Release	New Hardware Support
• 25.3	• Support for the following devices: — Intel® Ethernet Controller E810-XXV. — Intel® Ethernet Controller E810-XXV for backplane. — Intel® Ethernet Controller E810-XXV for SFP. — Intel® Ethernet Network Adapter E810-XXV-2. — Intel® Ethernet Network Adapter E810-XXV-2 for OCP 3.0. — Intel® Ethernet Network Adapter E810-C-Q2 for OCP 3.0. • Linux support for the following devices: — Intel® Ethernet Network Adapter E810-C-Q1 for OCP. • Microsoft* Windows* support for the following devices: — Intel® Ethernet Controller I225-IT. — Intel® Ethernet Controller I225-LM. — Intel® Ethernet Controller I225-LMvP. — Intel® Ethernet Controller I225-V.

1.2.2 Software Features

Release	New Software Support
• 25.3	• Support for Red Hat Enterprise Linux 7.9.

1.2.3 Firmware Features

Added support for devices listed in Section 1.2.1.

1.2.4 Removed Software Support

- Software Release 25.2 was the last release that supported the 32-bit Microsoft* Windows* 10 e1d driver. This affects devices based on the following controllers:
  - Intel® Ethernet Connection I217-LM
  - Intel® Ethernet Connection I217-V
  - Intel® Ethernet Connection I218-LM
  - Intel® Ethernet Connection I218-V
  - Intel® Ethernet Connection I219-LM
  - Intel® Ethernet Connection I219-V

Note: 64-bit Microsoft Windows 10 is still supported.

- The Linux e1000e driver has been removed from this release and all future releases. Support for Linux e1000e has been moved to the community. The latest drivers can be downloaded from SourceForge.

1.2.5 Removed Features

No previous features were removed for this release.

2.0 Fixed Issues

2.1 Intel® Ethernet 800 Series

2.1.1 General

- **Fixed:** The link might go down when Intel device drivers load using auto-speed (default) with non-specification compliant modules. Manually setting the speed might recover the link.
- **Fixed:** Applications using the Intel® Ethernet Controller E810 in a backplane configuration with no WoL, and no management during power-up might experience no link, which is fixed by any form of reset.

2.1.2 Firmware

- **Fixed:** The device firmware might enter recovery mode during stress testing of PCIe in-band resets where many resets are done back-to-back. This is not observed in normal operation but can be seen by running the PCI-SIG configuration verification test under stress.
- **Fixed:** RDMA MPI stress might crash with eight or more systems in cluster use.

2.1.3 Linux Driver

- **Fixed:** DPDK sample applications specify a packet size of 1518 bytes. At present the PF driver does not take into account the extra 4 bytes needed if a port VLAN is configured. In this situation (when a port VLAN is configured) full sized packets of 1500 do not reach the DPDK application. In order for this to work properly, initialize the max packet length within the DPDK application to account for the possible port VLAN (specify 1522 bytes).
- **Fixed:** When in an 8x10 configuration and using firmware LLDP, the 8TC to 4TC mapping in CEE mode does not work.
- **Fixed:** Operating system traffic is not passed when using a VLAN over bridge topology. To workaround this issue, set the `rx_vlan_filter` to off and traffic will resume as expected.
- **Fixed:** IEEE1588/PTP 1PPS does not work on Intel® E810 series adapters.
- **Fixed:** The `dmesg` log might show ICE_ERR_AQ_TIMEOUT errors due to know issues with the driver. The feature might not work as expected when using Release 25.2 driver which has now been fixed.

2.1.4 Linux RDMA Driver

- **Fixed:** When using RoCEv2 on newer kernels, some iSER operations might experience errors related to iSER's handling of work requests. To work around this issue, set the ice fragment_count_limit devlink parameter to 13. Refer to the Devlink Configuration section of the irdma driver README for details on setting the devlink parameter.

2.1.5 NVM Update Tool

- **Fixed:** Intel recommends not enabling non-default modes (delayed reboot and selective preservation) with this release. Generally when using these features, errors might be reported; however, after a reboot the update should have completed successfully.

2.1.6 FreeBSD Driver

- **Fixed:** When using the FreeBSD ice driver and the enable/disable LLDP agent status sysctl, enable/disable on one port might enable/disable on other ports.

2.1.7 Manageability

- **Fixed:** When using PLDM Type 2 monitoring and control messages to read thermal sensor information the following can be observed:
  - Inserting modules post boot will not be detected.
  - When sending a Enable Channel NC-SI command when the device is in low power with link off, the link is not turned on.
  - Module temperature value is incorrect.
- **Fixed:** When sending Enable Channel NC-SI command when the device is in low power with link off, the link is not turned on.

2.2 Intel® Ethernet 700 Series

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

- Some devices support auto-negotiation. Selecting this causes the device to advertise the value stored in its NVM (usually disabled).
- Packet drops might occur under receive stress.

  Devices based on the Intel® Ethernet Controller 800 Series are designed to tolerate a limited amount of system latency during PCIe and DMA transactions. If these transactions take longer than the tolerated latency, it can impact the length of time the packets are buffered in the device and associated memory, which might result in dropped packets. These packets drops typically do not have a noticeable impact on throughput and performance under standard workloads.

  If these packet drops appear to affect your workload, the following might improve the situation:

  1. Make sure that your system's physical memory is in a high-performance configuration, as recommended by the platform vendor. A common recommendation is for all channels to be populated with a single DIMM module.
  2. In your system's BIOS/UEFI settings, select the **Performance** profile.
     - For Microsoft Windows Server* operating systems:

       In the **Power Options** control panel, set the system to use the **High Performance Profile**.
     - For Linux operating systems:

       On RHEL 7.x/8.x, use the tuned power management tool to set the latency-performance profile.

       In other operating systems and environments, use the equivalent tool to set the equivalent profile.
     - For FreeBSD operating systems:

       Do not modify the ACPI, P-state, or C-state system controls (`sysctls`) to values intended to save power on the system.
     - For VMWare ESX, instructions to set power management to High Performance in the operating system can be found here:

       https://docs.vmware.com/en/VMware-vSphere/6.7/com.vmware.vsphere.resmgmt.doc/GUID-F48D75C7-2461-4643-8A3A-B0383146F3AA.html#GUID-F48D75C7-2461-4643-8A3A-B0383146F3AA
- 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 Intel's part.

- 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.
- At lower packet rates flow control and/or packet drops might be observed.
- In order for an Intel® Ethernet 800 Series-based adapter to reach its full potential, you 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.
- Currently, installing SLES distros via PXE boot is not supported on Intel® Ethernet Controller E810-XXVAM2 devices. This limitation will be resolved in a future software release.
- 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.
- If you are having link issues (including no link) at link speeds faster than 10 Gb/s, check your 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 your switch to use RS-FEC mode.

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

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

3.1.2 Firmware

- When performing a firmware update using the return to factory defaults option, if there are added TLVs that are needed for firmware functionality, firmware might go into recovery mode. To work around this issue, perform an update using the no preservation option and then override the needed PFA settings with information from factory settings.

3.1.3 Linux Driver

- 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.

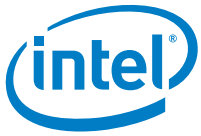

- Current limitations of minimum Tx rate limiting on SR-IOV VFs:
  - If DCB or ADQ are enabled on a PF then configuring minimum Tx rate limiting on SR-IOV VFs on that PF is rejected.
  - If both DCB and ADQ are disabled on a PF then configuring minimum Tx rate limiting on SR-IOV VFs on that PF is allowed.
  - 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.
- When using ice Linux driver and creating a flow director rule with ethtool the ethtool stat `fdir_sb_match` might not increment even though traffic is received matching on that rule.
- Creating MACVLAN interfaces might result in Device or Resource Busy message if SR-IOV is already enabled on a given port. To avoid this condition, create MACVLAN interfaces prior to enabling SR-IOV feature.
- On some server platforms running RHEL 7.8, there might be a call trace seen on the Intel IOMMU driver when the MTU is changed. A reboot of the system might be required in order to recover.
- 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 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.
- Intel's `iavf` driver in a Linux guest/virtual machine might occasionally experience driver crash/call trace during an install or uninstall on Windows Server 2016.
- Windows Server 2019 powershell command Get-IntelNetAdapterStatus -Status Virtualization -Name *e810* might result in an error message *The remote procedure call failed* when Linux virtual machine(s) are running.
  - When using `iproute2` commands or `ethtool` operations immediately after loading the module, there might be a race condition that might cause undefined behavior with the driver.
  - **Workaround:** Allow approximately five seconds after module load and before issuing `iproute2` or `ethtool` operations. The NetworkManager configuration might need to be updated to ignore ICE interfaces.

3.1.4 Windows Driver

- When changing the Locally Administered Address the Intel® Ethernet Controller E810 reflects the updated address but the switch might not show this address in its LLDP neighbors information.

- If the link is disabled on the link partner the Intel® Ethernet Controller E810 might not report the link is down after running Powershell or DMIX. If you disable then enable the Windows driver the issue should resolve.
- If you change the DCB map on the link partner while RDMA is enabled it might cause a reset on the device.

3.1.5 Windows RDMA Driver

- The Intel® Ethernet Network Adapter E810 might experience an adapter-wide reset (all ports) when the DCBx willing configuration change (in firmware managed mode) propagated from the switch removes a TC that was RDMA enabled. This typically occurs when removing a TC associated with UP0 because it's the default UP that RDMA based it's configuration on. The reset results in a temporary loss in connectivity as the adapter re-initializes.
- When connecting multiple rping clients to a persistent rping server, older kernels might experience a crash related to the handling of `cm_id` values in the kernel stack. With Intel® Ethernet Controller E810, this problem typically appears in the system log as a kernel oops and stack trace pointing to `irdma_accept`. The issue has been fixed in kernels 5.4.61 and later. For patch details, see https://git.kernel.org/pub/scm/linux/kernel/git/torvalds/linux.git/commit/drivers/infiniband/core/ucma.c?h=v5.9-rc2&id=7c11910783a1ea17e88777552ef146cace607b3c for more detail.
- 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 your 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 your 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 config file is not supported.
- In ESXi, the device serial number appears as N/A.
- Reboot might fail after an update if the PRESERVATION entry has NONE or SELECTIVE in config file. To workaround this issue, remove the entry and let the tool update using the default full preservation option.

3.1.8 FreeBSD Driver

- The FreeBSD ICE driver might show the wrong media type (25GBASE-AOC) for 25G-LR/SR optics when the advertise_speed parameter is set to 164 (1G+10G+25G).

3.1.9 Manageability

- When using the PLDM Firmware Update, the update might fail if there is CORER / GLOBR during the update as it resets the operation and the update need to be restarted. CORER/ GLOBR could be caused by the device driver or tools running on the host.

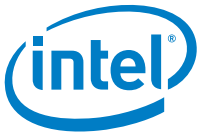

- The bandwidth/throughput might vary across different VFs if VF rate limiting is not applied.

  Workaround: To avoid this situation it is recommended to apply VF rate limiting.

- Occasionally, when a QSFP/ SFP+ module is inserted, PLDM PDR's are not created and a sensor is not accessible. If this occurs, unplug the module and try again. This has been observed only with a slow removal/insertion of the module.

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, you 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 Firmware

3.2.2.1 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

- On Windows Server 2012 R2 the host might suffer a BSOD after modifying ANS team type with tagged VLAN. It can be reproduced both with GUI or a Powershell command.

3.2.4 Linux Driver

- When the second subordinate interface is added to the bonding, its MAC address becomes the same as the first added subordinate and bond interface. Reloading the PF driver and creating VFs again is needed to restore the MAC addresses.
- On Kernel version 5.0.9 and higher, setting promiscuous mode on trusted VF leads to a periodic and endless update of this mode.
- Maximum number of SR-IOV VFs is 62 (instead of 63) per port. This behavior is observed when the RSS and VMDq settings are disabled. It is observed only when using ESX-ixgben driver.
- When two VF interfaces are in balance-tlb bonding mode, the ICMP traffic might be interrupted for a few sequences.

3.2.5 NVM Update Tool

- Minimum Security (Rollback) Revision (MinSREV) values and Security (Rollback) Revision (SREV) are not reported in unified (hexadecimal) format via the NVM Update Tool.

3.2.6 PROSet

- Start and Stop words are missing from the blinking button in the **Identify Adapter** tab. This behavior is observed only if the German language has been chosen.
- Use Default button in the UI is missing from the **Locally Administered Address** tab.
- VLAN status might be presented as Link Down in the **Teaming/VLANs tab** whereas the status is Enabled in the Windows Device Manager.

3.3 Intel® Ethernet 500 Series

- Maximum number of SR-IOV VFs is 62 (instead of 63) per port. This behavior is observed when the RSS and VMDq settings are disabled. It is observed only when using the ESX-ixgben driver.

3.4 Intel® Ethernet 300 Series

None for this release.

3.5 Intel® Ethernet 200 Series

None for this release.

3.6 Legacy Devices

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

http://www.intel.com/support/go/network/adapter/nicoscomp.htm

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 French German Italian Japanese	Spanish Simplified Chinese Traditional Chinese Korean Portuguese

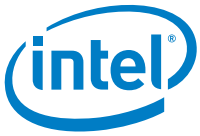

6.0 Related Documents

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	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/515715
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 at:

https://www.intel.com/content/www/us/en/support/products/36773/network-and-i-o/ethernet-products.html

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

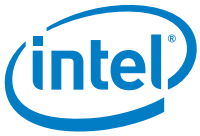

NOTE: This page intentionally left blank.

LEGAL

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

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

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

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

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

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

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

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

© 2020 Intel Corporation.