



Intel Ethernet Controller Products

Release 26.0 Release Notes

Intel Corporation

Revision 1.0
634103-001



Revision History

Revision	Date	Comments
1.0	December 2020	<ul style="list-style-type: none"><li data-bbox="576 411 865 436">• Initial release (Intel Public).

1.0 Overview

This document provides an overview of the changes introduced in the latest Intel® Ethernet controller/adaptor 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 26.0.

Controller Family	Windows* Driver Version	Linux* Driver Version	FreeBSD Driver Version	NVM Version
Intel® Ethernet 800 Series	1.6.21.0	1.3.2	0.28.1	2.3x
Intel® Ethernet 700 Series Intel® Ethernet Network Connection X722	1.14.67.0	2.14.x	1.13.x	8.2 N/A
Intel® Ethernet 500 Series Intel® Ethernet Controller X550	4.1.229.x	5.10.x	3.3.x	N/A 3.1
Intel® Ethernet 300 Series	12.18.11.x	5.5.x	2.5.x	N/A
Intel® Ethernet 200 Series¹	12.19.0.x	5.5.x	2.5.x	N/A
Intel® Ethernet 225 Series³	1.01.x	N/A	N/A	N/A
Intel® Ethernet Controller 82598³	2.4.36.x	5.10.x	3.3.x	N/A
Intel® Ethernet Controller 8257x³	9.16.10.x	5.5.x	2.5.x	N/A

1. Linux driver version update only 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>

1.2 New Features

1.2.1 Hardware Support

Release	New Hardware Support
• 26.0	• No new hardware was introduced in this release.

1.2.2 Software Features

Release	New Software Support
• 26.0	<ul style="list-style-type: none"> • Aggregated receive error counters for the i40e driver. • Aggregated receive error counters for the ixgbe driver. • i40e FreeBSD driver — When the VLAN_HWFILTER flag is enabled on an interface and there aren't enough hardware resources available to filter all the configured VLANs, the driver disables hardware filtering and enables the operating system or other software to filter VLAN tags.

1.2.3 Removed Features

Release	Hardware/Feature Support
• 26.0	No features were removed in this release.

1.2.4 Firmware Features

Release	New Firmware Support
• 26.0	• No new firmware features introduced in this release.

2.0 Fixed Issues

2.1 Intel® Ethernet 800 Series

None for this release.

2.2 Intel® Ethernet 700 Series

2.2.1 PROSet

- **Fixed:** When an Intel® Ethernet controller has connection and it's bounded to a virtual adapter the connection test results states "TCP/IP protocol not configured" instead of "Cannot run this test because the device is bound to a virtual NIC".
- **Fixed:** In the Adapter Information tab, *Negotiated link speed* text is truncated. Only first word is shown. It's exists only when the German language is used.
- **Fixed:** Team setting discrepancy observed between DMIX and PROSet ACU. A different default value of the check time (in seconds) is being used.

2.2.2 Linux Driver

- **Fixed:** With DCB enabled, when dynamically switching LLDP/DCBX from firmware to software ownership using `ethtool --set-priv-flags ethX disable-fw-lldp` on command, an issue results in `dmesg` error messages and incorrect DCB configuration. **Possible Workarounds:** 1. Reload the Linux driver and/or 2. Shut down and re-enable the switch port connected to the affected adapter port.
- **Fixed:** Using all queues to create Traffic Classes (TCs) via the Application Device Queues (ADQs) feature might cause loss of network connection.
- **Fixed:** `ip_link_spoof` option is not available when using the SR-IOV Hypervisor management interface (VF-d).
- **Fixed:** For the XDP redirect feature, when using the `ethtool -L combined [N]` option, traffic being received on the Tx side is almost half of the packets received on Rx. There is no corresponding packet drop happening when checking `netstat`.
- **Fixed:** For Precision Time Protocol (PTP), there is no synchronization when SMA1 (GPIO_4) is input for the 1PPS signal.
- **Fixed:** When VF is disabled via SR-IOV Hypervisor Management Interface (VF-d) it cannot be enabled again.

- **Fixed:** Running two VM's with multiple VF's might cause a Reboot with a General Protection Fault if the driver is unloaded while VF's are running traffic.
- **Fixed:** Express Data Path — On the newest kernels, system hangs might be observed after the eBPF program loads.
- **Fixed:** Auto negotiation cannot be disabled on the Intel® Ethernet Controller V710 for 5GBASE-T.

2.2.3 Linux RDMA Driver

- **Fixed:** There is approximately a 25% chance that a system crashes during RDMA CM connection and simple ping-pong test (`ucmatose`) when the number of connections is greater than 16000.

2.3 Intel® Ethernet 500 Series

- **Fixed:** On Kernel version 5.8.5 or higher call trace is observed during loading `xdp_redirect` or `xdp_tx`.
- **Fixed:** Inventory traffic can be lost on Intel® Server System R2208WFQ with Intel® Ethernet Network Connection OCP X527.

2.3.1 PROSet

- **Fixed:** When an Intel® Ethernet controller has connection and it's bounded to a virtual adapter the connection test results states "TCP/IP protocol not configured" instead of "Cannot run this test because the device is bound to a virtual NIC".
- **Fixed:** In the Adapter Information tab, *Negotiated link speed* text is truncated. Only first word is shown. It's exists only when the German language is used.

2.4 Intel® Ethernet 300 Series

2.4.1 PROSet

- **Fixed:** When an Intel® Ethernet controller has connection and it's bounded to a virtual adapter the connection test results states "TCP/IP protocol not configured" instead of "Cannot run this test because the device is bound to a virtual NIC".
- **Fixed:** In the Adapter Information tab, *Negotiated link speed* text is truncated. Only first word is shown. It's exists only when the German language is used.

2.5 Intel® Ethernet 200 Series

None for this release.

3.0 Known Issues

3.1 Intel® Ethernet 800 Series

3.1.1 General

- 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.
- 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/BLOPR2101MB093051C80B1625AAE3728551CA4A0@BLOPR2101MB0930.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.

- 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

- When using RDE operation status might return operation failed after a second update of a property.

3.1.3 Linux Driver

- 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.
- A minimal performance regression was observed on Intel network devices operating under Red Hat Enterprise Linux release version 8.2. This regression was not seen with versions 8.1 or 8.3.
- 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.
- 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.

3.1.4 Windows Driver

- 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 may 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).
- Under heavy stress on Windows Server 2019, customers might see Event 260: "Failed to move RSS queue from VMQ.." in the system event viewer. This is due to the driver not moving a particular queue in time for it to move to a new CPU. This is a non-fatal message and traffic is still processed as expected.

- On some platforms, an Intel® Ethernet E810 OCP-based device might not retain link when the platform is powered down. The device regains link when the platform is powered up. Sending a Magic Packet to such a device does not wake the system.
- 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.
- 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. This issue can be resolved by disabling and then enabling the Windows driver.

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.
- 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

- After a system reboot, an Intel® Ethernet Network Adapter E810 RDMA device in RoCEv2 mode might occasionally become active with a missing or incorrect GID. To correct the GID value, unload and reload the `irdma` driver.
- When using Intel MPI in Linux, Intel recommends to enable only one interface on the networking device to avoid MPI application connectivity issues or hangs. This issue affects all Intel MPI transports, including TCP and RDMA. To avoid the issue, use `ifdown <interface>` or `ip link set down <interface>` to disable all network interfaces on the adapter except for the one used for MPI. OpenMPI does not have this limitation.
- In RoCEv2 mode, running the `perftest ib_send_bw` benchmark application with high QP count, `rdma_cm`, and the duration option might occasionally result in a Destroy QP Cmd Error and a device reset.

- 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.

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 workaroud 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

- 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

- OID_GEN_RCV_ERROR doesn't increment when stream has not allowed preamble.
- NCSI version advertised during MCTP discovery has been changed.
- 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.
- Some odata.id properties are wrong.

3.2.1.1 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.

3.2.2 Windows Driver

- OID_GEN_RCV_ERROR doesn't increment when stream has not allowed preamble.

- `OID_GEN_RCV_ERROR` double counts some packets when traffic doesn't pass Cyclic Redundancy Check (CRC).

3.2.3 Linux Driver

- The `rx_errors` counter double counts some packets when traffic doesn't pass Cyclic Redundancy Check (CRC).
- On Kernel version 5.0.9 and higher, setting promiscuous mode on trusted VF leads to a periodic and endless update of this mode.
- When two VF interfaces are in balance-tlb bonding mode, the ICMP traffic might be interrupted for a few sequences.

3.2.4 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.
- VLAN status might be presented as Link Down in the **Teaming/VLANs** tab whereas the status is Enabled in the Windows Device Manager.

3.2.5 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 Intel® Ethernet 500 Series

3.3.1 General

- VXLAN switch creation on Windows Server 2019 Hyper V might fail.

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

- `OID_GEN_RCV_ERROR` doesn't increment when stream has not allowed preamble.
- `OID_GEN_RCV_ERROR` double counts some packets when traffic doesn't pass Cyclic Redundancy Check (CRC).

3.3.4 Linux Driver

- The `rx_errors` counter double counts some packets when traffic doesn't pass Cyclic Redundancy Check (CRC).

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 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	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/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.

© 2021 Intel Corporation.