



Intel[®] Ethernet Controller Products

29.0 Release Notes

February 2024

Revision 1.0

Revision History

Revision	Date	Comments
1.0	February 2024	<ul style="list-style-type: none">• Initial release.• Updated Linux drivers and Firmware/NVM versions.• Fixed issue for Linux* driver when creating new VLAN devices or setting the port VLAN.• Added Azure Stack HCI 23H2 OS Support in section 1.1.2.

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	5
1.3 Operating Systems Supported	6
1.3.1 Linux*	6
1.4 Windows Server	6
1.4.1 Windows Client	8
1.4.2 FreeBSD	9
1.4.3 ESXi Drivers	9
1.5 NVM Versions Supported	10
1.6 DDP Versions Supported	10
2.0 Fixed Issues	11
2.1 Intel® Ethernet 800 Series Network Adapters	11
2.1.1 Intel® Ethernet 810 Series	11
2.1.2 Intel® Ethernet 820 Series	12
2.2 Intel® Ethernet 700 Series Network Adapters	13
2.2.1 Linux Driver	13
2.2.2 Firmware/NVM/NVM Update	13
2.2.3 Windows Driver	14
3.0 Known Issues	15
3.1 Intel® Ethernet 800 Series Network Adapters	15
3.1.1 Intel® Ethernet 810 Series	15
3.1.2 Intel® Ethernet 820 Series	21
3.2 Intel® Ethernet 700 Series Network Adapters	24
3.2.1 General	24
3.2.2 Windows Driver	24
3.2.3 Intel® Ethernet Controller V710-AT2/X710-AT2/TM4	24
3.2.4 Linux Driver	24
3.2.5 VMware Driver	25
3.2.6 Firmware/NVM/NVM Update	25
3.3 Intel® Ethernet 500 Series Network Adapters	25
3.4 Legacy Devices	25
4.0 NVM Upgrade/Downgrade 800 Series/700 Series and X550	26
5.0 Languages Supported	26
6.0 Related Documents	26
6.1 Feature Support Matrix	26
6.2 Specification Updates	27
6.3 Software Download Package	27
6.4 SourceForge Ethernet Drivers and Utilities	27
6.5 Intel Product Security Center Advisories	27



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
29.0	<ul style="list-style-type: none">None for this release.

1.1.2 Software Features

Release	New Software Support
29.0	<ul style="list-style-type: none">Azure Stack HCI 23H2 OS Support.

1.1.3 Firmware Features

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

1.2 Removed Features

Release	Hardware/Feature Support
29.0	<ul style="list-style-type: none">None for this release.

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) 9.3, 9.2
- Red Hat Enterprise Linux 8.9, 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	1.13.7	4.10.6	1.13.43
Intel® Ethernet 700 Series	2.24.6	4.10.6	1.13.43
Intel® Ethernet 10 Gigabit Adapters	5.19.9	4.18.9	Not Supported
Intel® Ethernet Gigabit Adapters	5.15.7	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 Azure Stack HCI

Table 2. Supported Operating Systems: Windows Server

Driver	Windows Server 2022	Windows Server 2019	Windows Server 2016
Intel® Ethernet 800 Series			
icea	1.14.104.0	1.14.104.0	1.14.104.0
scea	1.13.261.0	1.13.261.0	Not Supported

Table 2. Supported Operating Systems: Windows Server [continued]

Driver	Windows Server 2022	Windows Server 2019	Windows Server 2016
Intel® Ethernet 700 Series			
i40ea	1.18.370.0	1.18.370.0	1.18.369.0
i40eb	1.18.370.0	1.18.370.0	1.18.369.0
Intel® Ethernet Adaptive Virtual Function			
iavf	1.14.202.0	1.14.202.0	1.14.202.0
v40e	Not Supported	Not Supported	Not Supported
Intel® Ethernet 10 Gigabit Adapters and Connections			
ixs	4.1.252.0	4.1.251.0	4.1.246.0
sxa	4.1.252.0	4.1.251.0	4.1.249.0
sxb	4.1.252.0	4.1.251.0	4.1.249.0
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	2.1.4.2	1.1.4.42	Not Supported
Intel® Ethernet Gigabit Adapters and Connections			
e1r	14.0.2.0	14.0.2.0	12.16.7.0
v1q	Not Supported	1.4.7.3	1.4.7.3

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	1.14.104.0	1.14.104.0	Not Supported
Intel® Ethernet 700 Series			
i40ea	1.18.370.0	1.18.370.0	Not Supported
Intel® Ethernet Adaptive Virtual Function			
iavf	1.14.202.0	1.14.202.0	1.14.202.0
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	2.1.4.2	1.1.4.42	Not Supported
Intel® Ethernet Gigabit Adapters and Connections			
e1r	14.0.2.0	14.0.2.0	12.16.7.0
e1d	12.19.2.56	21H2: 12.19.2.56 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 14.0
- FreeBSD 12.4

Table 4. Supported Operating Systems: FreeBSD

Driver	PF Driver	VF Driver	RDMA Driver
Intel® Ethernet 810/820 Series	1.39.13	3.0.33	1.2.36
Intel® Ethernet 700 Series	1.13.5	3.0.33	1.2.36
Intel® Ethernet 10 Gigabit Adapters	3.3.36	1.5.37	Not Supported
Intel® Ethernet Gigabit Adapters	2.5.30	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	4.40
820 Series	
E822	2.28
E823-C	2.28
E823-L	2.28
700 Series	
X710	9.40
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

- Previously, upon creating new VLAN devices or setting the port VLAN for the VF device, and the irdma module is also loaded, a system hang might occur. Try to unload the irdma module ("modprobe -r irdma").
- Previously, 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) could result in starvation of the node/leaf (0 BW on Tx), with a lower priority set, and trigger a Tx hang detection.
- Resolved an issue when DCB configuration was 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 reported 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 would repeat forever.

- Previously, with the 810 Series 3.2 NVM in the Intel® Ethernet Network Adapter E810-CQDA2 card, if the 810 Series 2.2 *javf* driver was installed, a fatal error was generated related to **pci-aspm.h**, and the installation failed.
- Resolved an issue where 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 would be concentrated on the first queue of interface.
- Fixed a firmware issue that could result in a firmware reset when some, but not all, ports have flow control enabled.
- Previously, there were unexpected FW logs during some special environments, such as tx_timeout or NVM update.
- Improved driver's robustness when bringing interface up/down with link-down-on-close flag is set and recoverable non-critical error occurs.
- Corrected an issue where upon creating new VLAN devices or setting the port VLAN for the VF device, and the irdma module is also loaded, a system hang might occur.

2.1.1.2 Firmware/NVM/NVM Update

- Previously, the PLDM Type 5 Command "Activate Firmware" could potentially cause a subsequent PLDM Type 2 event.
- Resolved an issue where during an ActiveFirmware request coming from Update Agent (BMC) PLDM Type2, a PlatformEventMessage was generated by the Firmware Device (Network Adapter). As a result, the response did not happen, and Update Agent (BMC) treated it as Update Failed.
- Fixed an issue where the UEFI HII interface could not save the setting of virtualization mode.

- Previously, in PLDM Platform Monitoring and Control, the `sensor_data_size` field of SFP thermal sensor PDR was containing wrong value: unsigned instead of signed char/byte.
- Previously, during ActiveFirmware request coming from Update Agent (BMC) PLDM Type2 `PlatformEventMessage` was generated by our Firmware Device (Network Adapter). As a result, response did not happen and Update Agent (BMC) treated it as Update Failed.
- Firmware now acquires semaphore before "Load Factory Settings to PFA" and releases it immediately after.
- Firmware now returns `PLDM_ERROR_NOT_READY` Completion Code instead of generic `0x01`(PLDM Base: Error) when firmware is not fully ready after platform reboot.
- Previously, when BMC subscribed to PLDM events, the NIC issued two equivalent redfish events for the different origins (Network Adapter and Port) caused by the same trigger (link state change for given port). This issue caused unnecessary redfish event flood. The resolution was to change firmware design in order to avoid sending multiple redfish events for the same trigger.
- Previously, in Network Controller Sideband Interface after `EnableChannel` command invoked, the OEM command `IANA=0x02A2 GetLldp` occasionally returned LLDP TLVs with missing mandatory values for `chassis_id` and `port_id`.
- Set `packageID` SDP to correct values for E810-CQDA1.
- Now the patched property of `NetworkDeviceFunction.Ethernet.MACAddress` is persistent during reboot. Users can restore to `PermanentMACAddress` with `ResetSettingsToDefaults`.

2.1.1.3 Application Device Queues (ADQ)

- Previously, ADQ dynamically registered and unregistered 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 were to be expected when ADQ queue groups were created or destroyed for the first time.

2.1.2 Intel® Ethernet 820 Series

2.1.2.1 Firmware/NVM/NVM Update

- Resolved an issue where promiscuous mode did not see all packets; it saw only those packets arriving over the wire; that is, not sent from the same physical function (PF) but a different virtual function (VF).
- Previously, the DCB-MAP Configuration was not displaying on the SUT Interface from the Extreme Switch after enabling the firmware Mode in the SUT.
- Previously, using the EPCT tool to change port configuration required 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.
- Resolved an issue where the **Get LLDP** command (0x28) response used to contain only 2 TLV types. The new implementation requires a third TLV type, which contains a copy of the whole LLDP frame payload.
- Resolved an issue where the BMC was unable to receive packets when DHCPv6 server to client multicast filter is enabled.

2.1.2.2 Linux Driver

- Previously, it had been observed in certain scenarios where inner VLAN traffic would concentrate on a single queue, creating packet drops.
- Resolved an issue where drop action was not properly offloaded in hardware when setting was done using Open vSwitch.

2.1.2.3 Windows Driver

- Previously, with a Windows host and Linux virtual machine (VM), the last transmit (Tx) queue might not increment when there are multiple Tx queues.
- Fixed an issue when using hardware timestamp an incorrect behavior was observed on the receive (Rx) side. Timestamp was received and no error messages were reported, however its value was 0 and therefore was considered as faulty.

2.1.2.4 FreeBSD Driver

- Previously, when a driver was loaded with an empty cage, an Admin Queue (AQ) error was recorded instead of the expected AHS link messages.

2.1.2.5 LANConf Tool

- Resolved a limitation of the LANConf tool that only allowed the user to program ID EEPROM on Quad 0. This also meant that each attempt to program ID EEPROM on Quad 1 would cause an update on Quad 0.

2.2 Intel® Ethernet 700 Series Network Adapters

2.2.1 Linux Driver

- Resolved an issue where the *DMARemappingCompatible* registry was missing in the previous driver.
- Resolved an issue that failed to detect temperature status with latest OOB driver.

2.2.2 Firmware/NVM/NVM Update

- Previously, for *XXV710DA2_9p30_CFGID12p0_OEMGEN_J65979* the configuration files did not contain mention of *0x8000DECD*, which is the original factory program image.
- Resolved an issue where POST to the *NetworkAdapter.ResetSettingsToDefault* looked successful from postman, but after reboot the adapter's Port 1 LLDP setting did not reset to default.
- Previously, there were cases of Fortville disappearing from the PCIe bus during reboot stress testing reported on Icelake-based platforms. These platforms performed a PCIe Hot Reset before asserting PERST# as part of the warm reboot cycle. PERST# is asserted at approximately the end of the Hot Reset, but due mostly to the Fortville LCB implementation, there was a variation in the relative timing of about 65 us, so many edge cases could occur.
- Fixed a wrong demsg message from **i40e** driver of unsupported SFP+ module type was detected when eeupdate upgrade NVM of x710t21 or x710t41.

2.2.3 Windows Driver

- Previously, setting PHY debug mode with disabling link during initialization caused link to go down.
- Previously, during the NVM update the adapter was getting disabled for a long time due to lock in `OID_INTEL_FLASH_INFO`. As result, the driver was unable to access the NVM timed out quickly.
- Resolved an issue where a rare PCI speed degradation with Gen3 to Gen2 had been seen during extensive link tests after a reset was caused by a core reset on an ICX-D platform.
- NDIS i40eb (X722 NIC) component now supports Japanese in advanced tab language.
- Resolved an issue in Windows Server 2022 where the E1R driver didn't disable auto-negotiation when choosing a specific speed and duplex setting.

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.
- 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.
- When using hardware timestamp incorrect behavior is observed on the receive (Rx) side. The timestamp is received and no error messages are reported, however its value is 0 and therefore is considered as faulty. The transmit (Tx) side works fine.

3.1.1.2 Firmware/NVM/NVM Update

- The second part of the string from the option ROM indicates the correct slot (port) numbers.
- Host overwrites MTU configuration on startup with AQ commands.
- When trying to read Redfish resource of the Port schema, the resource data may contain field `AutoSpeedNegotiationEnable` set to `false` and the `LinkConfiguration.ConfiguredNetworkLink.ConfiguredLinkSpeedGbps` field set to multiple speeds. Those two values are conflicting in this case. **Workaround:** Ignore those values until issue resolved in future releases.
- When Firmware is operating in mode when LLDP is ON, the DCB-MAP is not reflecting as configured in both switch and back to back. The **workaround** for this behavior is to do the power cycle of the setup to see the assigned DCB-MAP is reflecting.
- PLDM Type 4 may return incorrect values for Part Number, Serial Number, Manufacturing Date/Time and SKU.

3.1.1.3 Linux

- 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 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.

- 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.
- 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 `iavf` 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/devink** which might result in errors.
 - **Workaround:** Update to the latest **devlink** version.
- 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.
- Attempting to unload ice driver with LLDP agent present may cause ice driver to crash, causing a kernel panic due to an issue with LLDP agent. The **workaround** for this is to disable LLDP prior to unloading ice driver.
- When plugging DAC cables into two ports at the same time on Clifton Channel, it is possible for one port to be unable to establish link. Workaround is to connect cables one-by-one with 5 second delay between each connection.
- The Intel® Ethernet Network Adapter E810-2C-Q2 may not be detected by Dell R750 systems.
- In order to use more queues for VF, the user is obligated to properly configure Global LUT by `echo 512 > /sys/class/net/$pf/device/virtfn0/rss_lut_vf_attr` or alternatively, configure VF LUT with

small amount of MSI-x and also properly configure distribution of MSI-x vectors by using `echo 16 > /sys/class/net/$pf/device/virtfn0/sriov_vf_msix_count`

- Starting a Linux VM can result in a Linux call trace, as communication is broken from the VF side. With Windows Host OS and Linux Guest OS (RHEL 8.8, 8.7 & 8.1) this behavior is not observed.
- The latest ICE driver limits the number of interrupts to the local core count, regardless of the number of queues created. As a result, certain scripts (`adqsetup.py`, `set_irq_affinity`, etc) may fail and errors are observed during interrupt affiliation.

Workaround: After a fresh driver load, run the following commands:

```
IFACE=<interface_name>

BDF=$(awk -F=' ' '($1 == "PCI_SLOT_NAME") {print $2}' /sys/class/net/$IFACE/device/uevent)

devlink resource set pci/$BDF path msix/msix_vf size 0
devlink resource set pci/$BDF path msix/msix_eth size 256
devlink dev reload pci/$BDF action driver_reinit
```

- Customer setups may not allow creation of 64 VF due to resource utilization, even on high-core systems (greater than 400 cores). When visible, configure 62. The problem is not visible on ~100 core systems.

We observe an issue with the core dump related with the reset done on VF. The problem identified with CrashDump is observed on 4.10.x or 4.9.5 IAVF drivers with the combination of 1.13.7 ICE driver, but is not visible on the 4.9.1 driver. For the High VF system, when the issue is visible, it is recommended to use 4.9.1, and use it until 4.10.x is released.

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. 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)

- When both remote and local IP addresses and ports are specified, the ADQ setup script uses an AND for local ports and addresses, an AND for remote ports and addresses, and an OR for the result of those two ANDs. As a result, filters with IP and port number are created separately (one with IP and one with port number), and traffic is expected to arrive on both queues with the current behavior. 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 *lavf* 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.
- The use of ntuple filters is incompatible with the in-line flow director and may result in traffic reverting to RSS selection of RX queues - even after the ntuple filters are deleted (by turning off the 'ntuple-filters' feature flag). If the user suspects this is the case, the system must be rebooted to restore proper operation of the in-line flow director. To avoid this issue when there are already ntuple filters on the system, delete the ntuple filters individually by ID before attempting to use the in-line flow director.

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.
- **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

- The 100 MB option, is visible in Windows* Device Manager. However, when it is selected, a link cannot be established.
- There is a limitation that the NVM update tool doesn't allow the user to program PHY FW on Quad 1.
- 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.

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.
- 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.
- 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
```

- 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.

3.1.2.4 FreeBSD Driver

- 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.

- When using virtualization, the adapter is repeatedly reset and SRIOV may become degraded at some point. After VM reboot all adapters are working correctly. This does not have an impact on Hyper-V as it is continuously working.

Based on this, we recommend to avoid stress testing the PF adapter reset (over 1000 times in a row).

- When user configures Hyper-V to use Encapsulated Task Offload with or without NVGRE, there is a chance the service will become unresponsive and crash. The system will recover in about 3 minutes with the feature enabled. Another workaround is that Hyper-V service can be restarted manually. This behavior is observed sporadically.

3.1.2.6 VMware Driver

- In combination with some specific VF-PF drivers timestamp may not work at all on Intel® Ethernet 820 Series devices.

3.2 Intel® Ethernet 700 Series Network Adapters

3.2.1 General

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.
- In some instances, a blue screen was observed while installing the driver with *DmaRemappingCompatible* set to 1, enabled Driver Verifier and RDMA.
- Dell SKU adapters may show 100 Mbps speed as an option, but this is incorrect.

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 to 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 PTP4I 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.
- NC-SI Intel OEM command can't work properly for X710 OCP adapters.
- In NUT code there is `fcntl` call to lock NVMupdate PID so two instances will be not run at the same time. It is done by opening `/var/run/nvmupdate.pid` with `O_CREAT|O_WRONLY` and `0600` mode at the begging of tool work. Then the `fcntl` call is used with lock type `F_WRLCK`. All customers need to add feature to its program to detect if NVMupdate is running by checking the `/var/run/nvmupdate.pid` and stop or decrease number of `ioctl`s sent during update.

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: <ul style="list-style-type: none"> – 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: <ul style="list-style-type: none"> – 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>



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

© Intel Corporation. Intel, the Intel logo, and other Intel marks are trademarks of Intel Corporation or its subsidiaries. Other names and brands may be claimed as the property of others.