



# **Intel Ethernet Controller Products**

**Release 26.1 Release Notes**

---

**Intel Corporation**

Revision 1.0  
635017-001



## Revision History

---

Revision	Date	Comments
1.0	March 2021	<ul style="list-style-type: none"><li data-bbox="576 413 738 434">• Initial release.</li></ul>

## 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 26.1.

Controller Family	Windows* Driver Version	Linux* Driver Version	FreeBSD Driver Version	NVM Version
<b>Intel® Ethernet 800 Series</b>	<b>1.7.29.0</b>	<b>1.4.11</b>	<b>0.28.6</b>	<b>2.4x</b>
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 Series <sup>1</sup>	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
<ul style="list-style-type: none"> <li>26.1</li> </ul>	<ul style="list-style-type: none"> <li>Linux and FreeBSD Support for the following devices:                             <ul style="list-style-type: none"> <li>Intel® Ethernet Network Adapter E810-2C-Q2</li> </ul> </li> </ul>

### 1.2.2 Software Features

Release	New Software Support
<ul style="list-style-type: none"> <li>• 26.1</li> </ul>	<ul style="list-style-type: none"> <li>• QinQ support</li> <li>• NVM Update using <b>devlink</b> <ul style="list-style-type: none"> <li>— Kernel versions 5.10 and later are needed for full devlink support using NVMLUpdate</li> </ul> </li> <li>• FreeBSD RDMA driver for Intel® Ethernet Controller 800 Series devices.</li> <li>• The Microsoft Windows download package has been split into a driver package and a Intel® PROSet package.</li> </ul> <p><b>Note:</b> You can still download the entire release all in one package.</p>

### 1.2.3 Removed Features

Release	Hardware/Feature Support
<ul style="list-style-type: none"> <li>• 26.1</li> </ul>	No features have been removed from this release.

### 1.2.4 Firmware Features

Release	New Firmware Support
<ul style="list-style-type: none"> <li>• 26.1</li> </ul>	<ul style="list-style-type: none"> <li>• No new firmware features introduced in this release.</li> </ul>

## 2.0 Fixed Issues

### 2.1 Intel® Ethernet 800 Series

#### 2.1.1 Linux Driver

- When the XDP program is loaded in native mode, a system crash might occur.
- On Intel Ethernet E810-2C-Q2, you may see errors when there are more than 128 queues per traffic class.
- The `rx_gro_dropped` statistics was removed and is no longer accounted for in `rx_dropped` netstats.
- Until the system is rebooted, you may see a loss of traffic following an NVM Update. Rebooting the system resumes traffic.
- The Intel ice driver interrupt moderation mechanism now utilizes the kernel DIMLIB library. This change has shown improvement across many workloads including large packet latency and UDP streams. Note that the ice driver from Release 25.6 (and earlier) does not have this change.
- Unified Ethtool Statistics was added to Linux driver starting from Release 25.5 as described in <https://sourceforge.net/p/e1000/wiki/UnifiedStats/>.

#### 2.1.2 Windows Driver

- 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.
- 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.

#### 2.1.3 Linux RDMA Driver

- 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.

#### 2.1.4 NVM Update Tool

- In ESXi, the device serial number appears as N/A.

#### 2.1.5 NVM

- A workaround was added to prevent Errata "Device Fails to Initialize PCIe Link After Reboot". See specification update document for more details <https://cdrdv2.intel.com/v1/dl/getContent/616943>

#### 2.1.6 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).

### **2.1.7 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.
- Downgrades to earlier releases using PLDM FW update method maybe be blocked.

### **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

- In 8x10G mode link may not be established on some ports if the link is not connected at power on.
- E810 currently supports a subset of 1000BASE-T SFP module types, which use SGMII to connect back to the E810. In order for the E810 to properly know the link status of the module's BASE-T external connection, the module must indicate the BASE-T side link status to the E810. An SGMII link between E810 and the 1000BASE-T SFP module allows the module to indicate its link status to the E810 using SGMII Auto Negotiation. However 1000BASE-T SFP modules implement this in a wide variety of ways, and other methods which do not use SGMII are currently unsupported in E810. Depending on the implementation, link may never be achieved. In other cases, if the module sends IDLEs to the E810 when there is no BASE-T link, the E810 may interpret this as a link partner sending valid data and may show link as being up even though it is only connected to the module and there is no link on the module's BASE-T external connection.
- 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/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.

- 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

- If a Virtual Function (VF) is not in trusted mode and eight or more VLANs are created on one VF, the VLAN that is last created might be non-functional and an error might be seen in `dmesg`.
- 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.
- When using a Windows Server 2019 RS5 Virtual Machine on a RHEL host, a VLAN configured on the VF using `iproute2` might not pass traffic correctly when an `ice` driver older than version 1.3.1 is used in combination with a newer AVF driver version.
- The user might see intermittent boot failures when booting from certain iSCSI targets.
- Unloading the ICE driver after a PCI reset causes an error in `dmesg`.
- When the Double VLAN Mode is enabled on the host, disabling then re-enabling a Virtual Function attached to a Windows guest might cause error messages to be displayed in `dmesg`. These messages will not affect functionality.
- With the current `ice` PF driver, there might not be a way for a trusted VF to enable unicast promiscuous and multicast promiscuous mode without turning on `ethtool --priv-flags with vf-true-promisc-support`. As such, the expectation is to not use `vf-true-promisc-support` to gate VF's request for unicast/multicast promiscuous mode.
- On an Intel Ethernet E810 2CQDA2 adapter in single VLAN mode, users might see packets with S-tag being dropped. This isn't seen with double VLAN.
- Repeatedly assigning a VF interface to a network namespace then deleting that namespace might result in an unexpected error message and might possibly result in a call trace on the host system.
- Receive hashing might not be enabled by default on Virtual Functions when using an older `iavf` driver in combination with a newer PF driver version.
- When Double VLAN is created on a Virtual Machine, `tx_tcp_cso` [TX TCP Checksum Offload] and `tx_udp_cso` [TX UDP Checksum Offload] statistics might not increment correctly.
- When changing the RSS queue value in VM while Double VLAN is configured, `rx_vlano` [rx-vlan-offload] might stop incrementing correctly.
- When using Double VLAN configuration with a specific non-Intel link partner, TCP traffic might fail to pass through the inner VLAN interface when the MTU of this interface matches the MTU of the outer VLAN interface. **Workaround:** change MTU of inner VLAN to be 4 bytes less than the MTU of the outer VLAN.



- 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.
- If setting the device MTU to either 4096 or 2048, applications using RoCEv2 UD might hang. **Workaround:** Set the MTU values to 4097 instead of 4096 and 2049 instead of 2048.
- A PCI reset performed on the host might result in traffic failure on VFs for certain guest operating systems.
- On RHEL 7.x and 8.x operating systems, it has been observed that the `rx_gro_dropped` statistic might increment rapidly when Rx traffic is high. This appears to be an issue with the RHEL kernels.
- When ICE interfaces are part of a bond with `arp_validate=1`, the backup port link status flaps between up and down. **Workaround:** It is recommended to not enable `arp_validate` when bonding ICE interfaces.
- Changing a Virtual Function (VF) MAC address when a VF driver is loaded on the host side might result in packet loss or a failure to pass traffic. As a result, the VF driver might need to be restarted.
- 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

- When a VM is running heavy traffic loads and is attached to a Virtual Switch with either SR-IOV enabled or VMQ offload enabled, repeatedly enabling and disabling the SR-IOV/VMQ setting on the vNIC in the VM may result in a BSOD.
- When updating from Release 25.6 to Release 26.1 using LANConf, the error "Shared flash write failed!" might be displayed at the end of the update. This does not indicate an update failure. Reboot the system to load the latest firmware.
- When a SR-IOV enabled Virtual Switch is assigned to 8 VMs and bi-directional traffic is ran continuously, it is possible that SR-IOV degradation might occur.
- Small performance limitations have been seen for some workloads on Windows when stressing both ports of the E810-2CQDA2 adapter.
- 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.
- 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

- .
- 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

- RDMA vLAN may not be supported in the current release.
- 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 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.
- 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.

## 3.2 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 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	<a href="https://cdrdv2.intel.com/v1/dl/getContent/630155">https://cdrdv2.intel.com/v1/dl/getContent/630155</a>
Intel® Ethernet 700 Series: – X710/XXV710/XL710 – X722 – X710-TM4/AT2 and V710-AT2	<a href="https://cdrdv2.intel.com/v1/dl/getContent/332191">https://cdrdv2.intel.com/v1/dl/getContent/332191</a> <a href="https://cdrdv2.intel.com/v1/dl/getContent/336882">https://cdrdv2.intel.com/v1/dl/getContent/336882</a> <a href="https://cdrdv2.intel.com/v1/dl/getContent/619407">https://cdrdv2.intel.com/v1/dl/getContent/619407</a>
Intel® Ethernet 500 Series	<a href="https://cdrdv2.intel.com/v1/dl/getContent/335253">https://cdrdv2.intel.com/v1/dl/getContent/335253</a>
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	<a href="https://cdrdv2.intel.com/v1/dl/getContent/616943">https://cdrdv2.intel.com/v1/dl/getContent/616943</a>
Intel® Ethernet 700 Series: – X710/XXV710/XL710 – X710-TM4/AT2 and V710-AT2	<a href="https://cdrdv2.intel.com/v1/dl/getContent/331430">https://cdrdv2.intel.com/v1/dl/getContent/331430</a> <a href="https://cdrdv2.intel.com/v1/dl/getContent/615119">https://cdrdv2.intel.com/v1/dl/getContent/615119</a>
Intel® Ethernet 500 Series – X550 – X540	<a href="https://cdrdv2.intel.com/v1/dl/getContent/333717">https://cdrdv2.intel.com/v1/dl/getContent/333717</a> <a href="https://cdrdv2.intel.com/v1/dl/getContent/515715">https://cdrdv2.intel.com/v1/dl/getContent/515715</a>
Intel® Ethernet 300 Series	<a href="https://cdrdv2.intel.com/v1/dl/getContent/333066">https://cdrdv2.intel.com/v1/dl/getContent/333066</a>
Intel® Ethernet 200 Series – I210 – I211	<a href="https://cdrdv2.intel.com/v1/dl/getContent/332763">https://cdrdv2.intel.com/v1/dl/getContent/332763</a> <a href="https://cdrdv2.intel.com/v1/dl/getContent/333015">https://cdrdv2.intel.com/v1/dl/getContent/333015</a>

## 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](http://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.