



Intel[®] Ethernet Controller Products

26.6 Release Notes

Ethernet Products Group

October 2021

Revision 1.0
648240-001



Revision History

Revision	Date	Comments
1.0	October 2021	Initial release.

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: Release 26.5 has been canceled. Release 26.6 is the first public release after Release 26.4.

Note: **Bold Text** indicates the main changes for Software Release 26.6.

Controller Family	Windows* Driver Version	Linux* Driver Version	FreeBSD Driver Version	NVM Version
Intel® Ethernet 800 Series	1.9.65.x,	1.6.x	1.30.x	3.0x
Intel® Ethernet 700 Series Intel® Ethernet Network Connection X722	1.16.62.x,	2.17.x	1.12.x	8.50
Intel® Ethernet Adaptive Virtual Function (Intel® Ethernet AVF)	1.12.9.x,	4.2.x	3.0.x	N/A
Intel® Ethernet 500 Series Intel® Ethernet Controller X550	4.1.239.x,	5.13.x	3.3.x	3.50
Intel® Ethernet 300 Series	12.18.12.x,	5.8.x	2.5.x	N/A
Intel® Ethernet 200 Series	12.19.1.x	5.8.x	2.5.x	2.00
Intel® Ethernet 225 Series	1.02.x	N/A	N/A	N/A
Intel® Ethernet Controller 82598	2.4.36.x	5.10.x	3.3.x	N/A
Intel® Ethernet Controller 8257x	9.16.10.x	5.5.x	2.5.x	N/A

1.2 New Features

1.2.1 Hardware Support

Release	New Hardware Support
• 26.6	• None for this release.

1.2.2 Software Features

Release	New Software Support
• 26.6	<ul style="list-style-type: none"> • Support for Microsoft* Windows* 10, version 21H2 • Support for SUSE* Linux Enterprise Server (SLES) 15 SP3 • Support for Intel® PROSet for Windows PowerShell* software on Microsoft Azure Stack HCI

1.2.3 Removed Features

Release	Hardware/Feature Support
<ul style="list-style-type: none"> • 26.6 	<ul style="list-style-type: none"> • SUSE Linux Enterprise Server 11 is no longer supported. • Canonical* Ubuntu* 16.04 is no longer supported. • Microsoft Windows e1q driver is no longer supported. The drivers may still be provided in your download package or on your install media for your convenience. This affects the following devices: <ul style="list-style-type: none"> – Intel® 82574L Gigabit Network Connection – Intel® Gigabit ET2 Quad Port Server Adapter – Intel® 82575EB Gigabit Backplane Connection – Intel® 82575EB Gigabit Network Connection – Intel® 82576 Gigabit Dual Port Network Connection – Intel® 82576 Gigabit Dual Port Server Network Connection – Intel® 82576NS Gigabit Ethernet Controller – Intel® 82576NS SerDes Gigabit Ethernet Controller – Intel® 82583V Gigabit Network Connection – Intel® Gigabit CT Desktop Adapter – Intel® Gigabit CT2 Desktop Adapter – Intel® Gigabit EF Dual Port Server Adapter – Intel® Gigabit ET Dual Port Network Connection – Intel® Gigabit ET Dual Port Server Adapter – Intel® Gigabit ET Quad Port Mezzanine Card – Intel® Gigabit ET Quad Port Server Adapter – Intel® Gigabit VT Quad Port Server Adapter

1.2.4 Firmware Features

Release	New Firmware Support
<ul style="list-style-type: none"> • 26.6 	<ul style="list-style-type: none"> • None for this release.

2.0 Fixed Issues

2.1 Intel® Ethernet 800 Series

2.1.1 General

2.1.2 Linux Driver

None for this release.

2.1.3 Windows Driver

None for this release.

2.1.4 Linux RDMA Driver

None for this release.

2.1.5 NVM Update Tool

None for this release.

2.1.6 NVM

None for this release.

2.1.6.1 Firmware

None for this release.

2.1.6.2 Manageability

None for this release.

2.1.7 FreeBSD Driver

None for this release.

2.1.8 Application Device Queues (ADQ)

None for this release.

2.2 Intel® Ethernet 700 Series

2.2.1 General

None for this release.

2.2.2 Linux driver:

- The Virtual Function MAC address is incorrect after removing bonding module in the balance tlb/alb modes. The Virtual Function that is added to the bonding tlb/alb modes as an inactive slave will have an incorrect MAC address after the currently active slave is down and bonding module is removed.

2.2.3 Intel® PROSet:

None for this release.

2.2.4 EFI Driver

None for this release.

2.2.5 NVM

- The XL710 cannot receive packets from the switch, causing a bond link down condition.

2.2.6 Windows driver:

None for this release.

2.2.7 Intel® Ethernet Flash Firmware Utility:

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

- The version of the Ethernet Port Configuration Tool available in Release 26.1 may not be working as expected. Please use the version of the tool from the latest release.
- On Intel Ethernet 800 Series adapters, DDPtool may not be able to display the DDP profile.
- 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 virtual functions (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 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 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

- For PLDM Type 2 the device reports the port thermal sensor PDR and status for all types of connectivity; even those without temp sensor (temperature and thresholds reported as zeros).
- BMC may not be able to get the ASIC temperature through NC-SI OEM command.
- Flow control settings have no effect on traffic, and counters do not increment with flow control set to TX=ON and Rx=OFF. However, flow control works fine with values set to TX=On RX=ON.

3.1.3 Linux Driver

- 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 may 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.
- On CentOS 7.2, ethtool may report Speed as Unknown and you may see warnings in dmesg log. This does not affect traffic or device functionality.
- QoS bandwidth shaping and priority tagging may not be functional in CentOS 7.2.
- For configurations in which a VF is attached to a VM, avoid setting the MTU value of the VF within the range of 9199-9202. When the MTU is set to one of these values the VM may experience a memory leak leading to kernel panic when TCP traffic is running from a link partner to the VF.
- Rapid unloading and loading of the irdma and ice drivers may cause a kernel panic.
- Interrupt Moderation settings reset to default when the queue settings of a port are modified using the "ethtool -L ethx combined XX" command.
- 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 double VLAN or Queue in Queue feature is enabled, the inner most traffic source might need to be limited to an MTU size of 1496 or less to avoid connection issues.
- 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.
- 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.
- It has been observed that when using ISCSI, the ISCSI initiator intermittently fails to connect to the ISCSI target.

- 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.
- 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 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.
- 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.
- 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.
- Some distros may contain an older version of `iproute2/devink` which may result in errors. **Workaround:** Please update to the latest `devlink` version.

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 might not initially be visible at the selection screen.
- 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.
- Small performance limitations have been seen for some workloads on Windows when stressing both ports of the E810-2CQDA2 adapter.
- 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.

3.1.5 Windows RDMA Driver

- The Intel® Ethernet Network Adapter E810 might experience an adapter-wide reset on all ports. When in firmware managed mode, a DCBx willing mode configuration change that is propagated from the switch removes a TC that was enabled by RDMA. This typically occurs when removing a TC associated with UP0 because it is the default UP on which RDMA based its configuration. The reset results in a temporary loss in connectivity as the adapter re-initializes.
- With a S2D storage cluster configuration running Windows Server 2019, high storage bandwidth tests might result in a crash for a BSOD bug check code 1E (KMODE_EXCEPTION_NOT_HANDLED) with `smbdirect` as the failed module. Customers should contact Microsoft via the appropriate support channel for a solution.

3.1.6 Linux RDMA Driver

- When adding the ports to a VM in passthrough mode, all ports on the device must be added to the same VM with the same instance numbers.
- 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 configuration file is not supported.
- After downgrading to Release 25.6 (and previous), a loss of traffic may result. Workaround: Unload and reload the driver to resume traffic. Rebooting the system would also help.

3.1.8 Application Device Queues (ADQ)

The code contains the following known issues:

- The latest RHEL and SLES distros have kernels with back-ported support for ADQ. For all other OS distros, you must use the LTS Linux kernel v4.19.58 or higher to use ADQ. The latest out-of-tree driver is required for ADQ on all Operating Systems.
- If the application stalls for some reason, this can cause a queue stall for application-specific queues for up to two seconds.
 - Workaround - Recommend configuration of only one application per Traffic Class (TC) channel.
- DCB and ADQ are mutually exclusive and cannot coexist. A switch with DCB enabled might remove the ADQ configuration from the device.
 - Workaround - Do not enable DCB on the switch ports being used for ADQ. Disable LLDP on the interface by turning off firmware LLDP agent using:

```
ethtool --set-priv-flags $iface fw-lldp-agent off
```
- Note (unrelated to Intel drivers): The 5.8.0 Linux kernel introduced a bug that broke the interrupt affinity setting mechanism.
 - Workaround - Use an earlier or later version of the kernel to avoid this error.
- The iavf driver must use Trusted mode with ADQ: Trusted mode must be enabled for ADQ inside a VF. If TC filters are created on a VF interface with trusted mode off, the filters are added to the software table but are not offloaded to the hardware.
- VF supports Max Transmit Rate only: the iavf driver only supports setting maximum transmit rate (max_rate) for Tx traffic. Minimum transmit rate (min_rate) setting is not supported with a VF.
- A core-level reset of an ADQ-configured VF port (rare events usually triggered by other failures in the NIC/iavf driver) results in loss of ADQ configuration. To recover, reapply ADQ configuration to the VF interface.
- VF errors occur when deleting TCs or unloading iavf driver in a VF: ice and iavf driver error messages might get triggered in a VF when TCs are configured, and TCs are either manually deleted or the iavf driver is unloaded. Reloading the ice driver recovers the driver states.
- Commands such as `tc qdisc add` and `ethtool -L` cause the driver to close the associated RDMA interface and reopen it. This disrupts RDMA traffic for 3-5 seconds until the RDMA interface is available again for traffic.
- When the number of queues is increased using `ethtool -L`, the new queues will have the same interrupt moderation settings as queue 0 (i.e., Tx queue 0 for new Tx queues and Rx queue 0 for new Rx queues). This can be changed using the `ethtool` per-queue coalesce commands
- To fully release hardware resources and have all supported filter type combinations available, the ice driver must be unloaded and re-loaded.
- VXLAN stateless offloads (checksum, TSO) and TC filters directing traffic to a VXLAN interface are not supported with Linux v5.9 or later.

- When ADQ is enabled on VFs, TC filters on the VF TCO (default TC) are not supported and will not pass traffic. It is not expected to add TC filters to TCO since it is reserved for non-filtered default traffic.

3.1.9 Manageability

- Intel updated the E810 FW to align the sensor ID design as defined by DMTF DSP2054 starting from Release 26.4. Previous versions of the E810 FW were based on draft version of the specification. As a result updating to the newer NVM with this FW will result in updating numbering for the thermal sensorsIDs and PDR handlers. Anyone using hard coded values for these will see changes. A proper description of the system through PLDM type 2 PDRs shall give a BMC enough information to understand what sensors are available, what they are monitoring and what their ID is.

3.2 Intel® Ethernet 700 Series

3.2.1 General

- Devices based on the Intel® Ethernet Controller XL710 (4x10 GbE, 1x40 GbE, 2x40 GbE) have an expected total throughput for the entire device of 40 Gb/s in each direction.
- The first port of Intel® Ethernet Controller 700 Series-based adapters display the correct branding string. All other ports on the same device display a generic branding string.
- In order for an Intel® Ethernet Controller 700 Series-based adapter to reach its full potential, 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.

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

- Incorrect *DeviceProviderName* is returned when using RDE *NegotiateRedfishParameters*. This issue has been root caused and the fix should be integrated in the next firmware release.

3.2.3 Windows Driver

None for this release.

3.2.4 Linux Driver

- On Kernel version 5.0.9 and higher, setting promiscuous mode on trusted VF leads to a periodic and endless update of this mode. It can be observed via `dmesg`.

3.2.5 Intel® PROSet

None for this release.

3.2.6 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.

3.2.7 NVM

- None for this release.

3.3 Intel® Ethernet 500 Series

3.3.1 General

- None for this release.

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

- None for this release.

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 [here](#).

4.0 NVM Upgrade/Downgrade 800 Series/700 Series and X550

Refer to the Feature Support Matrix (FSM) links listed in [Related Documents](#) for more detail. FSMs list the exact feature support provided by the NVM and software device drivers for a given release.

5.0 Languages Supported

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

This release supports the languages listed in the table that follows:

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

6.0 Related Documents

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

6.1 Feature Support Matrix

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

Device Series	Support Link
Intel® Ethernet 800 Series	https://cdrdv2.intel.com/v1/dl/getContent/630155
Intel® Ethernet 700 Series: – X710/XXV710/XL710 – X722 – X710-TM4/AT2 and V710-AT2	https://cdrdv2.intel.com/v1/dl/getContent/332191 https://cdrdv2.intel.com/v1/dl/getContent/336882 https://cdrdv2.intel.com/v1/dl/getContent/619407
Intel® Ethernet 500 Series	https://cdrdv2.intel.com/v1/dl/getContent/335253
Intel® Ethernet 300 Series	N/A
Intel® Ethernet 200 Series	N/A

6.2 Specification Updates

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

Device Series	Support Link
Intel® Ethernet 800 Series	https://cdrdv2.intel.com/v1/dl/getContent/616943
Intel® Ethernet 700 Series: – X710/XXV710/XL710 – X710-TM4/AT2 and V710-AT2	https://cdrdv2.intel.com/v1/dl/getContent/331430 https://cdrdv2.intel.com/v1/dl/getContent/615119
Intel® Ethernet 500 Series – X550 – X540	https://cdrdv2.intel.com/v1/dl/getContent/333717 https://cdrdv2.intel.com/v1/dl/getContent/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 Intel Product Security Center Advisories

Intel product security center advisories can be found at:

<https://www.intel.com/content/www/us/en/security-center/default.html>



LEGAL

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

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

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

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

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

Copies of documents 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.