

Product Brief

Intel® Ethernet Server Adapter I350

Quad-port Gigabit Ethernet server adapters designed with performance-enhancing features and power management technologies



Key Features

- High-performing bridgeless design supporting PCI Express (PCIe) 2.15GT/s
- Flexible I/O virtualization for port partitioning and quality of service (QoS) of up to 32 virtual ports
- Scalable iSCSI performance delivering cost-effective SAN connectivity
- Innovative power management features including Energy Efficient Ethernet (EEE) and DMA Coalescing for increased efficiency and reduced power consumption
- Reliable and proven Gigabit Ethernet technology from Intel Corporation

Overview

The quad-port Intel® Ethernet Server Adapter I350 builds on Intel's history of excellence in Ethernet products with this generation of PCIe GbE network adapters.

Built with the bridgeless Intel® Ethernet Controller I350, this network adapter introduces new levels of performance through industry-leading enhancements for both virtualized and iSCSI environments. Intel uses a suite of hardware assists that improve overall system performance by lowering the I/O overhead in virtualized environments. iSCSI simplifies SAN connectivity by eliminating the need for SAN-specific adapters or switches.

This adapter also includes power management technologies such as Energy Efficient Ethernet (EEE) and DMA Coalescing (DMAC). With advanced Power Management Technologies, customers can configure power options on the adapter and more effectively manage their power consumption.

The Intel® Ethernet I350 Server Adapters includes these featured technologies:

Flexible I/O Virtualization

The Intel® Ethernet I350 adapters include Intel® Virtualization Technology for connectivity (Intel VT-c) to deliver I/O virtualization and Quality of Service (QoS) features designed directly into the I350 controller on the adapter. I/O virtualization advances network connectivity models used in today's servers to more efficient models by providing Flexible Port Partitioning (FPP), multiple Rx/Tx queues, and on-controller QoS functionality that can be used in both virtual and non-virtual server deployments.

By taking advantage of the PCI-SIG SR-IOV specification, Intel® Ethernet products enable Flexible Port Partitioning (FPP). With FPP, virtual controllers can be used by the Linux host directly and/or assigned to virtual machines. With this port partitioning, administrators can create up to eight dedicated networks on a single Ethernet port for use in baremetal and virtualized server deployments.

In a bare-metal Linux server, host processes can be assigned to dedicated network resources to provide traffic isolation and balanced bandwidth allocation.

In a virtualized environment, a VM can be assigned to a virtual controller to reduce the CPU overhead seen when using a software-based network bridge, by offloading network traffic management to the Ethernet controller silicon.

Scalable iSCSI Performance

iSCSI uses Ethernet to carry storage traffic, extending the familiarity and simplicity of Ethernet to storage networking, without the need for SAN-specific adapters or switches. Intel® Ethernet adapters, with native iSCSI initiators built into Microsoft Windows, Linux, and VMware ESXi platforms, provide a simple, dependable, cost-effective way to connect to LANs and iSCSI SANs.

Intel® Ethernet adapters include hardware-based iSCSI acceleration features that do not require offloading to a proprietary TCP/IP stack. iSCSI acceleration uses large send offload, Receive Side Coalescing and transmit send offloads to help reduce latency and lower CPU utilization. To improve efficiency, MSI-X, Receive-side Scaling and Intel® Ethernet Flow Director, scale I/O processing across multiple CPU cores. Direct memory access (DMA), direct cache access (DCA) and header splitting improve network data processing efficiency, and data center bridging (DCB) supports multiple traffic classes that can be prioritized for iSCSI traffic.

These native initiators are broadly tested using multiple generations of operating systems, storage systems, and OS tools to help ensure reliability and ease of use. Standardizing on Intel Ethernet for iSCSI allows administrators to use a single initiator, TCP/IP stack, and a common set of management tools and IT policies.

In addition, Intel Ethernet adapters include a number of hardware features designed to accelerate iSCSI traffic and enhance data processing. For example, TCP segmentation offload, Receive Side Coalescing (RSC), and checksum offload capabilities help reduce processor usage, increase throughput, and deliver exceptional iSCSI performance.

iSCSI uses your existing Ethernet infrastructure to connect to remote storage units using the SCSI protocol encapsulated in standard TCP/IP packets. Intel iSCSI Remote Boot enables PCI Express-based Intel® Ethernet Server Adapters to boot from a remote iSCSI disk volume on an iSCSI-based storage area network (SAN).

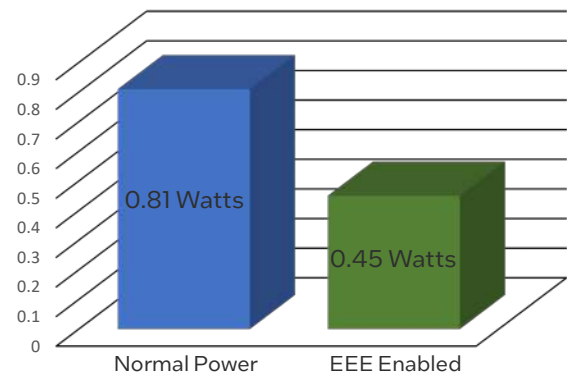
Finally, using native OS initiators, Intel Ethernet adapters enable support for the CRC-32 digest instruction set included with Intel® Xeon® processor products, which improves transmission reliability, delivering an enterprise-class iSCSI solution for the IT customer.

Power Management Technologies

Today, companies everywhere are looking for ways to decrease energy consumption across the enterprise to reduce costs and environmental impact, while at the same time solving increasingly important power density challenges. That's why Intel has introduced new, advanced Power Management Technologies (PMTs) with the Intel Ethernet I350 Server Adapter that enable enterprises to configure power options on the adapter and more effectively manage their power consumption.

Energy Efficient Ethernet (EEE)

The Intel Ethernet I350 Server Adapter supports the IEEE802.3az Energy Efficient Ethernet (EEE) standard so that, during periods of low network activity, EEE reduces the power consumption of an Ethernet connection by negotiating with a compliant EEE switch port to transition to a low power idle (LPI) state. This reduces the controller power to approximately 50 percent of its normal operating power, saving power on the network port and the switch port. As soon as increased network traffic is intelligently detected, the controller on the platform and the switch quickly come back to full power to handle the increased network traffic. EEE is supported for both 1000BASE-T and 100BASE-TX.



Power Savings with EEE

Source: Intel Labs
Energy Efficient Ethernet reduces the controller power to approximately 50 percent of its normal operating level.

DMA Coalescing

Another power management technology that can reduce power on the server platform is DMA Coalescing (DMAC). Typically, when a packet arrives at a server, DMA calls are made to transfer the packet within the server. These calls wake up the processor, memory and other system components from a lower power state in order to perform the tasks required to handle the incoming packet.

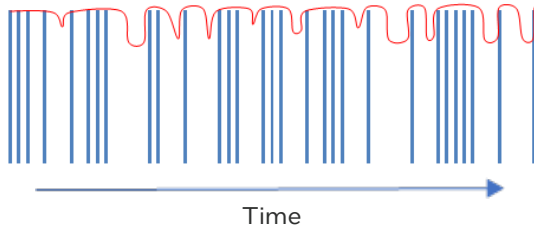
Based on the configurable DMAC settings, incoming packets are buffered momentarily before any DMA calls are made. This enables the controller to intelligently identify opportunities to batch multiple packets together so that when components are wakened from lower power states they can efficiently handle the batched packets at the same time. Platform components are able to remain in lower power states longer, which can dramatically reduce platform energy consumption. DMAC synchronizes DMA calls across all controller ports to ensure maximum power savings.

Software Tools and Management

Intel Ethernet adapters support Dell's Lifecycle Controller. The Lifecycle Controller is coupled with the Dell DRAC service processor to provide embedded system management. The Lifecycle Controller enables both local and remote access to manage initial setup and configuration of the BIOS settings on the platform, setup and configuration of Intel Ethernet adapters, update of all the platform firmware, and the deployment of the operating systems.

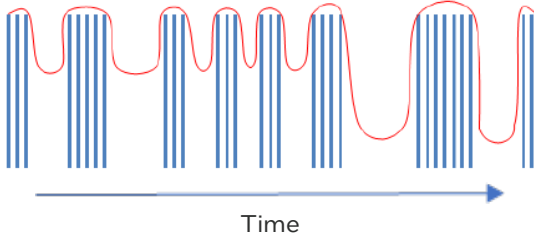
Intel® Advanced Network Services (Intel® ANS) include new teaming technologies and techniques such as Virtual Machine Load-Balancing (VMLB) for Hyper-V environments. Intel ANS also provides a variety of teaming configurations for up to eight ports, and support for teaming mixed vendors' server adapters. Intel ANS includes support for 802.1Q VLANs, making Intel ANS one of the most capable and comprehensive tools for supporting server adapter teaming.

Without DMA Coalescing



As shown by the red line, components have less time between DMA calls to reach and stay in lower power.

With DMA Coalescing



With more time between DMA calls, components can reach lower power states and remain in them longer.

Features	Description
General	
Intel® Ethernet Controller I350 with PCI Express 2.1 (5 GT/s) Support	<ul style="list-style-type: none"> Industry-leading non-bridged PCIe 2.1 quad-port 1GbE controller Enables customers to take full advantage of 1GbE by providing maximum bi-directional throughput per port on a single quad-port adapter
Halogen Free (Copper)	<ul style="list-style-type: none"> Leadership in an environmentally friendly ecosystem
Low-Profile	<ul style="list-style-type: none"> Enables higher bandwidth and throughput from standard and low-profile PCIe slots and servers
Ethernet	
IEEE 802.3 auto-negotiation	<ul style="list-style-type: none"> Automatic link configuration for speed, duplex, flow control
1Gb/s Ethernet IEEE 802.3, 802.3u, 802.3ab PHY specifications Compliant	<ul style="list-style-type: none"> Robust operation over installed base of CAT-5 twisted-pair cabling
Integrated PHY for 10/100/1000 Mb/s for multispeed, full, and half-duplex	<ul style="list-style-type: none"> Smaller footprint and lower power dissipation compared to multiple discreet MAC and PHY
IEEE 802.3x and 802.3z compliant flow control support with software-controllable Rx thresholds and Tx pause frames	<ul style="list-style-type: none"> Local control of network congestion levels Frame loss reduced from receive overruns
Automatic cross-over detection function (MDI/MDI-X)	<ul style="list-style-type: none"> The PHY automatically detects which application is being used and configures itself accordingly
IEEE 1588 protocol and 802.1AS implementation	<ul style="list-style-type: none"> Time-stamping and synchronization of time-sensitive applications Distribute common time to media devices
Power Management and Efficiency	
<1W S0-Max (state) 1000BASE-T Active 90 °C (mode) <400mW S0-Typ (state) 100BASE-T Active (mode)	<ul style="list-style-type: none"> Controller is designed for low power consumption
IEEE 802.3az - Energy Efficient Ethernet (EEE)	<ul style="list-style-type: none"> Power consumption of the PHY is reduced by approximately 50 percent link transitions to low power Idle (LPI) state as defined in the IEEE 802.3az (EEE) standard
DMA Coalescing	<ul style="list-style-type: none"> Reduces platform power consumption by coalescing, aligning, and synchronizing DMA Enables synchronizing port activity and power management of memory, CPU and RC internal circuitry
Smart Power Down (SPD) at S0 no link / Sx no link	<ul style="list-style-type: none"> PHY powers down circuits and clocks that are not required for detection of link activity
Active State Power Management (ASPM) Support	<ul style="list-style-type: none"> Optionality Compliance bit to help determine whether to enable ASPM or whether to run ASPM compliance tests to support entry to L0s
LAN disable function	<ul style="list-style-type: none"> Option to disable the LAN Port and/or PCIe Function. Disabling just the PCIe function but keeping the LAN port that resides on it fully active (for manageability purposes and BMC pass-through traffic).
Full wake up support <ul style="list-style-type: none"> Advanced Power Management (APM) Support (formerly Wake on LAN) Advanced Configuration and Power Interface (ACPI) specification v2.0c Magic Packet wake-up enable with unique MAC address 	<ul style="list-style-type: none"> APM - Designed to receive a broadcast or unicast packet with an explicit data pattern (Magic Pack) and assert a signal to wake up the system ACPI - PCIe power management based wake-up that can generate system wake-up events from a number of sources
ACPI register set and power down functionality supporting D0 and D3 states	<ul style="list-style-type: none"> A power-managed link speed control lowers link speed (and power) when highest link performance is not required
MAC Power Management controls	<ul style="list-style-type: none"> Power management controls in the MAC the PHY can be entered into a low-power state
Low Power Link Up - Link Speed Control	<ul style="list-style-type: none"> Enables a link to come up at the lowest possible speed in cases where power is more important than performance
Power Management Protocol Offload (Proxying)	<ul style="list-style-type: none"> Avoid spurious wake up events and reduce system power consumption when the device is in D3 low power state and system is in S3 or S4 low power states
Latency Tolerance Reporting (LTR)	<ul style="list-style-type: none"> Reports service latency requirements for memory reads and writes to the Root Complex for system power management
Manageability Features	
Management Component Transport Protocol (MCTP)	<ul style="list-style-type: none"> Baseboard management controller (BMC) communication between add-in devices using a standardized protocol
Dell Lifecycle Controller	<ul style="list-style-type: none"> Local and remote access to BIOS setup and configuration on the platform and adapter
Management Component Transport Protocol (MCTP)	<ul style="list-style-type: none"> Baseboard management controller (BMC) communication between add-in devices using a standardized protocol
Firmware Based Thermal Management	<ul style="list-style-type: none"> Can be programmed via the BMC to initiate thermal actions and report thermal occurrences
IEEE 802.3 MII Management Interface	<ul style="list-style-type: none"> Enables the MAC and software to monitor and control the state of the PHY
MAC/PHY Control and Status	<ul style="list-style-type: none"> Enhanced control capabilities through PHY reset, link status, duplex indication, and MAC Dx power state
Watchdog timer	<ul style="list-style-type: none"> Gives an indication to the manageability firmware or external devices that the chip or the driver is not functioning
Extended error reporting	<ul style="list-style-type: none"> Messaging support to communicate multiple types/severity of errors
Controller Memory Protection	<ul style="list-style-type: none"> Main internal memories are protected by error correcting code (ECC) or parity bits
Vital Product Data (VPD) Support	<ul style="list-style-type: none"> Support for VPD memory area

Features	Description
I/O Virtualization	
Virtual Machine Device queues (VMDq) ¹	<ul style="list-style-type: none"> Offloads the data-sorting functionality from the Hypervisor to the network silicon, improving data throughput and CPU usage Provides QoS feature on the Tx data by providing round-robin servicing and preventing head-of-line blocking Sorting based on MAC addresses and VLAN tags
Next-Generation VMDq (8 queues per port)	<ul style="list-style-type: none"> Enhanced QoS feature by providing weighted round-robin servicing for the Tx data Provides loopback functionality, where data transfer between the virtual machines within the same physical server need not go out to the wire and come back in, improving throughput and CPU usage Supports replication of multicast and broadcast data
Eight transmit (Tx) and receive (Rx) queue pairs per port	<ul style="list-style-type: none"> Supports VMware NetQueue and Microsoft VMQ
Flexible Port Partitioning: 8 Virtual Functions per port	<ul style="list-style-type: none"> Virtual Functions (VFs) appear as Ethernet Controllers in Linux OSs that can be assigned to VMs, Kernel processes or teamed using the Linux Bonding Drivers
Support for PCI-SIG SR-IOV specification	<ul style="list-style-type: none"> 8 Virtual Functions per Port
Rx/Tx Round-Robin Scheduling	<ul style="list-style-type: none"> Assigns time slices in equal portions in circular order for Rx/ Tx for balanced bandwidth allocation
Traffic Isolation	<ul style="list-style-type: none"> Processes or VMs can be assigned a dedicated VF with VLAN support
Traffic Steering	<ul style="list-style-type: none"> Offloads sorting and classifying traffic in to VF or queues
VM to VM Packet forwarding (Packet Loopback)	<ul style="list-style-type: none"> On-chip VM-VM traffic allows PCIe speed switching between VM
MAC and VLAN anti-spoofing	<ul style="list-style-type: none"> Enables anti-spoofing filter on MAC addresses and VLAN for VFs
Malicious driver detection	<ul style="list-style-type: none"> Monitors queues and VFs for malformed descriptors that might indicate a malicious or buggy driver
Storm control	<ul style="list-style-type: none"> Limits to the broadcast or multicast traffic it can receive
Per-pool statistics, offloads, and jumbo support	<ul style="list-style-type: none"> Each Queue Pair or Pool has its own statistics, offloads and Jumbo support options
Independent Function Level Reset (FLR) for Physical and Virtual Functions	<ul style="list-style-type: none"> VF resets only the part of the logic dedicated to specific VF and does not influence the shared port
IEEE 802.1Q Virtual Local Area Network (VLAN) support with VLAN tag insertion, stripping and packet filtering for up to 4096 VLAN tags	<ul style="list-style-type: none"> Adding (for transmits) and removing (for receives) of VLAN tags with no VM involvement Filtering packets belonging to certain VLANs
IEEE 802.1Q advanced packet filtering	<ul style="list-style-type: none"> Lower processor utilization
Mirroring rules	<ul style="list-style-type: none"> Ability to reflect network traffic to a given VM or VLAN based on up to four rules
Support for Simple VEPA	<ul style="list-style-type: none"> Support for external VM switching
VF Promiscuous modes	<ul style="list-style-type: none"> VLAN, unicast, multicast
Stateless Offloads/Performance	
TCP/UDP, IPv4 checksum offloads (Rx / Tx/Large-send); Extended Tx descriptors for more offload capabilities	<ul style="list-style-type: none"> Improved CPU usage Checksum and segmentation capability extended to new standard packet type
IPv6 support for IP/ TCP and IP/UDP receive checksum offload	<ul style="list-style-type: none"> Improved CPU usage
Tx TCP segmentation offload (IPv4, IPv6)	<ul style="list-style-type: none"> Increased throughput and lower processor usage Compatible with large-send offload
Transmit Segmentation Offloading (TSO)	<ul style="list-style-type: none"> Large TCP/UDP I/O is segmented by the device to L2 packets according to the requested MSS
Interrupt throttling control	<ul style="list-style-type: none"> Limits maximum interrupt rate and improves CPU utilization
Legacy and Message Signal Interrupt (MSI) Modes	<ul style="list-style-type: none"> Interrupt mapping
Message Signal Interrupt Extension (MSI-X)	<ul style="list-style-type: none"> Dynamic allocation of up to 25 vectors per port
Intelligent interrupt generation	<ul style="list-style-type: none"> Enhanced software device driver performance
Receive Side Scaling (RSS) for Windows environment Scalable I/O for Linux environments (IPv4, IPv6, TCP/UDP)	<ul style="list-style-type: none"> Up to eight queues per port
Support for packets up to 9.5K Bytes (Jumbo Frames)	<ul style="list-style-type: none"> Improves the system performance related to handling of network data on multiprocessor systems Enables higher and better throughput of data
Low Latency Interrupts	<ul style="list-style-type: none"> Based on the sensitivity of the incoming data, the controller can bypass the automatic moderation of time intervals between the interrupts
Header/packet data split in receive	<ul style="list-style-type: none"> Helps the driver to focus on the relevant part of the packet without the need to parse it
PCIe 2.1 TLP Processing Hint Requester	<ul style="list-style-type: none"> Provides hints on a per-transaction basis to facilitate optimized processing of transactions that would target memory space
Descriptor ring management hardware for Transmit and Receive	<ul style="list-style-type: none"> Optimized descriptor fetch and write-back for efficient system memory and PCIe bandwidth usage

Remote Boot Options

Preboot eXecution Environment (PXE) flash interface support	<ul style="list-style-type: none"> Enables system bootup via the EFI (32 bit and 64 bit) Flash interface for PXE 2.1 option ROM
Intel® iSCSI Remote Boot for Windows, Linux, and VMware	<ul style="list-style-type: none"> Enables system bootup via iSCSI Provides additional network management capability
Intel Boot Agent software: Linux boot via PXE or BOOTP, Windows Deployment Services, or UEFI	<ul style="list-style-type: none"> Allows networked computer to boot using a program code image supplied by a remote server Complies with the Pre-boot eXecution Environment (PXE) Version 2.1 Specification

Adapter Product Features

Plug and play specification support	Standard
Intel® I/O Acceleration Technology (Intel® I/OAT)	Extreme system throughput
Ships with full-height bracket installed; low-profile bracket included in package (T2 and T4)	Streamlines installation
Cable distance	Copper: up to 100 m

Specifications

Connectors	RJ45 (Copper)
IEEE standard/network topology	EEE 802.3/10BASE-T, 100BASE-TX, 1000BASE-T
Cabling: Copper	CAT3 or higher for 10BASE-T operation CAT5 or higher for 100BASE-TX operation CAT5e or higher for 1000BASE-T operation
Data rate supported per port	10/100/1000 Mbps (Copper)
Bus type	PCI Express v2.1 (5 GT/s)
Bus width	4-lane PCI Express; operable in x4, x8 and x16 slots
Interrupt levels	INTA, INTB, INTC, INTD, MSI, MSI-X
Hardware certifications	FCC B, UL, CE, VCCI, BSMI, CTICK, KCC
Controller	Intel® Ethernet Controller I350
Power consumption (typical)	I350T4 5.0 W
Operating temperature	0 °C to 55 °C (32 °F to 131 °F)
Storage temperature	-40 °C to 70 °C (-40 °F to 158 °F)
Storage humidity	90% non-condensing relative humidity at 35 °C
Connect Speed LED Indicators	Not illuminated=10 M b/s; green=100 M b/s; amber=1 G b/s (Copper)

Physical Dimensions

Length	13.54 cm (5.33 in.)
Width	6.89 cm (2.71 in.)
Full-height end bracket	12.0 cm (4.725 in.)
Low-profile end bracket	7.92 cm (3.117 in.)

Operating System/Architecture Support

Operating System	IA32	X64	IPFI
Windows XP Professional SP3		▪	
Windows XP Professional SP3	▪		
Windows Vista SP2	▪	▪	
Windows 7 SP1	▪	▪	
Windows Server 2003 SP2	▪	▪	▪
Windows Server 2008 SP2	▪	▪	▪
Windows Server 2008 SP2 Core	▪	▪	
Windows Server 2008 SP2 (w/Hyper-V role)		▪	
Hyper-V Server 2008 SP2 (stand-alone version)		▪ ¹	
Windows Server 2008 R2 SP1		▪	▪
Windows Server 2008 R2 SP1 Core		▪	
Windows Server 2008 R2 SP1 (w/Hyper-V role)		▪	
Hyper-V Server 2008 R2 SP1 (stand-alone version)		▪ ¹	
Linux Stable Kernel version 2.6	▪	▪	
Linux RHEL 5.5	▪	▪	
Linux RHEL 6.0	▪	▪	
Linux SLES 10 SP3	▪	▪	
Linux SLES 11 SP1	▪	▪	▪
FreeBSD 8.0	▪	▪	
DOS NDIS 2	▪		
DOS ODI	▪		
EFI 1.1			▪
uEFI 2.1		▪	▪
VMware ESX 4.0 ²		▪	
VMware ESX 4.1 ²		▪	
VMware ESX M/N ²		▪	
Xen ³		▪	

Key: ▪ =affected
(blank)=OS not available on specified architecture

Product Order Codes for tFRU-less* Versions

Description	SKU	DPN	Intel MM#
I350-T4, LP	Factory installed: 540-BBCZ Customer kit: 540-BBDV	NM8TT	99AWH5
I350-T4, FH	Factory installed: 540-BBDD Customer kit: 540-BBDS	Y10X9	99AWH4

*The minimum software required for the Dell tFRU-less cards is Dell (FFV): 19.5, 20.0 or 20.5

Dell Tech Backing Information

Standard one-year warranty.

To see the full line of Intel Ethernet Network Adapters visit www.dell.com or contact your Dell Technologies sales representative.

1-Minimal Validation

2-VMware ESX drivers are not included on LAD SW Release CDs; they are only available from VMware's web site and they are released on a separate sc

3-SR-IOV validation only

The information contained in this document, including all instructions, cautions, and regulatory approvals and certifications, is provided by Intel and has not been independently verified or tested by Dell. Dell cannot be responsible for damage caused as a result of either following or failing to follow these instructions. All statements or claims regarding the properties, capabilities, speeds or qualifications of the part referenced in this document are made by Intel and not by Dell. Dell specifically disclaims knowledge of the accuracy, completeness or substantiation for any such statements. All questions or comments relating to such statements or claims should be directed to Intel. Visit www.dell.com for more information.

INFORMATION IN THIS DOCUMENT IS PROVIDED IN CONNECTION WITH INTEL® PRODUCTS. NO LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS IS GRANTED BY THIS DOCUMENT. EXCEPT AS PROVIDED IN INTEL'S TERMS AND CONDITIONS OF SALE FOR SUCH PRODUCTS, INTEL ASSUMES NO LIABILITY WHATSOEVER, AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO SALE AND/OR USE OF INTEL PRODUCTS INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT OR OTHER INTELLECTUAL PROPERTY RIGHT. UNLESS OTHERWISE AGREED IN WRITING BY INTEL, THE INTEL PRODUCTS ARE NOT DESIGNED NOR INTENDED FOR ANY APPLICATION IN WHICH THE FAILURE OF THE INTEL PRODUCT COULD CREATE A SITUATION WHERE PERSONAL INJURY OR DEATH MAY OCCUR.

Intel may make changes to specifications and product descriptions at any time, without notice. Designers must not rely on the absence or characteristics of any features or instructions marked "reserved" or "undefined." Intel reserves these for future definition and shall have no responsibility whatsoever for conflicts or incompatibilities arising from future changes to them. The information here is subject to change without notice. Do not finalize a design with this information. The products described in this document may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request. Contact your local Intel sales office or your distributor to obtain the latest specifications and before placing your product order. Copies of documents which have an order number and are referenced in this document, or other Intel literature, may be obtained by calling 1-800-548-4725, or by visiting Intel's Web Site at <http://www.intel.com>.

© 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.