PRODUCT BRIEF
Intel® Ethernet I350 Server Adapters
Network Connectivity

Intel® Ethernet I350 Server Adapters

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

Key Features

- Halogen-free dual- or quad-port gigabit Ethernet adapters with copper or fiber interface options
- Innovative power management features including Energy Efficient Ethernet (EEE) and DMA Coalescing for increased efficiency and reduced power consumption
- 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
- High-performing bridgeless design supporting PCI Express* Gen 2.0 5GT/s
- Reliable and proven Gigabit Ethernet technology from Intel Corporation

Overview

The new Intel® Ethernet Server Adapter I350 family builds on Intel's extended history of excellence in Ethernet products. Intel continues its market leadership with this new generation of PCIe* GbE network adapters. Built with the bridgeless Intel® I350 Ethernet Controller, these adapters represent the next step in the GbE networking evolution for the enterprise and data center by introducing new levels of performance through industry-leading enhancements for both virtualized and iSCSI Unified Networking environments. This new family of adapters also includes new power management technologies such as Energy Efficient Ethernet (EEE) and DMA Coalescing (DMAC).

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 bare-metal 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

Intel® Ethernet I350 Server Adapters with native iSCSI initiators built into Microsoft® Windows®, Linux®, and VMware® ESX platforms provide a simple, dependable, cost-effective way to connect to LANs and iSCSI SANs. 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 includes 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. Finally, using native OS initiators, an Intel® Ethernet I350 Server Adapter enables support for the CRC-32 digest instruction set included with Intel® Xeon® processor products, which improves transmission

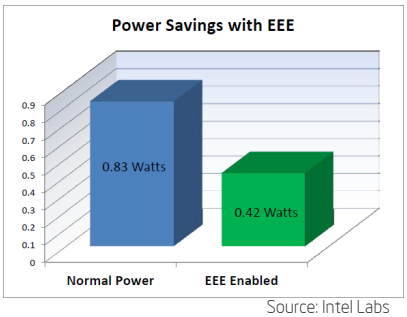

Source: Intel Labs

Energy Efficient Ethernet reduces the controller power to approximately 50% of its normal operating level.

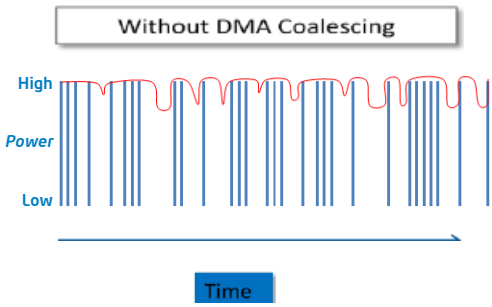

As shown by the red line, components have less time between DMA calls to reach and stay in lower power

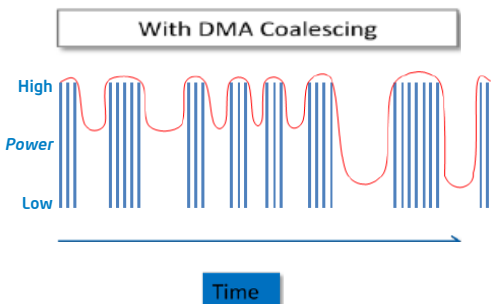

With more time between DMA calls, components can reach lower power states and remain in them longer.

reliability and thus delivers 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 family 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 family 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% 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.

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. This enables platform components 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.

These and the additional Power Management Technologies included with Intel® Ethernet will help you more effectively management your power challenges.

Software Tools and Management

In 1996, Intel Corporation introduced Intel® Adapter Fault Tolerance software, since then, Intel® Advanced Network Services (Intel® ANS), as they are now called, have continued evolving to include new teaming technologies and techniques such as Virtual Machine Load-Balancing (VMLB) for Hyper-V environments. Today, Intel ANS includes a variety of teaming configurations for up to eight adapters, support for mixed vendors server adapters teaming and includes support for 802.1q VLANs, making Intel ANS one of the most capable and comprehensive tools for supporting server adapter teaming.

Additionally, Intel® PROSet for Windows* Device Manager (DMIX) and PROsetCL extends driver functionality to provide additional reliability and Quality of Service features and configuration.

General Features

Features	Benefits
Intel® I350 Gigabit Ethernet Controller with PCI Express* V2.0 (5 GT/s) Support	▪ Industry-leading smallest non-bridged PCIe Gen2 quad-port 1 GbE controller ▪ Enables customers to take full advantage of 1 GbE by providing maximum bi-directional throughput per port on a single quad-port adapter
Halogen Free[1] (Copper)	▪ Leadership in an environmentally friendly ecosystem
Low-Profile (Dual and Quad Port Copper; Dual-Port Fiber) and Standard height (Quad-Port Fiber)	▪ Enables higher bandwidth and throughput from standard and low-profile PCIe slots and servers

Ethernet Features

Features	Benefits
IEEE 802.3* auto-negotiation	▪ Automatic link configuration for speed, duplex, flow control
1Gb/s Ethernet IEEE 802.3, 802.3u, 802.3ab PHY specifications Compliant	▪ Robust operation over installed base of Category-5 twisted-pair cabling
Integrated PHY for 10/100/1000 Mb/s for multispeed, full, and half-duplex	▪ 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	▪ Local control of network congestion levels ▪ Frame loss reduced from receive overruns
Automatic cross-over detection function (MDI/MDI-X)	▪ The PHY automatically detects which application is being used and configures itself accordingly
IEEE 1588 protocol and 802.1AS implementation	▪ Time-stamping and synchronization of time sensitive applications ▪ Distribute common time to media devices

Power Management and Efficiency

Features	Benefits
<1W S0-Max (state) 1000BASE-T Active 90oC (mode) <400mW S0-Typ (state) 100BASE-T Active (mode)	▪ Controller is designed for low power consumption
IEEE802.3az - Energy Efficient Ethernet (EEE)	▪ Power consumption of the PHY is reduced by approximately 50% link transitions to low power Idle (LPI) state as defined in the IEEE802.3az (EEE) standard
DMA Coalescing	▪ 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	▪ PHY powers down circuits and clocks that are not required for detection of link activity
Active State Power Management (ASPM) Support	▪ 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	▪ 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 ▪ 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	▪ 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	▪ A power-managed link speed control lowers link speed (and power) when highest link performance is not required
MAC Power Management controls	▪ Power management controls in the MAC the PHY can be entered into a low-power state
Low Power Link Up - Link Speed Control	▪ 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)	▪ 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)	▪ Reports service latency requirements for memory reads and writes to the Root Complex for system power management

I/O Virtualization Features

Features	Benefits
Eight transmit (Tx) and receive (Rx) queue pairs per port	• Supports VMware* NetQueue and Microsoft* VMQ
Flexible Port Partitioning: 32 Virtual Functions on Quad-port or 16 Virtual Functions on Dual-port	• Virtual Functions (VFs) appear as Ethernet Controllers in Linux OSes that can be assigned to VMs, Kernel processes or teamed using the Linux* Bonding Drivers
Support for PCI-SIG SR-IOV specification	• Up to 8 Virtual Functions per Port
Rx/Tx Round-Robin Scheduling	• Assigns time slices in equal portions in circular order for Rx/Tx for balanced bandwidth allocation
Traffic Isolation	• Processes or VMs can be assigned a dedicated VF with VLAN support
Traffic Steering	• Offloads sorting and classifying traffic in to VF or queues
VM to VM Packet forwarding (Packet Loopback)	• On-chip VM-VM traffic allows PCIe* speed switching between VM
MAC and VLAN anti-spoofing	• Enables anti spoofing filter on MAC addresses and VLAN for VFs.
Malicious driver detection	• Monitors queues and VFs for malformed descriptors that might indicate a malicious or buggy driver.
Storm control	• Limits to the broadcast or multicast traffic it can receive
Per-pool statistics, off loads, and jumbo support	• Each Queue Pair or Pool has its own statistics, off-loads and Jumbo support options
Independent Function Level Reset (FLR) for Physical and Virtual Functions	• 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	• 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	• Lower processor utilization
Mirroring rules	• Ability to reflect network traffic to a given VM or VLAN based on up to four rules
Support for Simple VEPA	• Support for external VM switching
VF Promiscuous modes	• VLAN, unicast, multicast

Stateless Offloads/Performance Features

Features	Benefits
TCP/UDP, IPv4 checksum offloads (Rx/ Tx/Large-send); Extended Tx descriptors for more offload capabilities	• Improved CPU usage • Checksum and segmentation capability extended to new standard packet type
IPv6 support for IP/TCP and IP/UDP receive checksum offload	• Improved CPU usage
Tx TCP segmentation offload (IPv4, IPv6)	• Increased throughput and lower processor usage • Compatible with large-send offload
Transmit Segmentation Offloading (TSO)	• Large TCP/UDP I/O is segmented by to the device it to L2 packets according to the requested MSS
Interrupt throttling control	• Limits maximum interrupt rate and improves CPU utilization
Legacy and Message Signal Interrupt (MSI) Modes	• Interrupt mapping
Message Signal Interrupt Extension (MSI-X)	• Dynamic allocation of up to 25 vectors per port
Intelligent interrupt generation	• Enhanced software device driver performance
Receive Side Scaling (RSS) for Windows environment Scalable I/O for Linux environments (IPv4, IPv6, TCP/UDP	• Up to eight queues per port
Support for packets up to 9.5K Bytes (Jumbo Frames)	• Improves the system performance related to handling of network data on multiprocessor systems • Enables higher and better throughput of data
Low Latency Interrupts	• 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	• Helps the driver to focus on the relevant part of the packet without the need to parse it
PCIe v2.1 TLP Processing Hint Requester	• Provides hints on a per transaction basis to facilitate optimized processing of
Descriptor ring management hardware for Transmit and Receive	• Optimized descriptor fetch and write-back for efficient system memory and PCIe bandwidth usage

Remote Boot Options

Features	Benefits
Preboot eXecution Environment (PXE) flash interface support	• Enables system boot up 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	• Enables system boot up via iSCSI • Provides additional network management capability
Intel Boot Agent software: Linux boot via PXE or BOOTP, Windows* Deployment Services, or UEFI	• 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

Manageability Features

Features	Benefits
Management Component Transport Protocol (MCTP)	• Baseboard management controller (BMC) communication between add-in devices using a standardized protocol
Firmware Based Thermal Management	• Can be programmed via the BMC to initiate Thermal actions and report thermal occurrences
IEEE 802.3 MII Management Interface	• Enables the MAC and software to monitor and control the state of the PHY
MAC/PHY Control and Status	• Enhanced control capabilities through PHY reset, link status, duplex indication, and MAC Dx power state
Watchdog timer	• Defined by the FLASHT register to minimize Flash updates
Extended error reporting	• Messaging support to communicate multiple types/severity of errors
Controller Memory Protection	• Main internal memories are protected by error correcting code (ECC) or parity bits
Vital Product Data (VPD) Support	• Support for VPD memory area

Adapter Product Features

Plug and play specification support	• Standard
Intel® I/OAT	• Extreme system throughput
Ships with full-height bracket installed; low-profile bracket included in package (T2, T4 and F2)	• Streamlines installation
Cable distance	• Copper: up to 100 m • Fiber: up to 300 m

Intel Backing

Limited Lifetime Warranty
90-day, money-back guarantee (U.S. and Canada)

Specifications

General

Connectors	• Rj45 (Copper) • LC Fiber Optic (Fiber)
IEEE standard/network topology	• EEE 802.3/10BASE-T, 100BASE-TX, 1000BASE-T
Cabling Copper	• Category-3 or higher for 10BASE-T operation • Category-5 or higher for 100BASE-TX operation • Category-5e or higher for 1000BASE-T operation
Fiber	• MMF 62.5/50 um

Technical

Data rate supported per port:	• 10/100/1000 Mbps (Copper), 1000 Mbps (Fiber)
Bus type	• PCI Express* 2.0 (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-processor	• Intel® I350 Gigabit Ethernet Controler
Power consumption (typical)	• Copper: I350T2 4.4 W I350T4 5.0 W • Fiber: I350F2 5.5 W I350F4 6.0 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 Mb/s; green=100 Mb/s; amber=1 Gb/s (Copper) • Green = 1 Gb/s. Not illuminated = no link (Fiber)

Specifications *(continued)*

Physical Dimensions

Copper T2 & T4; Fiber F2	
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.)
Fiber F4	
Length	• 13.54 cm (5.33 in.)
Width	• 11.12 cm (4.376 in)
Full-height end bracket	• 12.0 cm (4.725 in.)
Low-profile end bracket	N/A

Operating System/Architecture Support

OPERATING SYSTEM	IA32	X64	IPF1
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)		▪[2]	
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)		▪[2]	
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[3]		▪	
VMware ESX 4.1[3]		▪	
VMware ESX M/N[3]		▪	
Xen[4]		▪	

Product Codes

DUAL PORT COPPER	CODE
Intel® Ethernet Server Adapter I350-T2	I350T2
Bulk Pack – Order 5, Get 5 – RJ45 IEEE	I350T2BLK
QUAD PORT COPPER	**CODE**
Intel® Ethernet Server Adapter I350-T4	I350T4
Bulk Pack – Order 5, Get 5 – RJ45 IEEE	I350T4BLK
DUAL PORT FIBER	**CODE**
Intel® Ethernet Server Adapter I350-F2	I350F2
Bulk Pack – Order 5, Get 5 – LC Fiber Optic IEEE	I350F2BLK
QUAD PORT FIBER	**CODE**
Intel® Ethernet Server Adapter I350-F4	I350F4
Bulk Pack – Order 5, Get 5 – LC Fiber Optic IEEE	I350F4BLK

Key: *=affected; No=not affected; (blank)=OS not available on specified architecture

1-Itanium™ Product Family

2-Minimal Validation

3-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

4-SR-IOV validation only

Customer Support

Intel® Customer Support Services offers a broad selection of programs including phone support and warranty service. For more information, contact us at

support.intel.com/support/go/network/

(Service and availability may vary by country.)

For Product Information

To speak to a customer service representative regarding Intel products, please call 1-800-538-3373 (U.S. and Canada) or visit

support.intel.com/support/go/network/contact.htm

for the telephone number in your area. For additional product information on Intel Networking Connectivity products, visit

www.intel.com/go/ethernet

To see the full line of Intel Ethernet Controllers, visit: www.intel.com/go/ethernet
For more information, contact your Intel sales representative.

NOTE: Low Halogen applies only to halogenated flame retardants and PVC in components. Halogens are below 1,000ppm bromine and 1,000ppm chlorine.

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 go to http://www.intel.com/design/literature.htm.

Copyright © 2011 Intel Corporation. All rights reserved. Intel, the Intel logo, and Xeon are trademarks of Intel Corporation in the U.S. and other countries.

*Other names and brands may be claimed as the property of others.

Printed in USA 0811/SU Please Recycle 325511-002US