PRODUCT BRIEF
Intel® Ethernet Controller I350
Network Connectivity

Intel® Ethernet Controller I350 Family

Dual- and Quad-port Gigabit Ethernet Controllers with advanced power-saving and flexible I/O virtualization technologies.

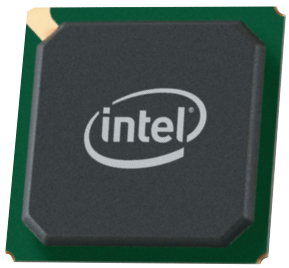

Key Features

- Halogen-free[1] dual- and quad-port gigabit Ethernet controller with multiple 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 on-controller quality of service (QoS) of up to 32 virtual ports
- Extended management support including MCTP, NC-SI or SMBus interfaces supporting IPMI pass-through and OS2BMC
- Scalable iSCSI performance delivering cost-effective SAN connectivity

Interface Options

- MDI (Copper) standard IEEE 802.3 Ethernet interface for 1000BASE-T, 100BASE-TX, 10BASE-T connections (IEEE 802.3, 802.3u, 802.3ab)
- SerDes interface for 1000BASE-SX/LX fiber connections
- 1000BASE-KX and 1000BASE-BX for backplane connections
- SGMII interface for SFP and external PHY connections

PCIe* Interface

- PCI Express* v 2.1 with 5.0 Gt/s and 2.5 Gt/s Support for x1, x2, x4 links widths

Overview

The Intel® Ethernet Controller I350 family builds on Intel's history of delivering Ethernet products with flexible design and in-box driver support. Run up to four 1 Gb ports with enhanced power-saving and market-leading flexible I/O virtualization including VMDq and SR-IOV. Intel's software drivers and support are unmatched for virtual or non-virtualized environments.

Power Management Technologies

Today, companies are looking to decrease energy consumption across the enterprise to reduce costs and environmental impact, while also solving increasingly important power density challenges. That's why Intel has introduced new, advanced Power Management Technologies (PMTs) with the I350 that enable enterprises to configure its power options and more effectively manage power consumption.

Energy Efficient Ethernet (EEE)

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

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

DMA Coalescing

Another PMT that will reduce power is DMA Coalescing (DMAC). Typically, when a packet arrives, 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 so that when components are wakened from lower power states they can efficiently handle all the batched packets at the same time. This enables components to remain in lower power states longer, which can dramatically reduce platform energy consumption. DMAC synchronizes DMA calls across all four ports to ensure maximum power savings.

These, and additional PMTs included with Intel® Ethernet, helps to more effectively manage your power challenges.

Multiple Integrated Interfaces

The family provides fully integrated gigabit Ethernet Media Access Control (MAC) and Physical-Layer (PHY), which has integrated power control components that can reduce board component cost and board layout space. The small package size increases board layout flexibility for all types of server and embedded designs. The family also provides fully integrated interface options to accommodate PHY front-access connectivity and SerDes back-plane connectivity in a single integrated package. The I350 Controller also has excellent thermal characteristics and operates at less than 1 W/port.

Flexible Design Configurations

The controller can be used for server system configurations such as rack-mounted or pedestal servers, in an add-on NIC, in LAN on Motherboard (LOM) designs and for blade servers. In the latter case, the I350 can support up to four SerDes ports in a LOM design or on a blade mezzanine card.

The single-chip, low-power I350 can provide significant BOM savings by reducing support components (bridge chips, crystals, and EEPROMS) required when compared with dual-port or multiple single-port GbE designs. This type of design makes the I350 excellent for embedded applications (switch add-on cards and network appliances).

Optimized Intelligent Offloads

Generation after generation the Intel Xeon® family of processors has demonstrated increased computing performance and increased integration of key server subsystems. For platform I/O, a balanced approach is to leverage the ever-increasing compute power of the processor and to implement the remaining accelerations in the network controller—this is what Intel refers to as "Intelligent Offloads." By balancing compute and offload, intelligent offloads can achieve optimal performance and efficiency.

The I/O intelligent offloads enable efficient multi-core performance in standard LAN environments, iSCSI storage connectivity, and virtualized deployments. LAN accelerations include TCP stateless offloads such as TCP/IP checksum, TCP segmentation, Receive Side Coalescing (RSC), and Receive Side Scaling (RSS). iSCSI accelerations benefit from CRC offloads as well as integrated host-based initiators running on the host processor. I/O Virtualization technologies such as VMDq (used in VMware* NetQueue and Microsoft* VMQ) provide increased throughput and reduced CPU usage by offloading the network processing in the hypervisors.

Furthermore, intelligent offloads provide the maximum flexibility of broad support and supporting features. In contrast, full offloads such as TOE (TCP/IP Offload Engines), are typically only supported in select environments and often require value-add system-level features to be disabled. Intelligent offloads strike the right balance of flexibility and efficiency, while maximizing performance over the lifetime of the platform.

Scalable iSCSI Performance

Intel® Ethernet Controllers 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 ensure reliability and ease of use. Standardizing on Intel Ethernet for iSCSI enables administrators to use a single initiator, a TCP/IP stack, and a common set of management tools and IT policies. In addition, Intel Ethernet includes hardware features to accelerate iSCSI traffic and enhance data processing. For example, TCP segmentation offload, Receive side coalescing (RSC), and checksum offload capabilities reduce processor usage, increase throughput, and deliver exceptional iSCSI performance. Finally, using native OS initiators, an Intel® Controller I350 enables support for the CRC-32 digest instruction set included with Intel Xeon processors, which improves reliability and delivers an enterprise-class iSCSI solution.

Flexible I/O Virtualization

The Intel® I350 includes Virtualization Technology for connectivity (VT-c) to deliver I/O virtualization and QoS features designed directly into the silicon. Using VT-c, network connectivity models used by servers today evolve to more efficient models by providing port partitioning, multiple Rx/Tx queues, and on-controller QoS functionality for both virtual and non-virtual servers.

At the core of Intel VT-c is a common QoS feature set that delivers a variety of capabilities used directly by an OS

or hypervisor or configured by an IT Administrator to meet a specific need.

VMDqs and PCI-SIG* SR-IOV are two of the common technologies used to enable the enhanced I/O virtualization functions. They share many of the same QoS features built into the silicon to provide native balanced bandwidth allocation and improved I/O scalability. These functions help reduce I/O bottlenecks and improve overall server performance by offloading functionality to the controller, which reduces the data processing associated with virtualization and enables process and application network segmentation in non-virtualized environments.

VMDq works in conjunction with VMware NetQueue or Microsoft Virtual Machine Queues (VMQ), in their respective hypervisors, to use the on-controller sorting and queuing functionality for traffic steering and Rx/Tx round-robin scheduling for balanced bandwidth allocation across multiple transmit and receive queues. These technologies enable the hypervisor to represent a network port as multiple ports that are assigned to the VMs to improve overall performance of I/O operations.

By taking advantage of the PCI-SIG SR-IOV specification, Intel Ethernet products enable virtual adapters that can be used by the Linux host directly and/or assigned to virtual machines. With this port partitioning, administrators can create multiple isolated 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.

Interfaces for Network Manageability

The controller family provides OS2BMC, SMBus and DMTF-defined Network Controller Sideband Interface (NC-SI) for BMC manageability. In addition, it introduces support for Management Component Transport Protocol (MCTP), a new DMTF standard, enabling a BMC to gather information about Intel Server Adapters that use the data rate, link speed, and error counts.

Software Tools and Management

In 1996, Intel Corporation introduced Intel® Adapter Fault Tolerance software, since then, Intel® Advanced Network Services (Intel® ANS), as it is now called, has 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.

External Interfaces	
PCI Express* Interface v 2.1	▪ 5.0 GT/s and 2.5 GT/s Support for x1, x2, x4 links widths (Lanes)
Network Interfaces	▪ MDI (Copper) standard IEEE 802.3 Ethernet interface for 1000BASE-T, 100BASE-TX, and 10BASE-T applications (802.3, 802.3u, and 802.3ab) ▪ Serializer-Deserializer (SERDES) to support 1000BASE-SX/LX (optical fiber - IEEE802.3) ▪ Serializer-Deserializer (SERDES) to support 1000BASE-KX (802.3ap) and 1000BASE-BX (PICMIG 3.1) for Gigabit backplane applications ▪ SGMII (Serial-GMII Specification) interface for SFP (SFP MSA INF-8074i)/external PHY connections
Management Interfaces	▪ Pass-Through (PT) Functionality via a sideband interface ▪ DMTF Network Controller Sideband Interface (NC-SI) ▪ Intel® System Management Bus (SMBus)

BOM Cost Reduction	
Features	**Benefits**
On-chip Linear Voltage regulator (LVR)	Removes the need for a higher cost on-board 1.8 V voltage regulator
On-chip Switched Voltage Regulator (SVR)	Removes need for a higher cost on-board 1.0 V voltage regulator
Autonomous on-die thermal management	Monitor on-die temperature and react when the temperature exceeds a pre-defined threshold

Ethernet Features	
IEEE* 802.3* auto-negotiation	▪ Automatic link configuration for speed, duplex, flow control
1 Gb/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 multi-speed, full, and half-duplex	▪ Smaller footprint and lower power dissipation compared to multiple discreet MAC and PHYs
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
Automatic cross-over detection function (MDI/MDI-X)	▪ Frame loss reduced from receive overruns
IEEE 1588 protocol and 802.1AS implementation	▪ The PHY automatically detects which application is being used and configures itself accordingly ▪ Time-stamping and synchronization of time sensitive applications ▪ Distribute common time to media devices
Power Management and Efficiency	
<1 W S0-Max (state) 1000BASE-T Active 90 °C <400 mW S0-Typ (state) 100BASE-T Active	▪ Controller is designed for low power consumption
IEEE802.3az - Energy Efficient Ethernet (EEE)	▪ Power consumption by 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)	▪ Optionality Compliance bit enables ASPM or runs 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	▪ Power-managed speed control lowers link speed/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)	▪ No spurious wake up and reduces power consumption in D3 low power state/S3 or S4 low power states
Latency Tolerance Reporting (LTR)	▪ Reports service latency requirements for memory reads and writes to the Root Complex

I/O Virtualization	
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* OSs that can be assigned to VMs, Kernel processes or teamed using the Linux* Bonding Drivers
Support for PCI-SIG SR-IOV specification	▪ Up to eight Virtual Functions per Port
Rx/Tx Round-Robin Scheduling	▪ Offloads sorting and classifying traffic in to VF or queues
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 usage
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 and Performance Features	
TCP/UDP, IPv4 checksum offloads (Rx/ Tx/Large-send); Extended Tx descriptors	▪ More offload capabilities and 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)	▪ Iinterrupt 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	▪ Up to eight queues per port
Scalable I/O for Linux environments (IPv4, IPv6, TCP/UDP)	▪ Improves the system performance related to handling of network data on multiprocessor systems
Support for packets up to 9.5 KB (Jumbo Frames)	▪ Enables higher and better throughput of data

Stateless Offloads and Performance Features *(continued)*	
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	
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	▪ Enables 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	
DMTF Network Controller Sideband Interface (NC-SI) Pass-through	▪ Supports pass through traffic between BMC and Controller's LAN functions ▪ Meets RMII Spec, Rev. 1.2 as a PHY-side device
Intel® System Management Bus (SMBus) Pass-through	▪ Enables BMC to configure the Controller's filters and management related capabilities.
Management Component Transport Protocol (MCTP)	▪ Used for baseboard management controller (BMC) communication between add-in devices
OS2BMC Traffic support	▪ Transmission and reception of traffic internally to communicate between the OS and local BMC
Private OS2BMC Traffic Flow	.▪BMC may have its own private connection to the network controller and network flows are blocked
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

Product Codes

I350-MM#	Brand Name	Description	Media	Forecast Name
915808	Intel® Ethernet Controller I350-AM4	17x17 Quad Port	tape and reel	NHI350AM4
915799	Intel® Ethernet Controller I350-AM4	17x17 Quad Port	tray	NHI350AM4
915809	Intel® Ethernet Controller I350-AM2	17x17 Dual Port	tape and reel	NHI350AM2
915801	Intel® Ethernet Controller I350-AM2	17x17 Dual Port	tray	NHI350AM2
915810	Intel® Ethernet Controller I350-AS4	17x17 Quad Port	tape and reel	NHI350AS4
915802	Intel® Ethernet Controller I350-AS4	17x17 Quad Port	tray	NHI350AS4
915811	Intel® Ethernet Controller I350-BT2	25x25 Dual Port	tape and reel	NHI350BT2
915804	Intel® Ethernet Controller I350-BT2	25x25 Dual Port	tray	NHI350BT2

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

For more information on the Intel® I350 Family, please visit **www.intel.com/go/ethernet**

Printed in USA 0711/SWU Please Recycle 325857-001US