Product Brief

Intel® Ethernet Server Adapter I350-T4V2

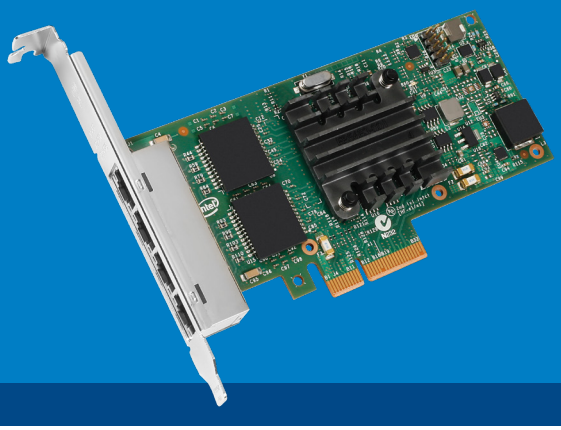

Quad-port Gigabit Ethernet server adapter designed with performance enhancing features and new power management technologies

Key Features

- Innovative power management features including Energy Efficient Ethernet (EEE), 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 (PCIe) 2.1 5GT/s
- Reliable and proven Gigabit Ethernet technology from Intel Corporation

Overview

Designed with the bridgeless Intel® Ethernet Controller I350, this 1GbE quad port adapter introduces new levels of performance through industry leading enhancements for both virtualized and iSCSI networking environments.

A suite of hardware assists improves 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 built-in iSCSI initiators for 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 built-in 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.

This adapter also 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.

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 PCIe-based Intel® Ethernet Server Adapters to boot from a remote iSCSI disk volume on an iSCSI-SAN.

Finally, using built-in OS initiators, Intel Ethernet Server 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. Advanced Power Management Technologies in the Intel Ethernet I350 Server Adapter 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. 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.

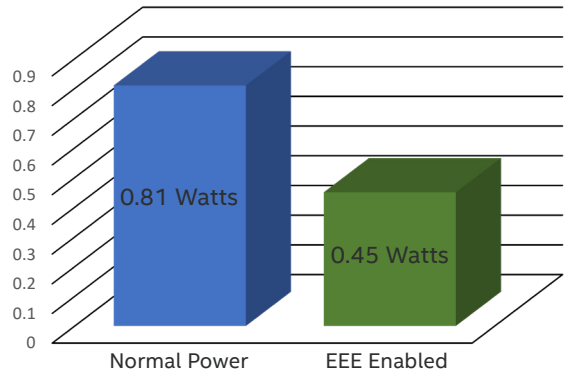

Power Savings with EEE

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.

Without DMA Coalescing

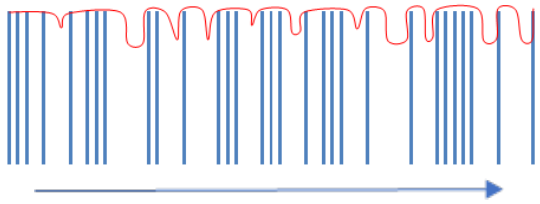

As shown by the red line, components have less time between DMA calls to reach and stay in lower power.

With DMA Coalescing

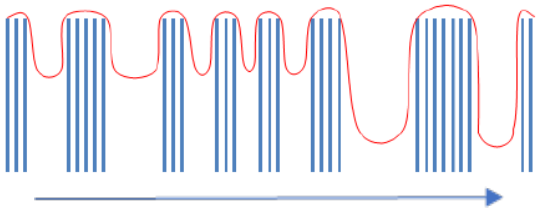

With more time between DMA calls, components can reach lower power states and remain in them longer.

Features	Description
General Features	
Intel® Ethernet Controller I350 with PCI Express 2.1 (5GT/s) Support	• Industry-leading smallest non-bridged PCIe Gen2 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[1] (Copper)	• Leadership in an environmentally firendly ecosystem.
Low profile (Quad Port Copper)	• Enables higher bandwidth and throughput from standard and low-profile PCIe slots and servers.
Ethernet Features	
IEEE 802.3 autonegotiation	• Automatic link configuration for speed duplex and flow control.
1Gb/s Ethernet IEEE 802.3, 802.3u, 802.3ab PHY specifications compliant	• Robust operation over installed base of CAT5 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 discrete MAC and PHYs.
IEEE 802.3x and IEEE 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	
<1W S0-Max (state) 1000BASE-T Active 90 °C (mode) <400mW S0-Typ (state) 100BASE-T Acvite (mode)	• Controller is designed for low power consumption.
IEEE 802.3az - 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 IEEE 802.3az (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]. – APM - Designed to receive a broadcast or unicast packet with an explicit data pattern (Magic Packet) and assert a signal to wake up the system • Advanced Configuration and Power Interface (ACPI) specification v2.0c. ACPI - PCIe power management based wake-up that can generate system wake-up events from a number of sources. • Magic Packet wake-up enable with unique MAC address.
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 /PHY enable the device to enter 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 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	
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	• 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 the 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 into VF or queues.
VM to VM Packet forwarding (Packet Loopback)	• On-chip VM-VM traffic enables PCIe speeds 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 in can receive.
Per-pool statistics, offloads, and jumbo frames support	• Each Queue Pair or Pool has its own statistics, off-loads and jumbo support options.

Features	Description
I/O Virtualization Features (continued)	
Per-pool statistics, offloads, and jumbo frames 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) ad 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 the 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 supports for IP/TCP and IP/UDP receives 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 the device to L2 packets according to the requested MSS.
Interrupt throttling control	• Limits maximum interrupt rate and improves CPU usage.
Legacy and Message Signal Interrupt (MSI)	• 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	• 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 faster and more accurate 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 2.1 TLP Processing Hint Requester	• Provides hints on a per transaction basis to facilitate optimized processing.
Descriptor ring management hardware for Transmit and Receive	• Optimized descriptor fetch and write-back for efficient system memory and PCIe bandwidth usage.
Remote Boot Option Features	
Preboot Execution Environment (PXE) flash interface support	• Enables system boot via the EFI (32-bit and 64-bit). • Flash interface for PXE 2.1 option ROM.
Intel® Ethernet 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 Servies, or UEFI	• Enables networked computer to boot using a program code image supplied bya remote server. • Complies with the PXE, Version 2.1 Specification.
Manageability Features	
Management Component Transport Protocol (MCTP) over SMBus and PCIe	• Used for baseboard management controller (BMC) communication between add-in devices.
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.

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 and installed; low-profile bracket included in package (T2, T4 and F2)	Streamlines installation
Cable Distance	Copper: up to 100m

Specifications

Connector	RJ45
IEEE standard/ network topology	EEE 802.3 (10BASE-T, 100BASE-TX, 1000BASE-T)
Cabling	CAT3 or higher for 10BASE-T operation CAT5 or higher for 100BASE-TX operation CAT5e or higher for 1000BASE-T operation Fiber: MMF 62.5/50 um
Data rate supported per port	10/100/1000 Mbps
Bus Type	PCIe 2.1 (5GT/s)
Bus Width	4-lane PCIe; 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)	Copper: I350T4V2 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	Copper - Speed LED: Not illuminated = 10Mb/s; green = 100Mb/s; amber = 1Gb/s

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)

Product Order Code

Configuration	Cisco Product ID
Quad Port	Server installed: UCSC-PCIE-IRJ45 Spare adapter: UCSC-PCIE-IRJ45=
Cisco Servers Supported*	C220 M6, C240 M6, C220 M5, C240 M5, C480 M5, S3260 M5

*Servers supported as of the date of this publication. For up-to-date server compatibility, please check: https://ucshcltool.cloudapps.cisco.com/public/

Supported Operating Systems

For a complete list of supported network operating systems for Intel® Ethernet Network Adapters visit: intel.com/support/EthernetOS

Printed in USA 0622/ED/123E Please Recycle ♻