Product Brief

Intel® 82576 Gigabit Ethernet Controller

Network Connectivity

Intel® 82576 Gigabit Ethernet Controller

High-performance, Dual-Port Gigabit Network Connectivity for multi-core and virtualized platforms

- High-performing, PCI Express* 10/100/1000 Ethernet connection
- Dual-port, single-chip configuration simplifies designs
- Improves performance and throughput by reducing the I/O overhead due to the software virtualization layer
- Provides near-native performance and data protection with direct connectivity to the virtual machines
- Enhanced support for passthrough traffic to board management controller

The Intelligent Way to Connect

The Intel® 82576 Gigabit Ethernet Controller is a compact component with two fully integrated Gigabit Ethernet Media Access Control (MAC) and physical layer (PHY) ports. This device uses the PCI Express* 2.0 (2.5 GT/s) [x4, x2, and x1] interface. The Intel 82576 Gigabit Ethernet Controller provides two IEEE 802.3* Ethernet interfaces for 1000BASE-T, 100BASE-TX, and 10BASE-T applications. Both ports also integrate a Serializer-Deserializer (SerDes) interface to support 1000BASE-SX/LX/BX (optical fiber), Gigabit Ethernet backplane applications and external PHYs connected either on board or via SFP connectors using SGMII signaling. Intel 82576 Gigabit Ethernet Controller provides hardware assists for auto sensing and switching between Copper and SerDes/SGMII interfaces.

High-Performance Design Features

The Intel 82576 Gigabit Ethernet Controller provides high-performance dual-port gigabit connectivity in a multi-core platform as well as in a virtualized platform. In a multi-core platform it supports different technologies including Intel® QuickData Technology, MSI-X, Low Latency Interrupts and others, that help accelerate the data across the platform thereby improving application response times. In a virtualized environment it supports Intel® Virtualization Technology[1] for Connectivity that helps improve the I/O performance by reducing the I/O overhead on a virtualized platform.

The I/O technologies on a multi-core platform make use of the multiple queues and multiple interrupt vectors available on the Controller. These queues and interrupt vectors help in load balancing the data and the interrupts amongst themselves in order to lower CPU utilization and improve the overall system performance. Depending upon the latency sensitivity of the data, the Low Latency Interrupts feature automatically modulates the time interval between the interrupts, lowering CPU utilization levels.

Intel Virtualization Technology for Connectivity is a suite of hardware assists provided on the I/O Controller that improves the overall system performance by lowering the I/O overhead on a virtualized environment. The technologies include Intel® I/O Acceleration Technology (Intel® I/OAT), Virtual Machine Device Queues and PCI-SIG SR-IOV implementation. Intel I/OAT comprises different stateless offloads (checksum and TCP Segmentation offload), MSI-X and other features to enhance data acceleration across the virtualized platform. Virtual Machine Device Queues help improve data processing by offloading the data sorting functionality from the software virtualization layer to the I/O Controller. PCI-SIG SR-IOV implementation helps direct connectivity to the virtual machines in order to provide them with near-native performance. It also provides data isolation between the virtual machines.

Intel 82576 Gigabit Ethernet Controller also provides security features like IPsec offload and LinkSec.*[1] IPsec offload significantly improves the system performance by delivering near-line-rate throughput with reduced CPU utilization. This performance gain is achieved by offloading the IPsec functionality on to the I/O silicon. IPsec provides protection between the host and destination devices of a communication network. LinkSec is a new IEEE industry-standard feature that helps in providing data protection in the network. The IEEE 802.3ae and IEEE 802.3af protocols help to provide this protection between two devices in a hop-by-hop manner along the communication path.

On-Board Management Features

The Intel 82576 Gigabit Ethernet Controller enables network manageability implementations required by IT personnel for remote control and alerting (IPMI, KVM Redirection, Media Redirection) by sharing the LAN port and providing standard interfaces to a Board Management Controller (BMC). The communication to the BMC is available either through an on-board System Management Bus (SMBus) port or through the DMTF defined NC-SI. The Controller provides filtering capabilities to determine which traffic is forwarded to the BMC. The Controller also supports the new IEEE 1588 Precision Time Control standard. This standard provides a time synch capability that synchronizes the internal clocks according to a network master clock.

Device Configuration

The Intel 82576 Gigabit Ethernet Controller can be configured using an EEPROM, but can also be used in an EEPROM-less configuration. The internal PHYs can be controlled using an internal IEEE 802.3 MDIO register set. External PHYs can be controlled either using an IEEE 802.3 MDIO interface or using a 2-wire interface as defined in the SFP module specification.

The Intel 82576 Gigabit Ethernet Controller package is a 25 mm x 25 mm, 576-pin Flip-Chip Ball Grid Array (FC-BGA). Intel 82576 Gigabit Ethernet Controller is pin compatible with Intel® 82575 Gigabit Ethernet Controller.

Features	Benefits
PCI Express* Features	
PCI Express* 2.0 (2.5GT/s)	• Supports x4/x2/x1 lanes • Supports configurable completion timeout
Vital Product Data (VPD) support	• Allows OEMs to keep proprietary inventory information in the EEPROM of the device
Compatible extensions to PCI power management and ACPI	• Efficient power management
Gigabit MAC/PHY Advanced Features	
Wide, pipelined internal data path architecture	• Low-latency data handling • Superior direct memory access (DMA) transfer-rate performance
Support for transmission and reception of packets up to 9.5 KBytes (Jumbo Frames)	• Enables higher and better throughput of data
IEEE 802.3* auto-negotiation	• Automatic link configuration for speed, duplex, and flow control • Improves performance and reliability
IEEE 802.3* compliant flow-control support with software-controllable pause times and threshold values	• Frame loss from receive overruns reduced • Control over the transmissions of pause frames through software or hardware triggering
Supported cable length is more than 100 meters	• Reliable operation at greater distances
Integrated PHY for 10/100/1000 Mbps	• Smaller footprint, lower power dissipation compared to multi-chip MAC and PHY solutions
IEEE 802.3 PHY compliance and compatibility	• Robust operation over installed base of Category-5 twisted-pair cabling
Built-in cable diagnostics and adjustments for cable faults	• Improved end user troubleshooting • Tolerance of common wiring faults

Features	Benefits
I/O features for Multi-core processor servers	
Intel® QuickData Technology	• DMA Engine: enhances data acceleration across the platform (network, chipset, processor), thereby lowering CPU utilization • Direct Cache Access (DCA): allows the processor to pre-fetch the data from memory, thereby avoiding cache misses and improving service level agreements (SLA) of the applications
MSI-X support	• Allows load balancing of interrupt handling between different CPUs/cores, thereby minimizing the overhead of data interrupts
Low Latency Interrupts	• Based on the sensitivity of the incoming data it can automatically moderate the time intervals between the interrupts
Header split and replication in receive	• Helps the driver to focus on the relevant part of the packet without the need to parse it
Multiple Queues (8 queues/port)	• Network packet handling without waiting or buffer overflow, providing efficient packet prioritization
Tx/Rx IP, SCTP, TCP, and UDP checksum offloading (IPv4, IPv6) capabilities (IPv4, IPv6)	• Lower processor utilization • Checksum and segmentation capability extended to new standard packet type
Tx TCP segmentation offload (IPv4, IPv6)	• Increased throughput and lower processor utilization • Compatible with large send offload feature (in Microsoft Windows* XP)
Receive Side Scaling for Windows environment and Scalable I/O for Linux* environments (IPv4, IPv6, TCP/UDP)	• This technology enables the affinitization of the interrupts to the processor cores in order to improve the CPU utilization rate
IPsec offload	• Offloads IPsec capability on to the silicon hardware instead of the software to significantly improve I/O throughput and CPU utilization rate
LinkSec*[2]	• Layer 2 data protection solution that provides encryption and authentication ability between two individual devices (routers, switches, etc)
IPv6 offloading	• Checksum and segmentation capability extended to new standard packet type
Advanced packet filtering	• 24 exact-matched packets for unicast and multicast frames • 4096-bit hash filter for unicast frames • Lower processor utilization • Promiscuous (unicast and multicast) transfer mode support • Optional filtering of invalid frames
IEEE 802.1q virtual local area network (VLAN) support with VLAN tag insertion, stripping and packet filtering for up to 4096 VLAN tags	• Ability to create multiple VLAN segments
Virtualization Features	
Virtual Machine Device Queues (VMDq) (8 queues/port available for this functionality)	• Offloads the data sorting functionality from the Hypervisor to the silicon, thereby improving data throughput and CPU utilization • 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	• 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 on to the wire and come back in. This improves throughput and CPU utilization • Supports replication of multicast and broadcast data
PC-SIG SR-IOV implementation (8 virtual functions per port)	• Provides an implementation of the PCI-SIG standard for I/O Virtualization. The physical configuration of each port is divided into multiple virtual functions. Each virtual function is assigned to an individual virtual machine directly by bypassing the virtual switch in the Hypervisor, thereby resulting in near-native performance. It is also integrated with Intel® VT[1] for Directed I/O (VT-d) to provide data protection between virtual machines by assigning separate physical address in the memory to each virtual machine.
Manageability Features	
IEEE 1588 Precision Time Control Protocol	• Time synch capability – synchronizes internal clocks according to a network master clock
On-board microcontroller	• Implements pass through manageability via a sideband interface to a Board Management Controller (BMC) via NC-SI or SMBus
DMTF NC-SI passthrough	• Industry standard for BMC interface • Allows fast data rates (up to 100Mb/s full duplex) • Better capabilities (video redirection) • Extended filtering capabilities
SMBus passthrough	• Supports passthrough over the SMBus interface • Data rates of up to 400 KHz • Better capabilities (video redirection) • Allows serial redirection and IPMI traffic redirection to BMC
Advanced filtering capabilities (IPv4, IPv6)	• Supports extended L2, L3 and L4 filtering for traffic routing to BMC • Supports MAC address, VLAN, ARP, IPv4, IPv6, RMCP UDP ports, UDP/TCP ports filtering • Supports flexible header filtering • Allows the BMC to share the MAC address with the host OS
Preboot eXecution Environment (PXE) flash interface support	• Enables system boot up via the LAN (32-bit and 64-bit) • Flash interface for PXE image

Features	Benefits
Manageability Features (continued)	
Simple Network Management Protocol (SNMP) and Remote Network Monitoring (RMON) statistic counters	• Easy system monitoring with industry-standard consoles
SDG 3.0, Wired for Management (WfM) 3.0 and PC2001 compliant	• Remote network management capabilities through Desktop Management Interface (DMI) 2.0 and SNMP software
Wake-on-LAN support	• Packet recognition and wake-up for LAN on motherboard applications without software configuration
iSCSI boot	• Enables system boot up via iSCSI • Provides additional network management capability
MDIO – internal or external management interface	• Enables the software to monitor and control the PHY
Watchdog timer	• Used to give an indication to the manageability firmware or external devices that the chip or the driver is not functioning
Additional Device Features	
Dual Integrated SerDes	• Supports backplane and fiber-based applications as well as copper-based applications via the SGMII interface
SFP (SGMII and 2-wire) interface	• Allows seamless connection to industry-standard SFP pluggable copper and fiber modules
Fiber/Copper auto switch	• Auto detection and switching between copper/fiber interfaces
Four outputs on each port that directly drive LEDs	• Link, speed, and activity indications (10, 100, and 1000 Mbps) on each port
Internal phase-locked loop (PLL) for clock generation can use a 25 MHz crystal	• Lower component count and reduced system cost
JTAG (IEEE 1149.1*) silicon test access port built-in	• Simplified testing using boundary scan • Supports the IDCODE instruction
Characteristics	
Electrical	
Targeted power dissipation (in active link state)	2.4 W
Environmental	
Operating temperature	• 1000BASE-T, 0° to 55° C (with thermal management) • 1000BASE-SX/LX (or SerDes backplane), 0° to 70° C • Storage temperature, 65° C to 140° C
Physical	
Implemented in 90nm complementary metal-oxide semiconductor (CMOS) process	• Offers lowest geometry to minimize power and size while maintaining quality and reliability
Package	• 25 mm x 25 mm 576-pin Flip-Chip Ball Grid Array (FC-BGA) package • Pin compatible with Intel® 82575 Gigabit Ethernet Controller

Order Codes

82576EB (Lead-free) JL82576EB

For more information, contact your Intel sales representative.

To see the full line of Intel Ethernet Controllers, visit www.intel.com/network.

[1] Intel® Virtualization Technology requires a computer system with an enabled Intel® processor, BIOS, virtual machine monitor (VMM) and, for some uses, certain platform software enabled for it. Functionality, performance or other benefits will vary depending on hardware and software configurations and may require a BIOS update. Software applications may not be compatible with all operating systems. Please check with your application vendor.

[2] Feature available on the silicon today and the ecosystem availability in 1H09.

Printed in USA 0608/TAR/OCG/XX/PDF ♻ Please Recycle 320025-001US