

Product Brief

Intel® 82575EB Gigabit Ethernet Controller

Network Connectivity



Intel® 82575EB Gigabit Ethernet Controller

High-performance, Dual-Port Gigabit Network Connectivity for Server and Embedded System Designs

- High-performing, PCI Express* 10/100/1000 Ethernet connection
- Dual-port, single-chip configuration simplifies designs
- Improves flexibility and performance in a virtualized environment in multi-core systems by lowering interrupt overhead and providing sorted multiple data queue paths
- Enhanced support for pass-through traffic to board management controller

The Intelligent Way to Connect

The Intel® 82575EB Gigabit Ethernet Controller is a single, 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 Gbps) [x2 and x4] interface. The Intel 82575EB 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 (optical fiber), Gigabit Ethernet backplane applications, and external PHYs connected either on board or via SFP connectors using SGMII signaling. Intel 82575EB Gigabit Ethernet Controller provides hardware assists for auto sensing and switching between copper and SerDes/SGMII interfaces.

High-Performance Design Features

The Intel 82575EB Gigabit Ethernet Controller for PCI Express is designed for high performance and low memory latency. Wide internal data paths eliminate performance bottlenecks by efficiently handling large address and data words. The controller efficiently handles packets with minimum latency by combining parallel and pipelined logic architectures that are optimized for independent transmit and receive queues. The controller includes advanced interrupt-handling features, including MSI-X support and uses efficient ring-buffer descriptor data structures with up to 64 packet descriptors per queue cached on chip. The controller also allows efficient routing of packets to the target destination in a virtualized environment using multiple hardware queues. Based on the data type being transferred, the interrupts can be handled either in an aggregation mode or sending interrupts immediately upon receipt. A large 48 KByte per port on-chip packet buffer maintains superior performance.



Direct Cache Access enables the I/O device to activate a pre-fetch engine in the CPU that loads the data into the CPU cache ahead of time, thereby eliminating cache misses and reducing CPU loads. In addition, by using hardware acceleration, the controller can offload certain tasks from the host such as transmit and receive checksum calculations for Transmission Control Protocol (TCP), User Datagram Protocol (UDP), and Internet Protocol (IP); header and data splitting of received packets; and TCP segmentation of transmit packets.

Hardware-assisted Virtualization

As IT focuses on reducing costs by deploying virtualization and consolidating servers, this new environment also creates unique challenges for the I/O infrastructure. Recognizing these challenges, Intel has designed smart acceleration capabilities into Ethernet Controllers that optimizes networking performance on virtual servers. The Intel 82575EB Gigabit Ethernet Controller adds Virtual Machine Device queue (VMDq) technology that offloads data sorting and data copying from the virtual machine monitor (VMM) software layer to the hardware. This improves overall throughput and CPU utilization on virtualized servers. VMDq technology also ensures transmit fairness and prevents head-of-line blocking to deliver enhanced latency performance.

On-Board Management Features

The Intel 82575EB 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 host. For low-cost implementation, the internal controller supports ASF 2.0.

Device Configuration

The Intel 82575EB Gigabit Ethernet Controller can be configured using the 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 using either an IEEE 802.3 MDIO interface or using a 2-wire interface as defined in the SFP module specification. Both of the ports support the Wake on LAN feature.

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

Features

Benefits

PCI Express* Features

PCI Express* 2.0 (2.5 Gbps)

- Compatible extensions to PCI power management and ACPI
- Wake on LAN feature supported

- Supports x4/x2 lanes
- Supports configurable completion timeout

- Efficient power management

Gigabit MAC/PHY Advanced Features

Intel® I/O Acceleration Technology (Intel® I/OAT)

Wide, pipelined internal data path architecture

MSI-X support

Mechanism available for reducing interrupts generated from Tx/Rx operations

Low-latency interrupts

Four optimized Transmit (Tx) and Receive (Rx) queues per port

Caches up to 64 packet descriptors per queue

Dual 48 KB configurable Rx and Tx first-in/first-out (FIFO) buffers

Support for transmission and reception of packets up to 9.5 KBytes (Jumbo Frames)

Programmable host memory receive buffers size per queue (1 KByte to 127 KBytes) and cache line size (64 Bytes or 128 Bytes)

Descriptor ring management hardware for Tx/Rx optimized descriptor fetching and write-back mechanisms

IEEE 802.3* auto-negotiation

IEEE 802.3* compliant flow-control support with software-controllable pause times and threshold values

Supported cable length is more than 100 meters

Integrated PHY for 10/100/1000 Mbps (full- and half-duplex)

IEEE 802.3 PHY compliance and compatibility

Built-in cable diagnostics and adjustments for cable faults

- Accelerated TCP I/O for improved CPU utilization

- Low-latency data handling
- Superior direct memory access (DMA) transfer-rate performance

- Minimizes the overhead of interrupts
- Allows load balancing of interrupt handling between different cores/CPUs

- Maximizes system performance and throughput

- Provides the ability to toggle between interrupt aggregation and non-aggregation mode based on the type of data being transferred

- Network packet handling without waiting or buffer overflow
- Efficient packet prioritization

- Efficient use of PCI Express bandwidth

- No external FIFO memory requirements
- FIFO size adjustable to application
- Error detection and correction for FIFO data

- Enables higher and better throughput of data

- Efficient use of PCI Express bandwidth and memory resources

- Simple software programming model
- Efficient use of system memory and PCI Express bandwidth

- Automatic link configuration for speed, duplex, and flow control
- Improves performance and reliability

- Frame loss from receive overruns reduced
- Control over the transmissions of pause frames through software or hardware triggering

- Reliable operation at greater distances

- Smaller footprint, lower power dissipation compared to multi-chip MAC and PHY solutions

- Robust operation over installed base of Category-5 twisted-pair cabling

- Improved end-user troubleshooting
- Tolerance of common wiring faults

Host Offloading Features

VMDq

Direct Cache Access (DCA)

Header split and replication in receive

- Allows the efficient routing of packets to the correct target machine in a virtualized environment using multiple hardware queues

- Enables the I/O device to activate a pre-fetch engine in the CPU that loads the data into the CPU cache ahead of time, before use, eliminating cache misses and reducing CPU load

- Helps the driver to focus on the relevant part of the packet without the need to parse it

Features

Benefits

Host Offloading Features (continued)

Tx/Rx IP, TCP, and UDP checksum offloading (IPv4, IPv6) capabilities	<ul style="list-style-type: none">• Lower processor utilization• Checksum and segmentation capability extended to new standard packet type
Tx TCP segmentation offloading (IPv4, IPv6)	<ul style="list-style-type: none">• Increased throughput and lower processor utilization• Compatible with large send offload feature (in Microsoft Windows* XP)
IPv6 offloading	<ul style="list-style-type: none">• Checksum and segmentation capability extended to new standard packet type
Receive Side Scaling for Windows environment and Scalable I/O for Linux* environments (IPv4, IPv6, TCP/UDP)	<ul style="list-style-type: none">• Multiple Rx queues
Advanced packet filtering	<ul style="list-style-type: none">• 16 exact-matched packets (unicast or multicast)• 4096-bit hash filter for multicast 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	<ul style="list-style-type: none">• Ability to create multiple VLAN segments
Double VLAN	<ul style="list-style-type: none">• Enables a mode where all received and sent packets have at least one VLAN tag in addition to the regular tagging, which may be optionally added. This mode is used for networks where the switches add an additional tag containing switching information.

Manageability Features

On-board microcontroller	<ul style="list-style-type: none">• Implements pass-through manageability via a sideband interface to a Board Management Controller (BMS) via NC-SI or SMBus• Implements ASF 2.0 for low cost, standalone management solution
DMTF NC-SI pass through	<ul style="list-style-type: none">• Industry standard for BMC interface• Allows fast data rates (up to 100 Mb/s full duplex)• Better capabilities (video redirection)• Extended filtering capabilities
SMBus pass through	<ul style="list-style-type: none">• Supports pass through over the SMBus interface• Data rates of up to 400 KHz• Allows serial redirection and IPMI traffic redirection to BMC
Advanced filtering capabilities (IPv4, IPv6)	<ul style="list-style-type: none">• 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
Alerting Standards Format 2.0	<ul style="list-style-type: none">• Standard alerting capability to notify IT of system events
Preboot eXecution Environment (PXE) flash interface support	<ul style="list-style-type: none">• Enables system boot up via the LAN (32 bit and 64 bit)• Flash interface for PXE image
Simple Network Management Protocol (SNMP) and Remote Network Monitoring (RMON) statistic counters	<ul style="list-style-type: none">• Easy system monitoring with industry-standard consoles
SDG 3.0, Wired for Management (WfM) 3.0 and PC2001 compliant	<ul style="list-style-type: none">• Remote network management capabilities through Desktop Management Interface (DMI) 2.0 and SNMP software
Wake on LAN support	<ul style="list-style-type: none">• Packet recognition and wake-up for LAN on motherboard applications without software configuration
iSCSI boot	<ul style="list-style-type: none">• Enables system boot up via iSCSI• Provides additional network management capability
MDIO - internal management interface	<ul style="list-style-type: none">• Enables the software to monitor and control the PHY
Watchdog timer	<ul style="list-style-type: none">• 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	<ul style="list-style-type: none">• Supports backplane and fiber-based applications as well as copper-based applications via the SGMII interface
SFP (SGMII and 2-wire) interface	<ul style="list-style-type: none">• Allows seamless connection to industry-standard SFP pluggable copper and fiber modules
Fiber/copper auto switch	<ul style="list-style-type: none">• Auto detection and switching between copper/fiber interfaces
Four outputs on each port that directly drive LEDs	<ul style="list-style-type: none">• Link 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	<ul style="list-style-type: none">• Lower component count and reduced system cost
JTAG (IEEE 1149.1*) silicon test access port built-in	<ul style="list-style-type: none">• Simplified testing using boundary scan• Supports the IDCOD instruction
No heat sink	<ul style="list-style-type: none">• Lower cost• Better airflow• Better thermal characteristics

Features

Benefits

Characteristics

Electrical

Typical targeted power dissipation (in active link state) • 2.4 W, 0.8 W and 0.3 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

Order Code

82575EB (Lead-free) JL82575EB

For more information, contact your Intel sales representative.

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

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 <http://www.intel.com/>.

Copyright © 2007 Intel Corporation. All rights reserved.

Intel, the Intel logo, Intel. Leap ahead., and the Intel. Leap ahead. logo 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

0707/TAR/OCG/PP/800

 Please Recycle

317554-001US

