intel®
Leap ahead

# 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

| Features | Benefits |
|---|---|
| **PCI Express\* Features** | |
| PCI Express\* 2.0 (2.5 Gbps) | ▪ Supports x4/x2 lanes<br>▪ Supports configurable completion timeout |
| ▪ Compatible extensions to PCI power management and ACPI<br>▪ Wake on LAN feature supported | ▪ Efficient power management |
| **Gigabit MAC/PHY Advanced Features** | |
| Intel® I/O Acceleration Technology (Intel® I/OAT) | ▪ Accelerated TCP I/O for improved CPU utilization |
| Wide, pipelined internal data path architecture | ▪ Low-latency data handling<br>▪ Superior direct memory access (DMA) transfer-rate performance |
| MSI-X support | ▪ Minimizes the overhead of interrupts<br>▪ Allows load balancing of interrupt handling between different cores/CPUs |
| Mechanism available for reducing interrupts generated from Tx/Rx operations | ▪ Maximizes system performance and throughput |
| Low-latency interrupts | ▪ Provides the ability to toggle between interrupt aggregation and non-aggregation mode based on the type of data being transferred |
| Four optimized Transmit (Tx) and Receive (Rx) queues per port | ▪ Network packet handling without waiting or buffer overflow<br>▪ Efficient packet prioritization |
| Caches up to 64 packet descriptors per queue | ▪ Efficient use of PCI Express bandwidth |
| Dual 48 KB configurable Rx and Tx first-in/first-out (FIFO) buffers | ▪ No external FIFO memory requirements<br>▪ FIFO size adjustable to application<br>▪ Error detection and correction for FIFO data |
| Support for transmission and reception of packets up to 9.5 KBytes (Jumbo Frames) | ▪ Enables higher and better throughput of data |
| Programmable host memory receive buffers size per queue (1 KByte to 127 KBytes) and cache line size (64 Bytes or 128 Bytes) | ▪ Efficient use of PCI Express bandwidth and memory resources |
| Descriptor ring management hardware for Tx/Rx optimized descriptor fetching and write-back mechanisms | ▪ Simple software programming model<br>▪ Efficient use of system memory and PCI Express bandwidth |
| IEEE 802.3\* auto-negotiation | ▪ Automatic link configuration for speed, duplex, and flow control<br>▪ 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<br>▪ 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 (full- and half-duplex) | ▪ 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<br>▪ Tolerance of common wiring faults |
| **Host Offloading Features** | |
| VMDq | ▪ Allows the efficient routing of packets to the correct target machine in a virtualized environment using multiple hardware queues |
| Direct Cache Access (DCA) | ▪ 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 |
| Header split and replication in receive | ▪ Helps the driver to focus on the relevant part of the packet without the need to parse it |

# Features

# Benefits

### Host Offloading Features (continued)

| Features | Benefits |
|---|---|
| Tx/Rx IP, TCP, and UDP checksum offloading (IPv4, IPv6) capabilities | • Lower processor utilization<br>• Checksum and segmentation capability extended to new standard packet type |
| Tx TCP segmentation offloading (IPv4, IPv6) | • Increased throughput and lower processor utilization<br>• Compatible with large send offload feature (in Microsoft Windows* XP) |
| IPv6 offloading | • 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) | • Multiple Rx queues |
| Advanced packet filtering | • 16 exact-matched packets (unicast or multicast)<br>• 4096-bit hash filter for multicast frames<br>• Lower processor utilization<br>• Promiscuous (unicast and multicast) transfer mode support<br>• 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 |
| Double VLAN | • 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

| Features | Benefits |
|---|---|
| On-board microcontroller | • Implements pass-through manageability via a sideband interface to a Board Management Controller (BMS) via NC-SI or SMBus<br>• Implements ASF 2.0 for low cost, standalone management solution |
| DMTF NC-SI pass through | • Industry standard for BMC interface<br>• Allows fast data rates (up to 100 Mb/s full duplex)<br>• Better capabilities (video redirection)<br>• Extended filtering capabilities |
| SMBus pass through | • Supports pass through over the SMBus interface<br>• Data rates of up to 400 KHz<br>• 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<br>• Supports MAC address, VLAN, ARP, IPv4, IPv6, RMCP UDP ports, UDP/TCP ports filtering<br>• Supports flexible header filtering<br>• Allows the BMC to share the MAC address with the host OS |
| Alerting Standards Format 2.0 | • Standard alerting capability to notify IT of system events |
| Preboot eXecution Environment (PXE) flash interface support | • Enables system boot up via the LAN (32 bit and 64 bit)<br>• Flash interface for PXE image |
| 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<br>• Provides additional network management capability |
| MDIO – internal 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

| Features | Benefits |
|---|---|
| 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 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<br>• Supports the IDCODE instruction |
| No heat sink | • Lower cost<br>• Better airflow<br>• 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)<br>▪ 1000BASE-SX/LX (or SerDes backplane), 0° to 70° C<br>▪ 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