


Product Brief

Intel® 82599 10 Gigabit Ethernet Controller

Network Connectivity



Intel® 82599 10 Gigabit Ethernet Controller

Transforming the data center through a reliable and unified 10GbE network

- PCI Express* 2.0, dual port 10 Gigabit Ethernet controller along with Intel® Xeon® processor 5500 series servers deliver the performance that demanding applications require.
- Industry-leading product innovations for I/O virtualization that are ideal for the dynamic data center.
- Unified networking support simplifies the network infrastructure by allowing LAN, SAN, and IPC traffic to share the same data center fabric.
- Energy-efficient design with integrated SFI/KR PHYs that creates new opportunities for Blades and LAN-on-motherboard (LOM) solutions.
- High-volume stable architecture with broad operating system support.



- Advanced features such as Receive Side Coalescing, Intel® Ethernet Flow Director, and support for PCI Express* 2.0 help deliver new levels of 10GbE performance
- Intel® Virtualization Technology¹ for Connectivity (Intel® VT-c) delivers industry-leading I/O virtualization capabilities that provide flexibility, scalability and mobility
- Advanced unified networking capabilities, including support for lossless Ethernet, iSCSI acceleration, NAS, and Fibre Channel over Ethernet (FCoE) offloads help consolidate data center networks onto a single Ethernet fabric

An Intelligent Solution

A number of trends are driving change in the data center. Information growth continues unabated, server virtualization requires multi-core servers with higher bandwidth and networked storage, and power and cooling remains a significant challenge. These trends are driving IT managers to look for new, optimized solutions that support the dynamic data center. From a network perspective, IT managers recognize how the trend towards 10 Gigabit Ethernet (10GbE) will benefit them; they want to consolidate multiple GbE links into a single 10GbE network to lower cost and complexity.

Intel's third-generation 10GbE controller, the Intel® 82599 10 Gigabit Ethernet controller continues to build on the innovative trends set by its predecessor and pushes the envelope even further.

The Intel 82599 10 Gigabit Ethernet controller along with Intel® Xeon® processor 5500^A series-based servers deliver unmatched performance scalability that signals a new paradigm in the data center – a paradigm of intelligent solutions that help manage growth while controlling costs; a paradigm of advanced virtualization features; and a paradigm of a simple, cost-effective, unified network.

Best Choice for Virtualization

The Intel 82599 10 Gigabit Ethernet controller includes Intel® Virtualization Technology for Connectivity (Intel VT-c) to deliver outstanding performance in virtualized server environments. Intel VT-c includes hardware optimizations that help reduce I/O bottlenecks and improve the overall server performance. These technologies are Virtual Machine Device Queues² (VMDq) and

Virtual Machine Direct Connect (VMDc). VMDq improves data processing by offloading the sorting and queuing functionality to the I/O controller from the VMM. VMDc provides direct connectivity to the VMs to deliver near-native performance and VM scalability. VMDc also provides flexibility with mobility by enabling VM migration between physical servers. VMDc is based on the industry-standard PCI-SIG SR-IOV (Single Root I/O Virtualization).

Unified Networking

Intel 82599 10 Gigabit Ethernet controller reduces cost and complexity of the data center by combining LAN and SAN traffic onto a single Ethernet fabric. Customers can use iSCSI, NAS or FCoE to carry storage traffic over Ethernet. In order to meet SAN requirements for guaranteed packet delivery, the controller implements capabilities such as enhanced transmission selecting and priority flow control. The controller accelerates iSCSI traffic by implementing key stateless offloads such as TCP segmentation offload (TSO) and Receive Side Coalescing (RSC). It also supports the trusted iSCSI initiators from Microsoft, Linux,* and VMware operating systems and provides a robust iSCSI remote boot implementation. Further, the Intel 82599 10 Gigabit Ethernet controller delivers a high performing FCoE solution that offloads the main data paths for I/O read and write commands. It also greatly reduces CPU processing on FCoE receive traffic by eliminating a data copy through Direct Data placement implementation.

Integrated Solution for LAN on Motherboard (LOM)

The Intel 82599 10 Gigabit Ethernet controller is a single-chip, dual-port 10GbE implementation in a 25x25 mm package. It reduces BOM cost and design complexity by integrating serial 10GbE PHYs and provides both SFI and KR interfaces. The device is designed for high performance and lower memory latency. Wide internal data paths eliminate performance bottlenecks by efficiently handling large address and data words. The controller also includes advanced interrupt-handling features and uses efficient ring-buffer descriptor data structures, with up to 64 packet descriptors. A large on-chip packet buffer maintains superior performance. The controller enables network manageability implementations required by IT personnel for remote control and alerting. The communication to the Board Management Controller (BMC) is available either through an on-board System Management BUS (SMBus) port or through the DMTF-defined NC-SI.

With industry-leading power consumption, a small footprint, and integrated PHYs, the controller is ideally suited for Server Blades, LAN on motherboard (LOM), NIC, and Mezzanine card implementations. The advanced features of the Intel 82599 10 Gigabit Ethernet controller along with the Intel Xeon processor 5500 series enable customers to scale volume servers to fully utilize and scale to 10GbE capacity.

Features

Host Interface Features

PCI Express* 2.0 (5 GT/s)

- Supports x8, x4, x2, x1 lanes
- Supports extended error reporting and completion timeout control

Compatible extensions to PCI power management and ACPI

- Efficient power management

End to End CRC (ECRC)

- Higher reliability on PCI bus traffic

Network Interface Features

XAUI, KX/KX4, BX, CX4

- Multiple interfaces for design flexibility

10GBASE-KR

- Serial 10GbE backplane interface for simpler blade implementation
- Support for Auto-negotiation and PCS layer

SFP+ MSA (SFI)

- Native support for SFI interface
- Saves BOM cost and reduces design complexity by integrating XAUI to SFI PHYs

NC-SI Interface

- Management interfaces for pass-through traffic to and from manageability controller

100 Mbps/1 Gbps/10 Gbps speeds

- Triple-speed support for backward-compatible implementations

10 Gigabit MAC Advanced Features

Dual configurable first-in/first-out (FIFO) buffers for each port: 160 KB Transmit (Tx) and 512 KB Receive (Rx)

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

Support for transmission and reception of packets up to 15.5 KBytes (Jumbo Frames) in basic mode and 9.5 Kbytes packets when DCB or Virtualization is enabled.

- Enables higher and better throughput of data

Programmable host memory receive buffer per queue (1 KByte to 16 KBytes) and cache line size (64 Bytes to 128 Bytes)

- Efficient use of PCI Express bandwidth

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

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

Features

Benefits

10 Gigabit MAC Advanced Features

Auto-negotiation support as defined in IEEE 802.3ap clause 73	<ul style="list-style-type: none">Provides automatic configuration between 1000BASE-KX, 10GBASE-KX4, and 10GBASE-KR modesImproves performance and reliability
IEEE 802.3*-compliant flow-control support with software-controllable pause times and threshold values	<ul style="list-style-type: none">Frame loss reduced from receive overrunsControl over the transmissions of pause frames through software or hardware triggering
Integrated LinkSec security engines (IEEE 802.1ae): AES-GCM 128bit; Encryption + Authentication; One SC x 2 SA per port. Replay Protection with Zero Window	<ul style="list-style-type: none">Provides end-to-end Layer 2 data protection between the host and the destination.It provides encryption and authentication at every hop along the network (i.e., servers, switches and clients)
Integrated IPsec security engines: AES GCM 128 bit; AH or ESP encapsulation; IPv4 and IPv6 (no option or extended headers)	<ul style="list-style-type: none">Provides end-to-end Layer 3 data protection between the host and the destinationIt offloads the encryption engine to provide line rate throughput in an IPsec environment
Time Sync (IEEE 1588, 802.1as)	<ul style="list-style-type: none">Lets networked Ethernet equipment synchronize internal clocks according to a network master clockEndpoint can then acquire an accurate estimate of the master time by compensating for link latency

LAN Performance Features

Tx/Rx IP, TCP, and UDP checksum offloading (IPv4, IPv6) capabilities (IPv4, IPv6)	<ul style="list-style-type: none">Lower processor utilizationChecksum and segmentation capability extended to new standard packet type
Tx TCP segmentation offload (IPv4, IPv6)	<ul style="list-style-type: none">Increased throughput and lower processor utilization
IPv6 offloading	<ul style="list-style-type: none">Checksum and segmentation capability extended to new standard packet type
TSO Interleaving for Reduced Latency	
TCP Receive Side Coalescing (RSC) for up to 32 flows/port	<ul style="list-style-type: none">Lowers CPU utilization by reducing overhead
MSI-X support	<ul style="list-style-type: none">Minimizes the overhead of interruptsAllows load balancing of interrupt handling between different cores/CPU's
Mechanism available for reducing interrupts generated from Tx/Rx operations	<ul style="list-style-type: none">Maximizes system performance and throughput
Low latency interrupts	<ul style="list-style-type: none">Provides the ability to toggle between interrupt aggregation and non-aggregation mode based on the type of data being transferred
Receive Side Scaling for Windows environments and Scalable I/O for Linux* environments (IPv4, IPv6, TCP/UDP)	<ul style="list-style-type: none">Multiple Rx queues
Intel® Ethernet Flow Director Supports advanced filters that direct receive packets by their flows to different queues	<ul style="list-style-type: none">Enables tight control on routing a flow in the platform. Matches flows and CPU cores for flow affinitySupports multiple parameters for flexible flow classification and load balancing
Optimized queues: 128 Transmit (Tx) and Receive (Rx)	<ul style="list-style-type: none">Network packet handling without waiting or buffer overflowEfficient packet prioritization
Header split and replication in receive	<ul style="list-style-type: none">Helps the driver to focus on the relevant part of the packet without the need to parse it
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">Used for systems where the switches add an additional tag containing switching information

Virtualization Features

Virtual Machine Data Queue ² (VMDq)	<ul style="list-style-type: none">Allows the efficient routing of packets to the correct target virtual machine in a virtualized environment using multiple hardware queues
Virtual Machine Direct Connect (VMDc)	<ul style="list-style-type: none">Enables direct I/O connectivity to the virtual machines bypassing the virtual switch in the VMMEnables VM migration capability across two physical serversBased on industry-standard PCI-SIG SR-IOV
Packet Loopback	<ul style="list-style-type: none">Enables the I/O silicon to route the data packets between the VMs within the same server instead of having to go out on the wire and come back to the same serverLowers CPU utilization and improves latency

Unified Networking Features

Enhanced Transmission Selection (draft IEEE 802.1az)	<ul style="list-style-type: none">Resource allocation per virtual link to provide differentiation among different traffic types (LAN, SAN, and IPC)
Priority Flow Control (draft IEEE 802.1Qbb)	<ul style="list-style-type: none">Enables finer grain control of traffic for virtual links associated with Priority group. Enables "no-drop" behavior on Ethernet for critical storage traffic
FCoE Transmit Segmentation	<ul style="list-style-type: none">Transmit segmentation enables the FCoE initiator to transmit multiple FCoE packets up to a complete FC sequence with a single header in host memory (single instruction), thus reducing CPU overhead

Features

Benefits

Unified Networking Features

FCoE Tx/Rx CRC offload	<ul style="list-style-type: none">▪ Offloads receive FC CRC integrity check while tracking the CRC bytes and FC padding bytes
FCoE Coalescing and Direct data placement	<ul style="list-style-type: none">▪ Saves CPU cycles by reducing the data copy and also minimizes CPU processing by posting only the packet's headers that are required for software
iSCSI Acceleration	<ul style="list-style-type: none">▪ Lower processor utilization using TCP checksum offloading▪ Increased throughput through TCP Segmentation Offload▪ TCP Receive-Side Coalescing
iSCSI boot	<ul style="list-style-type: none">▪ Enables system boot up via iSCSI▪ Provides additional network management capability

Manageability Features

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▪ Supports 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
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
Wake-on-LAN support	<ul style="list-style-type: none">▪ Packet recognition and wake-up for LAN on motherboard applications without software configuration
MDIO - internal management interface	<ul style="list-style-type: none">▪ Enables the MAC and software to monitor and control the PHY

Additional Device Features

Four outputs on each port that directly drive LEDs	<ul style="list-style-type: none">▪ Link and activity indications on each port
JTAG (IEEE 1149.1*) test access port built-in silicon	<ul style="list-style-type: none">▪ Simplified testing using boundary scan▪ Supports the IDCODE instruction

Characteristics

<i>Electrical</i> Typical targeted power dissipation	<ul style="list-style-type: none">▪ 5.1W (10GBase-KX4)▪ 5.5W (10GBase-KR)
<i>Environmental</i> Operating temperature	<ul style="list-style-type: none">▪ 0° to 70° C (with thermal management)
<i>Physical</i> Implemented in 65nm complementary metal-oxide semiconductor (CMOS) process	<ul style="list-style-type: none">▪ Offers lowest geometry to minimize power and size while maintaining quality and reliability
Package	<ul style="list-style-type: none">▪ 25 mm x 25 mm 576-pin Flip-Chip Ball Grid Array (FC-BGA) package

Order Codes

82599EB (Base SKU) - JL82599EB
82599ES (Serial SKU) - JF82599ES

To see the full line of Intel Ethernet Controllers, visit www.intel.com/network/connectivity
For more information, contact your Intel sales representative.

^Δ Intel processor numbers are not a measure of performance. Processor numbers differentiate features within each processor family, not across different processor families. See www.intel.com/products/processor_number for details.

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 VMDq requires a virtualization operating system that supports VMDq.

INFORMATION IN THIS DOCUMENT IS PROVIDED IN CONNECTION WITH INTEL® PRODUCTS. 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. INTEL MAY MAKE CHANGES TO SPECIFICATIONS, PRODUCT DESCRIPTIONS, AND PLANS AT ANY TIME, WITHOUT NOTICE.

Copyright © 2009 Intel Corporation. All rights reserved.

Intel, the Intel logo, and Xeon are trademarks of Intel Corporation in the U.S. and other countries.

* Other names and brands may be claimed as the property of others.

