


Product Brief

Intel® Gigabit ET, ET2, and EF Multi-Port Server

Adapters


Network Connectivity



Intel® Gigabit ET, ET2, and EF Multi-Port Server Adapters

Dual- and quad-port Gigabit Ethernet server adapters designed for multi-core processors and optimized for virtualization

- High-performing, 10/100/1000 Ethernet connection
- Reliable and proven Gigabit Ethernet technology from Intel Corporation
- Scalable PCI Express* interface provides dedicated I/O bandwidth for I/O-intensive networking applications
- Optimized for virtualized environments
- Flexibility with iSCSI Boot and choice of dual- and quad-port adapters in both fiber and copper



The Intel® Gigabit ET, ET2, and EF Multi-Port Server Adapters are Intel's third generation of PCIe GbE network adapters. Built with the Intel® 82576 Gigabit Ethernet Controller, these new adapters showcase the next evolution in GbE networking features for the enterprise network and data center. These features include support for multi-core processors and optimization for server virtualization.

Designed for Multi-Core Processors

These dual- and quad-port adapters provide high-performing, multi-port Gigabit connectivity in a multi-core platform as well as in a virtualized environment. In a multi-core platform, the adapters support different technologies such as multiple queues, receive-side scaling, MSI-X, and Low Latency Interrupts, that help in accelerating the data across the platform, thereby improving application response times.

The I/O technologies on a multi-core platform make use of the multiple queues and multiple interrupt vectors available on the network controller. These queues and interrupt vectors help in load balancing the data and interrupts amongst themselves in order to lower the load on the processors and improve overall system performance. For example, depending upon the latency sensitivity of the data, the low level latency interrupts feature can bypass the time interval for specific TCP ports or for flagged packets to give certain types of data streams the least amount of latency to the application.

Optimized for Virtualization

The Intel® Gigabit ET, ET2, and EF Multi-Port Server Adapters showcase the latest virtualization technology called Intel® Virtualization Technology for Connectivity (Intel® VT for Connectivity). Intel VT for Connectivity is a suite of hardware assists that improve overall system performance by lowering the I/O overhead in a virtualized environment. This optimizes CPU usage, reduces system latency, and improves I/O throughput. Intel VT for Connectivity includes:

- Virtual Machine Device Queues (VMDq)
- Intel® I/O Acceleration Technology¹ (Intel® I/OAT)

Use of multi-port adapters in a virtualized environment is very important because of the need to provide redundancy and data connectivity for the applications/workloads in the virtual machines. Due to slot limitations and the need for redundancy and data connectivity, it is recommended that a virtualized physical server needs at least six GbE ports to satisfy the I/O requirement demands.

Virtual Machine Device queues (VMDq)

VMDq reduces I/O overhead created by the hypervisor in a virtualized server by performing data sorting and coalescing in the network silicon.² VMDq technology makes use of multiple queues in the network controller. As data packets enter the network adapter, they are sorted, and packets traveling to the same destination (or virtual machine) get grouped together in a single queue. The packets are then sent to the hypervisor, which directs them to their respective virtual machines. Relieving the hypervisor of packet filtering and sorting improves overall CPU usage and throughput levels.

This generation of PCIe Intel® Gigabit adapters provides improved performance with the next-generation VMDq technology, which includes features such as loop back functionality for inter-VM communication, priority-weighted bandwidth management, and doubling the number of data queues per port from four to eight. It now also supports multicast and broadcast data on a virtualized server.

Intel® I/O Acceleration Technology

Intel I/O Acceleration Technology (Intel I/OAT) is a suite of features that improves data acceleration across the platform, from networking devices to the chipset and processors, which help to improve system performance and application response times. The different features include multiple queues and receive-side scaling, Direct Cache Access (DCA), MSI-X, Low-Latency Interrupts, Receive Side Scaling (RSS), and others. Using multiple queues and receive-side scaling, a DMA engine moves data using the chipset instead of the CPU. DCA enables the adapter to pre-fetch data from the memory cache, thereby avoiding cache misses and improving application response times. MSI-X helps in load-balancing I/O interrupts across multiple processor cores, and Low Latency Interrupts can provide certain data streams a non-modulated path directly to the application. RSS directs the interrupts to a specific processor core based on the application's address.

Single-Root I/O Virtualization (SR-IOV)

For mission-critical applications, where dedicated I/O is required for maximum network performance, users can assign a dedicated virtual function port to a VM. The controller provides direct VM connectivity and data protection across VMs using SR-IOV. SR-IOV technology enables the data to bypass the software virtual switch and provides near-native performance. It assigns either physical or virtual I/O ports to individual VMs directly.

This technology is best suited for applications that demand the highest I/O throughput and lowest-latency performance such as database, storage, and financial applications.

The PCI-SIG SR-IOV capability is a mechanism for devices to advertise their ability to be directly assigned to multiple virtual machines. This technology enables the partitioning of a PCI function into many virtual interfaces for the purpose of sharing the resources of a PCI Express* (PCIe) device in a virtual environment. These virtual interfaces are called Virtual Functions. Each virtual function can support a unique and separate data path for I/O-related functions within the PCI Express hierarchy. Use of SR-IOV with a networking device, for example, allows the bandwidth of a single port (function) to be partitioned into smaller slices that may be allocated to specific VMs, or guests, via a standard interface.

End-to-End Wired Security

The Intel® Gigabit ET, ET2, and EF Multi-Port Server Adapters are Intel's first PCIe adapters to provide authentication and encryption for IPsec and LinkSec. LinkSec is already designed into the

network adapter hardware. These adapters are future-proof and prepared to provide LinkSec functionality when the ecosystem supports this new technology.

IPsec provides data protection between the end-point devices of a network communication session. The IPsec offload feature is designed to offload authentication and encryption of some types of IPsec traffic and still delivers near line-rate throughput and reduced CPU utilization.

LinkSec is an IEEE industry-standard feature that provides data protection in the network. The IEEE 802.3ae and IEEE 802.3af protocols provide hop-to-hop data protection between two network devices in the transaction line between the host and destination. The two network devices must support the LinkSec technology. The network devices could be servers, switches, and routers.

Features

Benefits

General

Intel® 82576 Gigabit Ethernet Controller	<ul style="list-style-type: none"> Industry-leading, energy-efficient design for next-generation Gigabit performance and multi-core processors
Low-profile	<ul style="list-style-type: none"> Enables higher bandwidth and throughput from standard and low-profile PCIe slots and servers
iSCSI remote boot support	<ul style="list-style-type: none"> Provides centralized storage area network (SAN) management at a lower cost than competing iSCSI solutions
Load balancing on multiple CPUs	<ul style="list-style-type: none"> Increases performance on multi-processor systems by efficiently balancing network loads across CPU cores when used with Receive-Side Scaling from Microsoft or Scalable I/O on Linux*
Compatible with x4, x8, and x16 standard and low-profile PCI Express* slots	<ul style="list-style-type: none"> Allows each port to operate without interfering with the other
Multi-port design	<ul style="list-style-type: none"> Enables dual- or quad-port operation in almost any PCI Express server slot, except x1 slots
Support for most network operating systems (NOS)	<ul style="list-style-type: none"> Enables widespread deployment
RoHS-compliant ³	<ul style="list-style-type: none"> Compliant with the European Union directive 2002/95/EC to reduce the use of hazardous materials
Intel® PROSet Utility for Windows* Device Manager	<ul style="list-style-type: none"> Provides point-and-click management of individual adapters, advanced adapter features, connection teaming, and virtual local area network (VLAN) configuration

Features

Benefits

I/O Features for Multi-Core Processor Servers

Multiple queues & receive-side scaling	<ul style="list-style-type: none">▪ DMA Engine: enhances data acceleration across the platform (network, chipset, processor), thereby lowering CPU usage▪ Direct Cache Access (DCA): enables the adapter to pre-fetch the data from memory, thereby avoiding cache misses and improving application response time
MSI-X support	<ul style="list-style-type: none">▪ Minimizes the overhead of interrupts▪ Allows load balancing of interrupt handling between multiple cores/CPU's
Low Latency Interrupts	<ul style="list-style-type: none">▪ Based on the sensitivity of the incoming data it can bypass the automatic moderation of time intervals between the interrupts
Header splits 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
Multiple queues: 8 queues per port	<ul style="list-style-type: none">▪ Network packet handling without waiting or buffer overflow providing efficient packet prioritization
Tx/Rx IP, SCTP, TCP, and UDP checksum offloading (IPv4, IPv6) capabilities	<ul style="list-style-type: none">▪ Lower processor usage▪ Checksum 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 usage▪ Compatible with large send offload feature (in Microsoft Windows* Server OSs)
Receive and Transmit Side Scaling for Windows* environment and Scalable I/O for Linux* environments (IPv4, IPv6, TCP/UDP)	<ul style="list-style-type: none">▪ This technology enables the direction of the interrupts to the processor cores in order to improve the CPU utilization rate
IPsec Offload	<ul style="list-style-type: none">▪ Offloads IPsec capability onto the adapter instead of the software to significantly improve I/O throughput and CPU utilization (for Windows* 2008 Server and Vista*)
LinkSec	<ul style="list-style-type: none">▪ A Layer 2 data protection solution that provides encryption and authentication ability between two individual devices (routers, switches, etc.)▪ These adapters are prepared to provide LinkSec functionality when the ecosystem supports this new technology

Virtualization Features

Virtual Machine Device queues ² (VMDq)	<ul style="list-style-type: none">▪ Offloads the data sorting functionality from the Hypervisor to the network silicon, thereby improving data throughput and CPU usage▪ 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	<ul style="list-style-type: none">▪ 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 to the wire and come back in. This improves throughput and CPU usage.▪ Supports replication of multicast and broadcast data
IPv6 offloading	<ul style="list-style-type: none">▪ Checksum and segmentation capability extended to the new standard packet type
Advanced packet filtering	<ul style="list-style-type: none">▪ 24 exact-matched packets (unicast or multicast)▪ 4096-bit hash filter for unicast and multicast frames▪ Lower processor usage▪ Promiscuous (unicast and multicast) transfer mode support▪ Optional filtering of invalid frames
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
PCI-SIG SR-IOV Implementation (8 virtual functions per port) <small>Note: Requires a virtualization operating system and server platform that supports SR-IOV and VT-d.</small>	<ul style="list-style-type: none">▪ Provides an implementation of the PCI-SIG standard for I/O Virtualization. The physical configuration of each port is divided into multiple virtual ports. Each virtual port is assigned to an individual virtual machine directly by bypassing the virtual switch in the Hypervisor, resulting in near-native performance.

Features

Benefits

Manageability Features

On-board microcontroller	<ul style="list-style-type: none"> Implements pass through manageability via a sideband interface to a Board Management Controller (BMC) via SMBus
Advanced filtering capabilities	<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, and UDP/TCP ports filtering Supports flexible header filtering Enables the BMC to share the MAC address with the host OS
Preboot eXecution Environment (PXE) 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
iSCSI boot	<ul style="list-style-type: none"> Enables system boot up via iSCSI Provides additional network management capability
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
IEEE 1588 precision time control protocol	<ul style="list-style-type: none"> Time synchrony capability—synchronizes internal clocks according to a network master clock

Specifications

General

Product codes	E1G42ET E1G42ETBLK E1G44ET2 E1G44ET2BLK E1G42EF E1G42EFBLK	Intel® Gigabit ET Dual Port Server Adapter (Bulk Pack – Order 5, Get 5) Intel® Gigabit ET2 Quad Port Server Adapter (Bulk Pack – Order 5, Get 5) Intel® Gigabit EF Dual Port Server Adapter (Bulk Pack – Order 5, Get 5)
Connectors	RJ45 (ET Adapters) LC Fiber Optic (EF Adapter)	
IEEE standards/network topology	10BASE-T, 100BASE-T, 1000BASE-T (ET Adapters) 1000BASE-SX (EF Adapter)	
Cabling	Category-5, unshielded twisted pair (UTP) (ET Adapters) Shielded Cable is required for EMI compliance MMF 62.5/50 um (EF Adapter)	

Adapter Product Features

Intel® PROSet Utility	For easy configuration and management
Plug and play specification support	Standard
Intel® I/OAT ¹ including multiple queues & receive-side scaling	<ul style="list-style-type: none">
Ships with full-height bracket installed, low-profile bracket added in package	<ul style="list-style-type: none">
Cable distance	100 m in Category-5 for 100/1000 Mbps; Category-3 for 10 Mbps (ET Adapters) 275 m at 62.5 um; 550 m at 50 um (EF Adapter)
Receive Side Scaling	<ul style="list-style-type: none">
Direct Cache Access (DCA)	The I/O device activates 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

Specifications (continued)

Network Operating Systems (NOS) Software Support

Operating System	IA32	x64	IPF
Windows* Vista* SP1	▪	▪	N/A
Windows Server* 2003 SP2	▪	▪	▪
Windows* Unified Storage Solution 2003	▪	▪	▪
Windows Server* 2008	▪	▪	▪
Linux* Stable Kernel version 2.6	▪	▪	▪
Linux* RHEL 4	▪	▪	▪
Linux* RHEL 5	▪	▪	▪
Linux* SLES 9	▪	▪	▪
Linux* SLES 10	▪	▪	▪
FreeBSD* 7.0	▪	▪	▪
UEFI* 1.1	▪	▪	▪
VMware ESX* 3.x	▪	▪	▪

Intel Backing

Limited lifetime warranty

90-day, money-back guarantee (U.S. and Canada)

Advanced Software Features

Adapter fault tolerance (AFT)

Switch fault tolerance (SFT)

Adaptive load balancing (ALB)

Teaming support

IEEE 802.3ad (link aggregation control protocol)

Test switch configuration	Tested with major switch original equipment manufacturers (OEMs)
---------------------------	------------------------------------------------------------------

PCIe Hot Plug* Active peripheral component interconnect (PCI)

IEEE 802.1Q* VLANs

IEEE 1588 Precision Time Control Protocol	Time synch capability – synchronizes internal clocks according to a network master clock
-------------------------------------------	------------------------------------------------------------------------------------------

IEEE 802.3 2005* flow control support

Tx/Rx IP, TCP, and UDP checksum offloading (IPv4, IPv6) capabilities (Transmission control protocol (TCP), user datagram protocol (UDP), Internet protocol (IP))

IEEE 802.1p*

TCP segmentation/large send offload

MSI-X supports Multiple Independent Queues

Interrupt moderation

IPv6 offloading	Checksum and segmentation capability extended to new standard packet type
-----------------	---------------------------------------------------------------------------

Technical Features

Data rate supported per port	10/100/1000
Bus type PCI Express	2.0 (2.5 GT/s)
Bus width	4-lane PCI Express, operable in x4, x8 and x16 slots
Interrupt levels	INTA, MSI, MSI-X
Hardware certifications	FCC B, UL, CE, VCCI, BSMI, CTICK, MIC
Controller-processor	Intel® 82576
Typical power consumption	E1G42ET: 2.9W E1G44ET2: 8.4W E1G42EF: 2.2W
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
LED indicators	LINK (solid) and ACTIVITY (blinking)

Physical Dimensions

E1G42ET and E1G42EF	
Length	16.74 cm (6.59 in)
Width	6.81 cm (2.681 in)
E1G44ET2	
Length	16.74 cm (6.59 in)
Width	6.94 cm (2.733 in)
Full-height end bracket	12.00 cm (4.725 in)
Low-profile end bracket	7.92 cm (3.12 in)

On-Board Management Features

The Intel® Gigabit ET, ET2, and EF Multi-Port Server Adapters enable network manageability implementations required by information technology 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 through an on-board System Management Bus (SMBus) port. The adapter provides filtering capabilities to determine which traffic is forwarded to the host.

Companion Products

Consider these Intel products in your server and network planning:

- Intel® 10 Gigabit Server Adapters
 - Copper or fiber-optic network connectivity, up to two ports per card
- Intel® PRO/1000 Server Adapters
 - Copper or fiber-optic network connectivity, up to four ports per card
 - Solutions for PCI Express, PCI-X,* and PCI interfaces
- Intel® PRO/1000 Desktop Adapters for PCI Express and PCI interfaces
- Other Intel® PRO Desktop and Server Adapters
- Intel® Xeon® processors
- Intel® Server Boards

To see the full line of Intel Network Adapters for PCI Express, please visit:
www.intel.com/go/ethernet

¹ Intel® I/O Acceleration Technology (Intel® I/OAT) requires an operating system that supports multiple queues and receive-side scaling.

² VMDq requires a virtualization operating system that supports VMDq.

³ Lead and other materials banned in EU RoHS Directive are either (1) below all applicable substance thresholds or (2) an approved exemption applies.

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 © 2009-2011 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.



Mouser Electronics

Authorized Distributor

Click to View Pricing, Inventory, Delivery & Lifecycle Information:

[Intel:](#)

[E1G42ETBLK](#) [E1G42EF](#) [E1G44ET2](#) [E1G42EFBLK](#) [E1G44ET2BLK](#) [E1G42ET](#)



Компания «ЭлектроПласт» предлагает заключение долгосрочных отношений при поставках импортных электронных компонентов на взаимовыгодных условиях!

Наши преимущества:

- Оперативные поставки широкого спектра электронных компонентов отечественного и импортного производства напрямую от производителей и с крупнейших мировых складов;
- Поставка более 17-ти миллионов наименований электронных компонентов;
- Поставка сложных, дефицитных, либо снятых с производства позиций;
- Оперативные сроки поставки под заказ (от 5 рабочих дней);
- Экспресс доставка в любую точку России;
- Техническая поддержка проекта, помощь в подборе аналогов, поставка прототипов;
- Система менеджмента качества сертифицирована по Международному стандарту ISO 9001;
- Лицензия ФСБ на осуществление работ с использованием сведений, составляющих государственную тайну;
- Поставка специализированных компонентов (Xilinx, Altera, Analog Devices, Intersil, Interpoint, Microsemi, Aeroflex, Peregrine, Syfer, Eurofarad, Texas Instrument, Miteq, Cobham, E2V, MA-COM, Hittite, Mini-Circuits, General Dynamics и др.);

Помимо этого, одним из направлений компании «ЭлектроПласт» является направление «Источники питания». Мы предлагаем Вам помощь Конструкторского отдела:

- Подбор оптимального решения, техническое обоснование при выборе компонента;
- Подбор аналогов;
- Консультации по применению компонента;
- Поставка образцов и прототипов;
- Техническая поддержка проекта;
- Защита от снятия компонента с производства.



Как с нами связаться

Телефон: 8 (812) 309 58 32 (многоканальный)

Факс: 8 (812) 320-02-42

Электронная почта: org@eplast1.ru

Адрес: 198099, г. Санкт-Петербург, ул. Калинина, дом 2, корпус 4, литера А.