

Simple, Reliable Performance for iSCSI Connectivity

Intel® Ethernet Server Adapters deliver excellent results for SANs, without proprietary components

January 2010
Version 1.0

Simplify iSCSI Connectivity

The broad range of Intel® Ethernet Server Adapters provides the easiest and most cost-effective way to connect servers to iSCSI SANs. Intel Ethernet Server Adapters use the native iSCSI initiators integrated in the operating system, are optimized for virtualized environments, and include hardware acceleration features across the entire product line.

EXECUTIVE SUMMARY

Intel® Ethernet Server Adapters offer compelling value to IT organizations because they are designed to meet the growing requirements for server connectivity to storage area networks (SANs). This paper examines some of the trends that have led to that growth of Internet small computer system interface (iSCSI) SANs and the standards-based technologies that make Intel Ethernet Server Adapters ideally suited to connect servers to them.

Next, the paper demonstrates key performance-oriented features of Intel Ethernet Server Adapters, such as the use of operating system (OS)-native iSCSI initiators and hardware-based acceleration. It then introduces optimizations throughout the hardware and software virtualization stack that complete the picture of Intel® Ethernet as the iSCSI connectivity solution-of-choice for IT.

Finally, the paper introduces the outstanding throughput that Intel has achieved using this hardware without the expense and complexity of proprietary iSCSI offloads. In sum, the paper demonstrates how using Intel Ethernet Server Adapters for iSCSI presents a superior solution over proprietary solutions.

OVERVIEW

Data-driven business operations and regulatory data-retention requirements have driven up data-storage needs among businesses of all sizes, making SANs a standard component of IT infrastructures. In particular, virtualized servers benefit dramatically from SAN functionality, in areas such as virtual machine (VM) storage and live migration.

Alongside traditional Fibre Channel networks, a growing proportion of SAN adopters find iSCSI connectivity compelling, particularly because it uses Ethernet, which is a standardized, widely deployed, and well-understood technology. Using iSCSI to transmit block-level storage commands over Ethernet networks adds benefits in terms of cost and simplicity, and Intel is committed to providing the most advanced solutions in the industry for iSCSI and other unified networking technologies.

Intel is the industry's leading supplier of Ethernet adapters, and the broad product line of Intel Ethernet Server Adapters provides the state of the art in simple, reliable, cost-effective iSCSI connectivity. Optimizations for the native iSCSI support provided by the OS help ensure stability and simplicity by eliminating the need for proprietary hardware and software. At the same time, hardware-based acceleration features speed iSCSI traffic, and assists in the Ethernet controller increase throughput in virtualized servers.

Table of Contents

Executive Summary	1
Overview	1
Intel® Ethernet Server Adapters: Ideal for iSCSI	2
Native OS Support for Dependable Simplicity and Lower Costs	2
Extending the Value of Centralized Data Storage with Remote Boot ...	3
Consistent, Reliable Performance across the Product Family	3
Hardware-based Native iSCSI Initiator Acceleration	4
Optimizations for Multicore Servers	4
Platform Advancements	4
Advanced Support for Virtualization	5
Performance Results	5
Conclusion	7

Intel® Ethernet Server Adapters: Ideal for iSCSI

The ubiquity of Ethernet in the data center means that no new topology needs to be added to the environment to support iSCSI. The network connectivity to support iSCSI SANs can share existing server adapters, use dedicated adapters identical to the ones already in use for local area network (LAN) traffic, or a combination of both. Likewise, adding iSCSI to existing switching infrastructures does not require the addition of new types of equipment. The technologies used for the associated networking, such as IP address configuration and routing, are well understood by administrators.

IT organizations broadly prefer Gigabit Ethernet (GbE) and 10GbE Intel Ethernet Server Adapters for iSCSI because of the advantages they offer:

- **Native OS support for dependable simplicity and lower costs.** Avoiding proprietary initiators, TCP/IP stacks, and management tools reduces the number of vendors involved in iSCSI solutions, lowering complexity to help alleviate the possibility of human error and decrease operational overhead.
- **Consistent, reliable performance.** Using Intel Ethernet Server Adapters in iSCSI solutions combines high throughput, intelligent hardware-based offloads with native iSCSI initiators, and stable hardware while leaving processor cycles available for application workloads.
- **Next-generation virtualization support.** Intel Ethernet Server Adapters implement the hardware assists in Intel® Virtualization Technology for Connectivity (Intel® VT-c) to accelerate I/O in virtualized servers.

Intel provides all of this functionality across its Ethernet adapter product line and at a lower cost than dedicated, proprietary iSCSI Host Bus Adapters (HBAs). Using Intel Ethernet Server

Adapters for iSCSI SAN traffic offers mature, standards-based technology instead of unpredictable, proprietary alternatives.

Native OS Support for Dependable Simplicity and Lower Costs

Intel Ethernet Server Adapters are engineered to use the native iSCSI initiators, TCP/IP, and network stacks built into OSs, including Microsoft Windows*, Linux*, and VMware ESX*. Intel works with OS providers to ensure robust optimized support and ease of use. These native initiators are broadly tested using multiple generations of OSs, storage disk systems, OS configuration tools, and other applications to ensure reliability. That design precept removes the need for proprietary, special-purpose hardware or software that would add complexity to the environment and potentially interfere with management tools and other applications.

Standardizing on Intel Ethernet Server Adapters for iSCSI allows IT to use a single initiator, TCP/IP stack, and set of management tools and IT policies. Easier server provisioning, lower likelihood of human error, and simpler management enable lower capital and operational expenditures. That standardization also allows IT to avoid the disparate sets of tools that result from assembling solutions from a variety of vendors.

In addition, the native OS initiators, such as the Microsoft iSCSI initiator, have consistently delivered advanced features that either are not available from or have limited functionality in proprietary initiators. Examples include multipathing solutions (MPIO), the ability to establish multiple connections per session (MCS), and IPsec. Moreover, native OS initiator support for the CRC-32 Digest instruction set included in the Intel® Xeon® processor 5500 series (discussed later) avoids the transmission of data in a vulnerable state to a proprietary iSCSI offload adapter.

The majority of server iSCSI connections today use the integrated native OS initiators, and the integration of the initiator with the OS greatly simplifies the establishment of an iSCSI connection. For example, these are the simple setup steps required to use iSCSI connections in Microsoft Windows Server* 2008:

1. Install Windows Server 2008.
2. Set up your target(s).
3. From Control Panel, launch the iSCSI Initiator utility.
4. From the Discovery tab, add your target.
5. From the Targets tab, select the target, and then click Log On.
6. After logging on, from the Bound Volumes/Devices tab, bind all volumes.
7. Partition or format the disk as needed.
8. Start using your new network-connected iSCSI volume.

In contrast to that simplicity, offload engines and other proprietary solutions force SAN administrators to rely on hardware vendors to provide special-purpose iSCSI initiators. Many of these unproven initiators require frequent firmware updates and reboots, have unpredictable compatibility across different applications, and are supported on only a limited number of adapter types.

Moreover, most HBAs force the SAN administrator to rely on a proprietary TCP/IP stack, the critical link between the host server and the network. OS vendors have hardened and ensured the standards compliance of their native TCP/IP stacks over many years, making them robust and dependable. As a result, most SAN administrators are rightly reluctant to take the unnecessary risk of moving from trusted, OS-provided TCP/IP stacks to unfamiliar, proprietary ones.

Many management applications are also designed to use the native TCP/IP stack in Windows or Linux. Therefore, offloading to a proprietary stack can defeat certain optimizations that have been made to those management applications or even limit their reach in terms of their ability to manage particular systems. More specifically, issues may arise with adapter teaming, as well as security or monitoring applications. Not surprisingly, support for proprietary software stacks by open source OSs and hypervisors can be particularly limited.

A later section shows how Intel Ethernet Server Adapters use this native OS support and intelligent hardware offloads to deliver outstanding iSCSI performance.

Extending the Value of Centralized Data Storage with Remote Boot

Using Intel® iSCSI Remote Boot, an IT administrator can build an OS image and start the server from a disk maintained on the remote storage array instead of from a locally attached hard disk. Maintaining OS images on the SAN with Intel iSCSI Remote Boot enhances many of the benefits of the SAN, including greater centralized data storage and efficiency, as well as improved backup and disk utilization.

In the past, a network boot or pre-boot execution environment (PXE) server was required to manage the remote boot process. Intel Ethernet Server Adapters store boot code in their local memory and work with native boot support in the OS to attach directly to the target SAN storage unit during the power on and boot stage. After boot, the adapter passes control, authentication, and security credentials to the native OS iSCSI initiator.

This process can be accomplished without requiring additional management servers, further simplifying the topology. This approach is in contrast to that of some

iSCSI offload vendors that provide a proprietary boot feature, which requires administrators to accommodate a specialized, proprietary process that is likely not as mature as the native iSCSI boot provided by most OSs.

Decoupling the OS image from the physical server with Intel iSCSI Remote Boot enables a consolidated virtualized environment that can be centrally managed with the following benefits:

- **More efficient provisioning and maintenance.** Managing OS images from a centralized console is more efficient than doing so on distributed servers.
- **Reduced hardware and associated power requirements.** Decoupling storage from the physical server eliminates the need for RAID cards and local storage.
- **Improved disaster recovery.** Virtual disks can be replicated to another physical location, removing the dependence on systems at the primary site.

Intel has worked with major OS vendors to provide remote boot capabilities with native OS iSCSI initiators. Intel iSCSI Remote Boot is designed to work with most major OSs, including Microsoft Windows Server and most Linux distributions.

Consistent, Reliable Performance across the Product Family

Intel Ethernet Server Adapters provide enterprise-class drivers that work together with the native iSCSI initiators to deliver consistent, outstanding iSCSI performance while reducing processor utilization across GbE and 10GbE products. The breadth of Intel's offerings enables IT organizations to increase performance and lower costs by tailoring the network connectivity solution to their specific needs, assisted by advanced technologies.

Simple, Reliable Performance for iSCSI Connectivity

Hardware-based Native iSCSI Initiator Acceleration

Intel Ethernet Server Adapters include a number of hardware features that accelerate iSCSI traffic and improve data processing on multicore processor-based servers. Intelligent hardware offloads deliver iSCSI performance that exceeds the performance demonstrated by leading iSCSI HBAs that offload iSCSI processing to a proprietary engine.

Intel Ethernet Server Adapters deliver specific benefits for processing iSCSI traffic, as summarized in the following table:

Features	Benefits
Receive-side coalescing aggregates packets on incoming traffic flows to reduce processor interrupts and direct memory access (DMA) memory copies	Reduces processor utilization
TCP segmentation offload breaks up chunks of data to be sent on the network in a hardware-based process instead of in software	Increases throughput and lowers processor utilization
Checksum offload receives the segmented packets from the TCP offload block and calculates the checksum	Increases throughput and lowers processor utilization
Direct memory access enables data copy by the chipset instead of the processor	Moves data more efficiently through the server, improving throughput
Direct cache access allows the Ethernet controller to place data directly into processor cache, reducing cache misses	Improves application response times
Data center bridging enables better traffic prioritization over a single interface, as well as advanced means for shaping traffic on the network to decrease congestion.	Provides the no-drop Ethernet fabric needed for storage traffic

Optimizations for Multicore Servers

Intel Ethernet Server Adapters are designed to seamlessly scale workloads across multicore systems, whereas offload engine-based HBAs are constrained by the limited headroom of the offload engine itself. With the current very fast rate of growth in multicore processor performance, servers have much greater headroom to support the application over the life of the server.

Intel Ethernet Server Adapters also feature key capabilities to load-balance and distribute packet processing across multiple execution cores:

Features	Benefits
Receive-and-transmit queues in the controllers allow network traffic streams to be distributed into queues associated with specific processor cores	Improves processor efficiency
Receive-side scaling directs packets to different queues in the adapter without the need for reordering and selects a processor to handle packet flow	Improves processor efficiency
Extended Message-Signaled Interrupts (MSI-X) allow multiple interrupts to be handled simultaneously and load-balanced across multiple processor cores	Improves processor efficiency, reduces latency
Interrupt moderation reduces the number of interrupts the processor needs to handle, primarily through adapter support for Large Send Offload and Large Receive Offload	Improves processor efficiency, reduces latency

Platform Advancements

10GbE Intel Ethernet Server Adapters support data center bridging (DCB) in hardware to provide the lossless or no-drop Ethernet fabric needed for storage traffic. This set of standards enables better traffic prioritization over a single interface as well as an advanced means for shaping traffic on the network to decrease congestion.

iSCSI traffic also benefits from the CRC32-C instructions supported in the Intel Xeon processor 5500 series and the upcoming next-generation Intel Xeon processor-based server platform for the expandable server segments. Data integrity checking with CRC32 has historically been very processor-intensive, and many IT organizations chose not to use it. Because CRC32-C can be enabled without undue burden on the processor, SAN administrators can use it to improve data integrity while accelerating processing, relative to prior generations. Both Windows Server 2008 and ESX 3.0 and later include support for this feature.

Architectural enhancements on these new Intel® server platforms, including PCI Express* Gen 2, an integrated memory controller, and Intel® QuickPath Interconnect, greatly increase the

scalability of 10GbE. Previous-generation platforms did not scale beyond approximately 17 Gbps because of architectural limitations. Conversely, the architecture of the Intel Xeon 5500 series processor-based server allows it to scale up to nearly 50 Gbps or four GbE ports. This level of throughput scalability will be put to good use by today’s evolving dynamic, virtualized data centers.

Advanced Support for Virtualization

Virtualization drives the need for added SAN connectivity to allow the advantages that originate from remote boot, live migration, and related virtualized usage models. Intel Ethernet Server Adapters offer virtualization optimizations that directly benefit iSCSI implementations.

In conventional virtualization implementations, the hypervisor abstracts the I/O devices and shares those hardware resources among multiple VMs. To route traffic coming from a shared I/O device, a virtual switch (vSwitch) within the hypervisor sorts incoming packets based on the destination VM and then delivers them accordingly. This sorting and grouping done in the hypervisor consumes processor cycles, negatively impacting overall network I/O.

Virtual Machine Device Queues (VMDq) technology, a component of Intel VT-c, relieves the hypervisor of this overhead and accelerates I/O traffic in virtualized servers by performing data sorting in the Ethernet controller. VMDq technology makes use of multiple queues in the network device. With VMDq, data packets are sorted as they enter the network adapter, and packets going to the same destination VM are grouped together in the same queue. The packets are then sent to the vSwitch, which directs them to their respective destinations. VMDq is supported in VMware ESX and Microsoft Hyper-V*.

Performance Results: 1 Million IOPS

Testing in Intel labs achieved one million bi-directional average transmit-and-receive I/O operations per second (IOPS) at block sizes of 512 bytes,¹ without proprietary iSCSI offloads. Shown in Figure 1, this result using 10GbE Intel Ethernet Server Adapters and Intel Xeon processor 5500 series-based servers represents new levels of iSCSI performance.

This level of performance, attained using the native iSCSI initiator in Windows Server 2008 R2, is approximately twice that achieved using previous-generation solutions.²




Figure 1. iSCSI read/write performance on the Intel® 82599 10 Gigabit Ethernet Controller, expressed in terms of IOPS (I/O operations per second).¹

Simple, Reliable Performance for iSCSI Connectivity

In real-world implementations, typical IOPS performance is much lower than these benchmarks and is limited by a number of factors, including application threading, as well as storage target processor and disk speeds. At these lower IOPS levels, processor utilization is typically quite low, thanks in part to recent increases in processor performance.

In addition, Intel Ethernet Server Adapter assists such as receive-side coalescing and TCP segmentation offload help reduce processor utilization. Platform features including receive-side scaling and MSI-X help accelerate throughput.

By contrast, solutions based on proprietary offload engines may show little if any improvement in IOPS on new hardware platforms (although processor utilization may improve).

This characteristic reflects the inherent limitations to throughput placed on the overall solution by proprietary offload engines.

In addition to the outstanding IOPS results shown in Figure 1, the results in Figure 2 show full 10GbE line-rate throughput at I/O sizes of 8 KB to 64 KB, which correspond to the most prevalent real-world cases.¹ At these I/O sizes, processor utilization falls to relatively low levels.

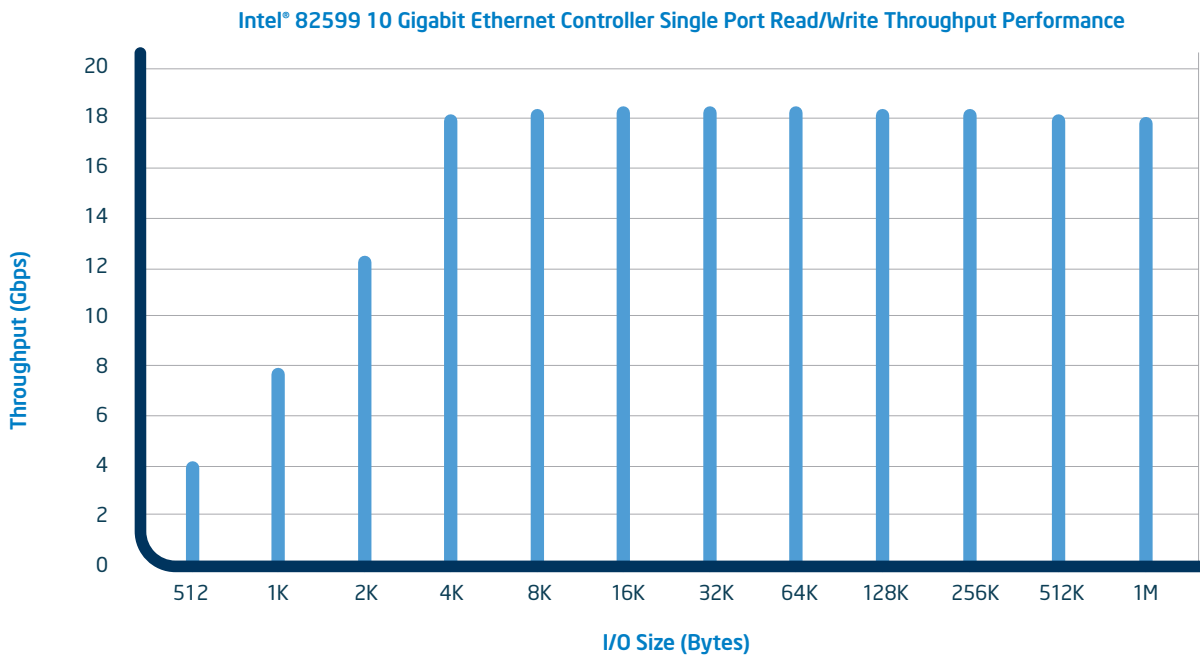


Figure 2. iSCSI read/write performance on the Intel® 82599 10 Gigabit Ethernet Controller, expressed in terms of throughput.¹

Conclusion

Intel remains committed to providing industry-leading support for unified networking. With its support for native initiators, hardware acceleration features, and optimizations for virtualized I/O, Intel Ethernet Server Adapters are the most reliable, cost-effective, and easy-

to-use solutions for iSCSI connectivity. These features enable Intel Ethernet Server Adapters to achieve new levels of iSCSI performance without the need for complicated, proprietary offloads. Intel Ethernet solutions are fully optimized through all layers of the solution stack, from the processor silicon out, including

servers and server adapters. Together with Intel's industry-leading spirit of collaboration with both commercial and open source software providers, the Intel Ethernet Server Adapter product family is broad and deep enough to meet all your connectivity needs.

Learn more: www.intel.com/go/ethernet

Contact your Intel reseller to buy: premierlocator.intel.com

Simple, Reliable Performance for iSCSI Connectivity

¹ **Test Configuration:** Iometer v. 2006.7.27, number of managers = 10 (1/VM), number of workers/manager = 5, (total number of workers) = 50, number of LUNs/worker = 1, number of outstanding I/Os = 90, IO size = 512B-1MB, MTU = 1500.

SUT: Intel® Shadycove, Intel® Xeon® processor W5580 (8 M Cache, 3.20 GHz, 6.40 GT/s Intel® QPI), Memory - 24 GB DDR3, BIOS - S5500.86B.01.10.0038, Windows Server® 2008 R2 x64, 10 VMs with 2 GB allocated RAM.

Network Configuration: Cisco Nexus® 5020, Niantic 82599EB Connected @ 10 Gbps.

Storage: Starwind 4.0 iSCSI Soft Target connected @ 10 Gbps.

² **Test Configuration:** Iometer 2004.07.30, number of managers = 1, total number of workers = 12, number of LUNs/worker = 1, number of outstanding I/Os = 4, IO size = 512B-64K, MTU = 1500.

SUT: Intel® Xeon® processor 5300 series (8 M Cache, 3.0 GHz), Memory - 8 GB 667MHz, Windows Server® 2008 x64.

Network Configuration: Cisco Catalyst® 6509, Intel 82598 10 Gigabit Controller Connected @ 10 Gbps.

Storage: Wintarget® iSCSI Soft Target connected @ 10 Gbps.

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 © 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.

Printed in USA

0110/BY/MESH/PDF

 Please Recycle

322580-001US

