White Paper
Intel® Server Adapters

iSCSI Simplifies Storage Area Networks

Today's medium-size and enterprise-level businesses face ever-increasing demands for more network storage capacity, both for regulatory compliance and for meeting evolving business needs. Of the various network storage architectures available, the storage-area network (SAN) based on the Internet Small Computer System Interface (iSCSI) standard offers the greatest advantages for flexible, scalable storage with easier manageability and lower total cost of ownership (TCO). Software iSCSI initiators, available in most operating systems today, allow you to use any Intel® Gigabit Server Adapter to handle local area network (LAN) traffic, iSCSI SAN traffic, or both. Now, with the addition of iSCSI Remote Boot support on Intel® Gigabit Server Adapters for PCI Express*, you can centralize SAN management without an expensive host bus adapter (HBA) while achieving high performance through the benefits of PCI Express and Intel® I/O Acceleration Technology.

Table of Contents

The Need for More Network Storage

Medium-size and enterprise-level business continue to experience expanding network storage capacity needs. This need for more capacity stems from a variety of causes. E-mail traffic volume, number of attachments, and attachment sizes continue to increase. Collaborative work environments, databases, and multimedia usage further increase traffic and storage requirements. Additionally, most businesses and governmental agencies are converting paper, film, and microfiche archives to digital storage and retrieval. Then there is the issue of regulatory compliance, such as the Sarbanes-Oxley Act, which mandates long-term storage of business records. The compliance impact on network storage needs is further compounded for global enterprises, which must comply with differing regulatory requirements in the various countries of their operation.

Various network storage architectures have evolved to meet expanding capacity needs. Most notably, these architectures include:

- Direct-attached storage (DAS)
- Network-attached storage (NAS)
- Storage-area networks (SANs)

While DAS, NAS, and SAN techniques do meet expanding storage needs in various ways, they each have their own distinct advantages and disadvantages. For example, DAS and NAS have the advantage of using common and familiar storage components and connection technologies, a Small Computer System Interface (SCSI) connection in the case of DAS and Ethernet in the case of NAS. However, DAS and NAS also have limited scalability, and DAS especially is expensive to manage and provision as the enterprise grows. By comparison, SAN is highly scalable, but until recently, SANs required specialized — and sometimes proprietary — hardware and management software. This made SANs expensive, both in terms of initial equipment purchase and installation and the specialized skills necessary for ongoing network storage management and expansion. That began to change in February 2003 when the Internet Small Computer System Interface (iSCSI) standard was ratified, opening the way for simpler, lower-cost implementation and management of SANs.

Additionally, a new storage paradigm, referred to as “unified storage,” is emerging, in which iSCSI SANs will play a prominent role. This white paper discusses the various network storage architectures, their uses, and the advantages provided by iSCSI SANs.

DAS, NAS, and SAN Approaches

DAS, shown on the left in Figure 1, is the simplest form of network storage. It consists of a disk drive attached directly to a server through a SCSI connection. It is important to note that the server-to-disk connection and input/output (I/O) communication protocol is SCSI, not iSCSI.

SCSI, the precursor to iSCSI, is one of the more common means of I/O communication between a server and external storage devices, such as tape drives, compact disk burners, and disk drives. SCSI commands transfer the storage data as blocks, which are the low-level, granular units used by storage devices. This is in contrast to file-oriented transfers commonly used for data transfer in a local area network (LAN).

DAS via SCSI is a relatively simple form of network storage, and it is advantageous for local storage needs such as an e-mail server or departmental server. However, DAS does have its limitations. It can be difficult to manage when its usage is stretched beyond local storage. Scaling up storage capacity is expensive since additional servers must be installed in order to add direct-attached disk drives. Also, since SCSI devices operate over parallel cable to a maximum of 12 meters, there is a distance limitation to DAS.

Additionally, DAS requires time-consuming operating system (OS) provisioning of each server or use of expensive imaging tools. This impedes disaster recovery and business continuity because new OS images must be built to replace or add servers.

NAS, shown at the center of Figure 1, is a file-based storage architecture that attaches directly to the LAN. NAS uses familiar Ethernet technology, and storage traffic is transmitted through files, which is the typical and familiar means of transferring data within a LAN.

Because NAS is a network device on the LAN, information technology (IT) staff can manage NAS with familiar networking tools and only minimum training in storage technology. NAS also has the advantage of being placed anywhere in the LAN up to the distance limitations of the LAN connection (up to 100 meters in CAT-5 copper and up to 10 kilometers in single-mode fiber-optic cabling). Additionally, it is relatively easy to add disk units in NAS to increase network storage capacity. However, because NAS traffic is also LAN traffic, scaling up NAS increases the LAN traffic burden. Also, with NAS the network traffic burden can become exceedingly heavy due to usage spikes or business growth (for example, an insurance company processing claims after a disaster or an online retailer storing orders for a new product line). Usage spikes may seriously decrease LAN performance, and slow storage access can impair application performance.

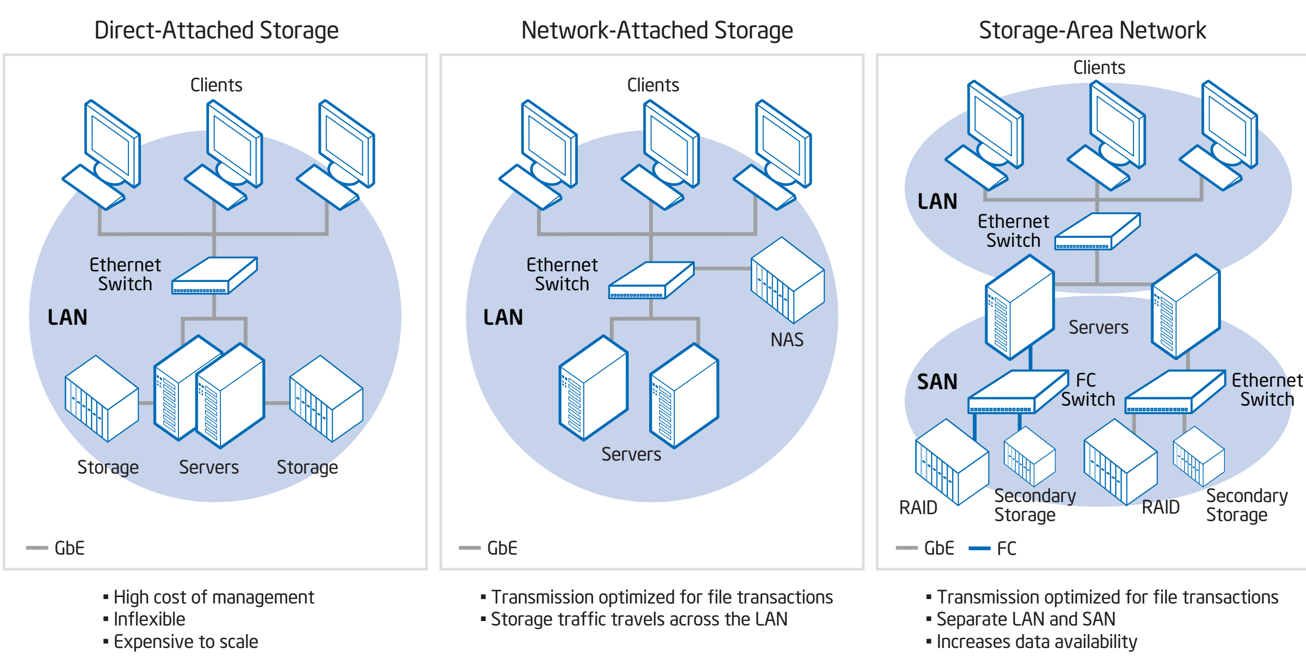

Figure 1. Evolution of network storage. The simplest, but least flexible, storage is direct-attached storage, shown on the left. The most flexible and scalable storage is the storage-area network, shown on the right.

As opposed to NAS, a SAN is a dedicated network of storage devices that transports storage traffic without burdening the enterprise LAN. This LAN and SAN separation is illustrated at the right in Figure 1, where redundant-array-of-independent-disks (RAID) technology is used as well for increased reliability and data accessibility. Also, notice in the Figure 1 Storage-Area Network diagram that two means of SAN connectivity are used – Fibre Channel (FC) connectivity and Gigabit Ethernet (GbE) connectivity through iSCSI.

Regardless of the SAN fabric – FC or GbE – the SAN architecture offers some distinct advantages for network storage. These advantages include increased performance, reliability, scalability, and storage management ease.

The SAN Advantages

From an overall network performance viewpoint, the SAN is a separate network segment or subnet. As a consequence, storage traffic is not traveling across the LAN or burdening the LAN, as is the case with NAS. Additionally, SANs provide the potential for centralized data management. This eliminates redundant file copies, which can rapidly consume disk space as well as cause reconciliation problems. SANs with appropriate management tools can also manage work loads so that high-demand applications with often-used data do not overload some servers while leaving others idle.

SANs are also easily and highly scalable. Growing network storage needs can be met by simply installing more storage devices and SAN resources. Again, as opposed to NAS, adding more storage capacity does not necessarily burden the LAN with additional storage traffic since the LAN and SAN are different network entities.

Additionally, SAN targets have a single IP address for all of the servers. This allows more efficient storage because the SAN controller can virtualize the physical storage location and spread data across the SAN disk array. In contrast, DAS dedicates physical disks to each server, often resulting in too much or too little storage capacity for the task.

SANs are also more flexible in terms of their physical location. SANs, especially iSCSI SANs, can be located long distances from the LAN. This simplifies data-center and storage centralization. It also allows geographical redundancy as part of your planning for quick disaster recovery.

iSCSI Makes SAN More Attractive

Historically, SANs have been implemented using a Fibre Channel (FC) fabric. This requires use of an FC host bus adapter (HBA), which connects the server to devices on the SAN, typically specialized FC switches or FC routers connecting to the various storage devices. The HBA serves the same purpose on the SAN side as a network interface card (NIC) serves on the LAN side. In other words, NICs and HBAs are both server adapters, where a NIC provides Ethernet connectivity and an HBA provides Fibre Channel (or other fabric) connectivity.

Because FC network devices are specialty items, they are more expensive than the corresponding Ethernet devices. Additionally, installing, maintaining, and operating an FC SAN requires IT staff with FC skills and knowledge.

The emergence of iSCSI is simplifying SANs substantially. This is because iSCSI encapsulates SCSI commands into TCP/IP packets and enables block-data transport over IP networks. This allows iSCSI SAN implementation with standard and familiar Ethernet switches, routers, and cabling. Additionally, the Ethernet and TCP/IP compatibility of iSCSI SANs extends their geographic location beyond the theoretical 10-kilometer maximum limit of Fibre Channel to anywhere in the world where Internet connectivity is available. This offers tremendous flexibility in providing iSCSI SAN availability to field offices, branch offices, and stranded servers. It also allows virtually unlimited options in iSCSI SAN location for optimum storage centralization and protection for disaster recovery. At the same time iSCSI and FC SANs can easily coexist, as indicated by the SAN diagram in Figure 1. This allows legacy FC SAN storage capacity to be readily expanded by using lower-cost, familiar Ethernet network components to add an iSCSI SAN to an existing FC SAN.

Additionally, iSCSI promises to play a significant role in the emerging universal storage paradigm. Universal storage, a SAN-like concept, allows external attachment of heterogeneous storage systems, partitioning of storage resources to maximize application quality of service, and moving data across tiers of storage to match business needs and achieve business continuity.

iSCSI Remote Boot Simplifies Further

For all of the iSCSI advantages, adoption of iSCSI SANs stalled initially. The problem was availability, cost, and performance of the iSCSI HBAs initially used for iSCSI SAN implementation. Just like Fibre Channel, an HBA was required for connection of the host, or server, to the SAN. This put iSCSI on par with Fibre Channel in terms of a costly HBA and a non-standard third-party stack or driver for its configuration and operation. (Today, software initiators and remote boot for iSCSI eliminate the need for an iSCSI HBA and allow use of standard NICs.)

Additionally, there were bandwidth and throughput performance concerns for iSCSI HBAs, as compared to Fibre Channel. To address those concerns, off-load techniques were proposed for iSCSI HBAs, the most notable of these being the TCP/IP Off-load Engine (TOE). While TOE improves throughput by freeing the host from TCP/IP processing, TOE also increases HBA cost and complexity. TOE usage may also require modification of the operating system and certain applications to make them TOE aware, which adds further complexity to installation and management.

Widespread support for iSCSI SANs did not begin to emerge until software initiators for iSCSI became available in most popular operating systems. Providing native iSCSI initiators in server OSs had the greatest impact on iSCSI adoption because IT staff no longer had to rely on HBAs and a network stack from a third-party vendor or deal with the related interoperability and standards compliance issues. With certified and standardized initiators in the OS, IT staff could use standard Ethernet switches and NICs (or LAN on motherboard), such as the Intel® Gigabit Server Adapter family, for iSCSI SAN implementation.

While iSCSI initiators provide a low-cost and simplified entry to iSCSI SAN implementation for small- or medium-size businesses, lack of the remote-boot capability available with HBAs is problematic for enterprise-level organizations. This is solved by recently added remote-boot capabilities, such as the Boot-from-SAN feature of Microsoft Windows Server* 2003 and the iSCSI Remote Boot capability being provided with all Intel Gigabit Server Adapters for PCI Express. This allows achieving, and enhancing, full SAN capabilities with the lower cost and simplified implementation provided by using standard Ethernet NICs and network components. Intel iSCSI Remote Boot capability also supports Red Hat Linux* 4, SUSE Linux* 9, and subsequent releases of those OSs.

Remote booting from the SAN offers a number of advantages and simplifications. These include reduced equipment and management costs, better security and performance, and easier consolidation of server and storage resources, and centralization of storage management. These advantages become particularly important in supporting centralized mass storage in the trend toward server consolidation and virtualization.

More specifically, remote boot provides the following advantages:

- Server Consolidation and Virtualization — Remote boot eliminates booting each server from its own direct-attached disk. Instead, servers can boot from an OS image on the SAN. This is particularly advantageous for using diskless rack-mount or blade servers in the high-density clusters often used in grid computing as well as for server consolidation and virtualization.
- Centralized Management — With OS images stored on the SAN, provisioning new servers and applying upgrades and patches to existing servers becomes easier to manage. When a server remotely boots from the SAN, it automatically acquires the latest upgrades and fixes. This also simplifies recovery from server failure because the server, or a spare server, can be booted and provisioned from centralized OS, application, and data images on the SAN.
- Improved Disaster Recovery — All information stored on a local SAN – including boot information, OS image, applications, and data – can be duplicated on a remotely located SAN for quick and complete disaster recovery. An iSCSI SAN provides even greater assurance of disaster protection and recovery. This is because iSCSI SANs can be located anywhere Internet connectivity is available, thus providing greater geographic separation and protection from local and region natural disasters, such as earthquakes and hurricanes.

Intel® I/OAT Speeds Up iSCSI SANs

Intel® I/O Acceleration Technology (Intel® I/OAT) is a newer technology now available on new Dual-Core and Quad-Core Intel® Xeon® processor-based servers and supported by all Intel Gigabit Server Adapters for PCI Express. Unlike TOE, which focuses only on one part of server I/O issues, Intel I/OAT provides system-wide I/O acceleration benefits. This includes reducing system overhead as a percentage of CPU usage, faster memory accesses, and highly efficient TCP stack and TCP/IP packet processing.

Since iSCSI SAN traffic consists of TCP/IP packets, Intel I/OAT provides an acceleration benefit for iSCSI SANs. This, combined with the higher bandwidth of the PCI Express serial bus, makes Intel Gigabit Server Adapters for PCI Express an excellent choice for iSCSI SAN implementation with either Dual-Core or Quad-Core Intel Xeon processor-based hosts. Also, with Intel I/OAT, there is no need to modify the operating system or applications, as is the case with the TOE solution provided on some HBAs.

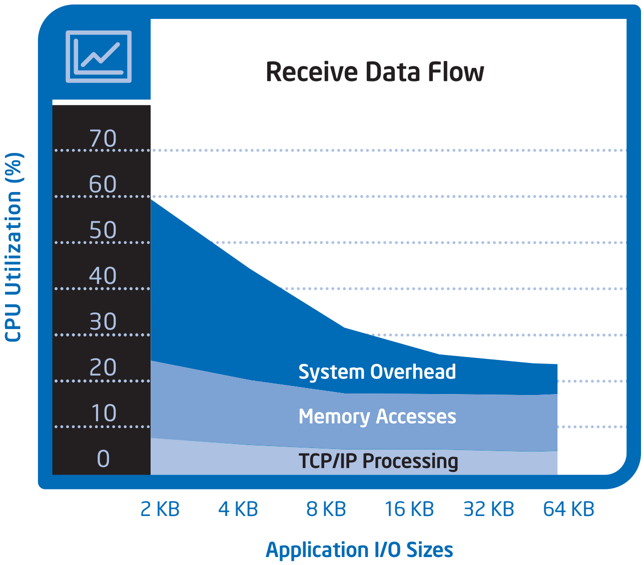

Figure 2. I/O overhead varies according to I/O size. I/O overhead, as measured by CPU utilization, varies according to application I/O packet size, with TCP/IP processing remaining fairly constant and tending to be the smaller part of CPU utilization compared to system overhead and memory accesses.

As a means of comparison between Intel I/OAT and TOE, consider the major sources of network I/O overhead shown by the chart in Figure 2. Notice that TCP/IP processing overhead as a percentage of CPU utilization is essentially constant regardless of application I/O packet size, while system overhead and memory accesses play larger roles in overall I/O overhead, especially for I/O sizes below 8 kilobytes (KB).

Since TOE only addresses TCP/IP processing, it addresses only one part of the overall I/O issue, which is also the smallest part of the I/O performance problem. Intel I/OAT address all three major I/O issues – system overhead, memory accesses, and TCP/IP processing – to provide enhanced I/O performance.

The performance advantage of Intel I/OAT has been born out by testing done by Intel as well as by testing done by an independent laboratory. Figure 3 shows the Intel test results, which compare iSCSI read/write throughput and percent CPU utilization for an Intel® PRO/1000 PT Dual Port Server Adapter versus an adapter from a leading iSCSI HBA vendor.[§] Tests for both adapters used the same Dual-Core Intel Xeon processor-based server and the same iSCSI SAN target and network configuration to ensure an equitable comparison between the Intel adapter with Intel I/OAT and the iSCSI HBA.

Microsoft Windows Server 2003* SP2 Iometer Dual-Port iSCSI Read/Write Test

Variable Buffer Performance Comparison for Intel® PRO/1000 PT Dual Port Server Adapter with Intel® I/OAT Enabled versus iSCSI HBA

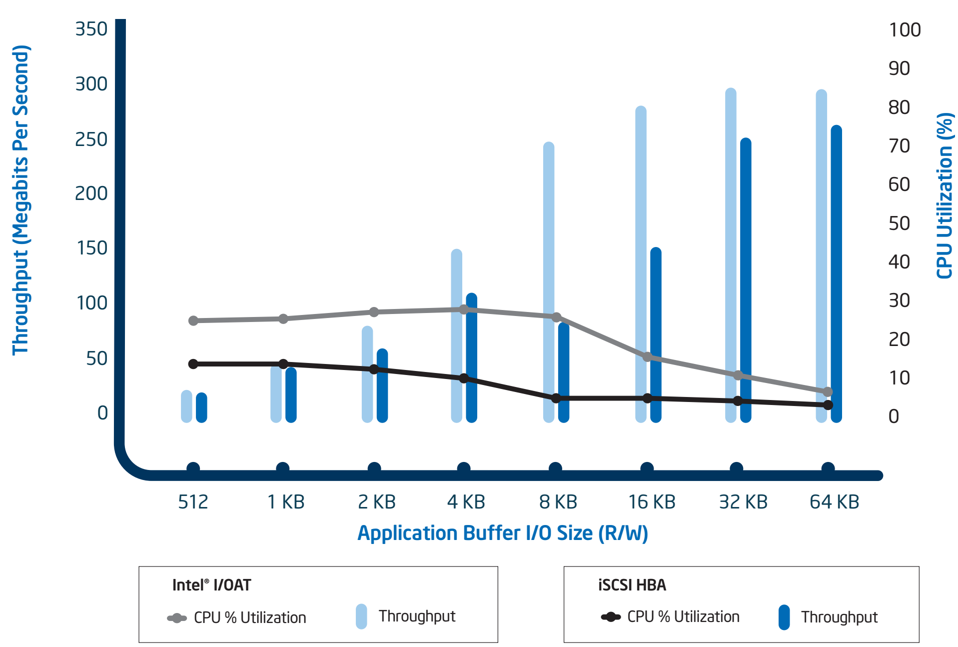

Figure 3. Performance comparison between Intel® I/OAT and a competing iSCSI HBA. For all I/O sizes, Intel I/OAT provides more throughput than the HBA and more than twice the throughput performance in the 8 KB to 16 KB range.

The bar graph in Figure 3 compares throughput performance across a range of I/O sizes, while the line graph compares percent of CPU utilization. Notice that the Intel adapter and Intel I/OAT provide higher throughput than the HBA across all tested I/O sizes, with throughput advantages of twice as much in the 8 KB to 16 KB I/O-size range. Also, CPU utilization diminishes noticeably for increasing I/O sizes; however, percent CPU utilization becomes much less of an issue with the greater headroom provided by new dual-core and quad-core processor-based servers. In the case of Figure 3, the higher CPU usage exhibited by Intel I/OAT is the result of a correspondingly higher throughput in most cases.

Similar independent testing produced similar results. An independent test facility showed that the Intel server adapter with Intel I/OAT demonstrated greater throughput than a competing HBA over all tested I/O packet sizes, with twice the performance of the HBA in some cases.

Typical SAN Architectures

The best-practice recommendation for SAN implementation is to provide a separate, dedicated port for the SAN, making the SAN a separate network segment. Additionally, for higher reliability, it is considered wise to provide an additional redundant SAN link for fail-over purposes. For even higher reliability and increased performance, still more ports can be used to gain the bandwidth and throughput benefits of adapter teaming and link aggregation.

When using HBAs for SAN implementation, providing extra ports for redundancy quickly drives costs up. It also consumes valuable slots on slot-constrained servers. In contrast, availability of iSCSI initiators and remote boot allows use of lower-cost Ethernet NICs to augment LAN on motherboard (LOM) server ports. Additionally, valuable server slots can be conserved by using Intel Gigabit dual-port or quad-port server adapters for PCI Express to provide both primary links and secondary links for link redundancy.

A basic SAN architecture for coexisting Fibre Channel and iSCSI was shown earlier in Figure 1. Figure 4 shows various other configurations of SANs for the enterprise data center using both high-availability (H/A) hosts with redundant links and non-H/A hosts. This includes iSCSI SANs in Figure 4a and 4c as well as iSCSI-to-FC bridging in Figure 4b and 4d.

Other iSCSI Issues Addressed

Even with iSCSI initiators and remote-boot capabilities, some concerns about iSCSI SANs continue to surface. One concern, especially from TOE proponents is that CPU utilization will be too high. Independent tests show that this is not the case for iSCSI implementations augmented by Intel I/OAT. Additionally, just like LAN traffic, new multi-core servers can easily handle iSCSI traffic.

Concerns about the bandwidth limitations of iSCSI over Gigabit Ethernet (GbE) have also been voiced, especially by proponents of the higher-bandwidth Fibre Channel technology. However, the truth is that GbE is quite adequate for many iSCSI SAN applications. For those cases where more bandwidth is truly needed, higher bandwidth can be achieved by using GbE link aggregation or by using Intel® 10 Gigabit Server Adapters for iSCSI SAN implementation over 10 Gigabit Ethernet (10GbE).

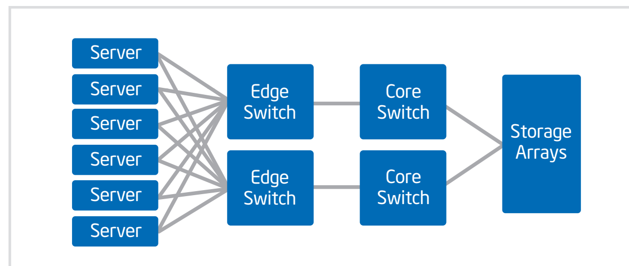

a. Enterprise Data Center—H/A Hosts

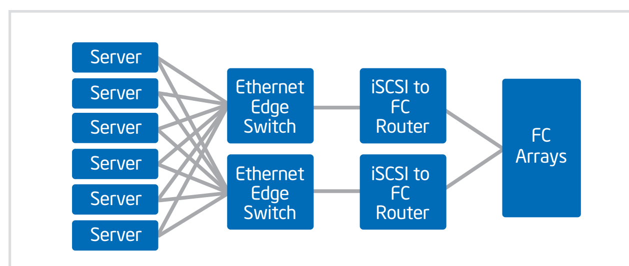

b. Enterprise Data Center—H/A Hosts, iSCSI-to-FC Bridging

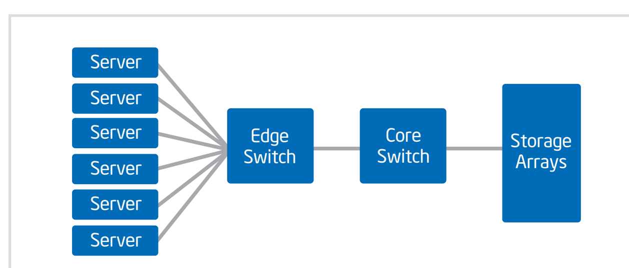

c. Enterprise Data Center—Non-H/A Hosts

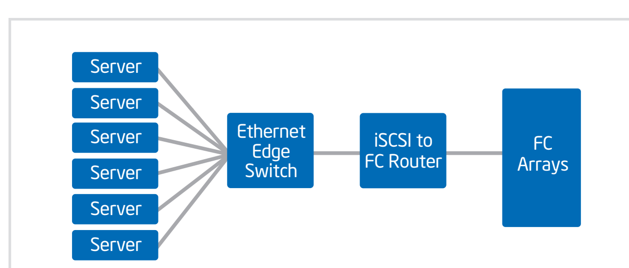

d. Enterprise Data Center—Non-H/A Hosts, iSCSI-to-FC Bridging

Figure 4. Example SAN implementations in an enterprise data center.

Yet another concern is iSCSI security. In reality, security is more a matter of how an IP network is implemented – firewalled, isolated and dedicated, and so forth – rather than any particular security mechanism in a protocol or fabric. However, iSCSI does have the additional advantage of providing Challenge-Handshake Authentication Protocol (CHAP) for authentication, which is also supported by Intel Gigabit Server Adapters.

The iSCSI SAN Advantage from Intel

The ability to create affordable, easy-to-manage iSCSI SANs did not really come to fruition until the advent of iSCSI software initiators and iSCSI remote-boot capabilities. The advent of iSCSI initiators allowed IT to use standard Ethernet NICs, rather than expensive HBAs, for iSCSI SAN implementation. The availability of remote-boot capability further enhanced iSCSI SAN implementation with the ability for centralized management of iSCSI SANs.

To meet the widest range of iSCSI implementation requirements, Intel provides a broad offering of Ethernet LOM and server adapter products. This includes the Intel Gigabit Server Adapters in single-port, dual-port, and quad-port configurations. Intel Gigabit Server Adapters are also available for fiber-optic or copper transmission media. Additionally, for the highest performance requirements, 10GbE capability is available with the family of Intel 10 Gigabit Server Adapters.

The broad range of server adapters and management tools from Intel allows you to use a common set of familiar tools for creating iSCSI SANs with network connectivity you can count on. For example, the Intel® PROSet Utility for Windows* Device Manager provides familiar point-and-click control over individual adapters, advanced adapter features, adapter teaming, virtual local area networks (VLANs), and adapter diagnostics. Such ease and simplicity of adapter configuration and management adds to the cost savings of using Intel Gigabit Server Adapters for SANs rather than the more expensive and complex solutions of iSCSI HBAs or TOE NICs. Plus, as shown by independent testing, Intel Gigabit Server Adapters with Intel I/OAT provide better performance for the I/O packet sizes critical to iSCSI SANs. For iSCSI SANs with higher performance and lower total cost of ownership, the choice is clear – Intel Gigabit Server Adapters with iSCSI Remote Boot capability.

For More Information

To find out more about iSCSI remote boot, visit www.intel.com/network/connectivity/products/iscsiboot.htm

To find out more about Intel Gigabit Server Adapters, visit www.intel.com/network/connectivity/products/server_adapters.htm

To find out more about Intel I/O Acceleration Technology, visit www.intel.com/go/ioat

To find out more about PCI Express, visit www.intel.com/network/connectivity/products/whitepapers/mesh_pcie_whitepaper.pdf

§Performance tests and ratings are measured using specific computer systems and/or components and reflect the approximate performance of Intel products as measured by those tests. Any difference in system hardware or software design or configuration may affect actual performance. Buyers should consult other sources of information to evaluate the performance of systems or components they are considering purchasing. For more information on performance tests and on the performance of Intel products, visit http://www.intel.com/performance/resources/limits.htm.

Printed in USA 0607/TAR/OCG/XX/PDF ♻ Please Recycle 317669-001US