11 Reasons to Choose QLogic iSCSI HBAs Over Software Initiators and TOE

iSCSI Offload HBAs Offer Superior Data Protection, Higher Performance, and Greater Manageability

Executive Summary

Storage area network (SAN) Administrators are under constant pressure to reduce costs and do more with less, particularly in the area of storage systems. iSCSI technology offers an alternative to Fibre Channel (FC) for implementing SANs enabled by Ethernet at a significantly lower total cost of ownership. As an emerging technology, however, it requires SAN Administrators to make more choices to implement a solution that meets their needs.

For example, one choice that SAN Administrators need to make when implementing an iSCSI based SAN is the type of technology that will be used for connecting systems to the SAN. iSCSI offers three connectivity choices: 1) A generic Ethernet NIC coupled with a software based iSCSI initiator specific to the required operating system, 2) A general purpose TCP offload engine (TOE) coupled with a specific operating system based iSCSI initiator or, 3) A fully integrated iSCSI offload based HBA.

This paper compares and contrasts the three options for connecting a system to an iSCSI SAN, and describes why a full iSCSI offload HBA is the best method for connecting a server or workstation to an iSCSI storage network in terms of performance, critical data protection, and ease of management.

Key Findings

Full iSCSI offload HBAs offer SAN Administrators what they need in a storage adapter, including:

- **Predictable Performance:** Full iSCSI offload HBAs consistently have low CPU overhead; cost of ownership is lowered by providing performance headroom for growing user application requirements.
- **Reliability:** Only full iSCSI offload based HBA implementations can offer iSCSI digest reliability at line speeds without impacting the performance of host applications.
- **Manageability:** A single hardware and management platform offers the flexibility to support a heterogeneous operating system.





For more information on Virtualization Performance, you can reach us at: Website: www.virtualizationperformance.com Sales: (866) 496-4560 Tel: (813) 486-5870 Fax: (813) 436-5301

Introduction

SAN Administrators prefer to install systems that offer scalability to ensure that a system is not obsolete the moment it is installed. In addition, making changes to the hardware configuration is not practical once the system is in service. Ideally, the system should provide all the required performance for its life cycle the moment it is installed. Scheduling down-time to install components to increase performance is simply not an option due to the risk of causing a failure in an otherwise working system.

Software Based Initiators: This strategy does not work well with iSCSI software based initiators. Software initiator vendors advise starting with a software based system, and then migrating to a full iSCSI offload HBA or TCP offload engine when required; however, this is exactly the opposite of what a SAN Administrator would prefer to do to reduce the risk in providing highly available storage solutions.

TCP Offload Engines: TCP offload engines meet the needs for providing performance headroom and avoiding potentially disruptive system hardware changes later in the life of the system. However, due to their focus only on TCP/IP offload, they do not provide all the potential iSCSI specific hardware offloads and configurability that can further improve overall iSCSI performance.

Full Offload iSCSI HBAs: Full offload iSCSI HBAs are the best choice in terms of performance. They provide the initial performance boost of TCP offload engines when the system in initially installed and the best combination of TCP/IP and ISCSI offloads, as well as iSCSI specific configurability of the network parameters needed to optimize iSCSI performance.

Performance Comparison

There are two primary measures of disk I/O performance: **Throughput** is the maximum data transfer rate of a device, and **IOPS** measures the number of I/O requests per second that a device can service. The industry standard tool for measuring disk I/O performance is IOMeter. It generates a series of disk I/O, and can be configured to generate a mix of read and write operations in random or sequential access patterns. iSCSI throughput will be compared for the iSCSI initiator combinations found in the following table.

Physical Network Interface	iSCSI Initiator	Manufacturer	Model
Full iSCSI Offload HBA	Offload HBA	QLogic	QLA4050C
TCP Offload Engine (TOE enabled)	SW	HP	NC370T
Network Interface Card (TOE disabled)	SW	HP	NC370T

Reason #1: Higher Throughput

The following figure shows each adapter's maximum throughput using an equal ratio of 32KB read and write I/O requests. To test for maximum throughput, a large I/O request size, such as 32K bytes, removes the effects of CPU performance from the test results.

Results show that the full iSCSI offload HBA can sustain bi-directional throughput **more than 32% higher** than a TCP offload engine and **5% higher** than a NIC. In addition, at maximum throughput, the CPU utilization on the iSCSI HBA is one-half that of the TOE and one and one-fifth that of the NIC.



SN0130907-00 Rev A



Reason #2: More Efficient IOPS

IOPS (the number of I/O requests per second) represent the total number of I/Os that a system can service per second. Taken in conjunction with the host's CPU utilization, IOPS represent the efficiency of the iSCSI adapter and protocol stack. The patterns used by many applications include a large number of smaller I/O requests rather than a smaller number of large I/O requests; this test closely models the I/O activity of an actual application and can roughly estimate performance impacts of each initiator type.

The following figure shows the full iSCSI offload HBA sustaining the highest number of IOPS, the generic NIC and the TCP offload engine a lesser number.

The key difference of interest to SAN Administrators is the how efficient each technology is at processing each I/O operation. When considering the number of I/Os per percentage point of CPU utilization, the iSCSI HBA is clearly more efficient at processing I/Os than the TOE engine or NIC. When systems are running using a TOE engine or NIC, fewer cycles are available for application processing.



iSCSI HBA: Highest IOPS with lowest CPU Overhead



Reason #3: Scalability

The CPU efficiency advantage of a full iSCSI offload HBA becomes even more important when multiple adapters are installed into a single system. Even on high powered CPUs, the CPU overhead incurred to run iSCSI sessions simultaneously on two or more NICs can overwhelm the processor, leaving either no CPU cycles left for the application or greatly reduced IOPS and throughput performance. In contrast, with a full iSCSI offload HBA handling the TCP/IP stack and the iSCSI processing, this disadvantage does not exist; rather, the CPU can focus on the needs and requirements of the application.

Reliability

Reason #4: Reliability

The number one criterion for data storage and storage area networks is that data should never be corrupted. The TCP protocol used by iSCSI includes simple checksums to protect against data corruption; however, checksums are not reliable enough to completely protect iSCSI data against errors. According to Stevens "TCP/IP" Illustrated Volume 1

"... if your data is valuable, you might not want to trust either the UDP or the TCP checksum completely, since these are simple checksums and were not meant to catch all possible errors."

To remedy this problem, Header and Data Digests were added by the iSCSI Working Group as a more robust mechanism for ensuring data integrity compared to TCP checksums. However, iSCSI Header and Data Digest calculations are very CPU intensive. Only a full iSCSI offload HBA has the logic built into the ASIC to accelerate these calculations. General purpose NICs and TOEs do not have this innate capability; therefore, the calculations must be performed by the host CPU (if desired). If these calculations are performed by the host CPU, both throughput and IOPS performance will further degrade, potentially slowing application performance to an unacceptable level. This decrease in performance makes general purpose TOE and NICs with software initiators impractical for use where data integrity is important.





Manageability

Manageability is a key factor when selecting an iSCSI initiator, since the cost of ownership will be partially defined by the amount of management effort required to initially install and then monitor and maintain the iSCSI SAN.

Reason #5: Broadest OS Support

Nearly all modern networks are heterogeneous environments with Microsoft Windows machines for desktops and e-mail servers; Linux machines for web servers and AIX; HP-UX or Solaris machines for database servers. To optimize SAN usage, these servers must connect to the same SAN, and ideally be managed in the same manner with the same set of tools. Most operating system vendors now provide a software based initiator for their systems; however, the interfaces and methods to manage these initiators vary, making the process to learn and stay current difficult. The same problem exists for TOE engine based iSCSI initiators, which rely on the software initiator to implement the iSCSI protocol. A full iSCSI offload HBA with full support for heterogeneous operating systems designed to serve the iSCSI SAN market can best meet the needs of SAN Administrators who need an iSCSI initiator with a uniform storage interface that uses their existing management tools.

Operating System	QLA4050C Full Offload HBA	HP NC370T TOE	Software Initiator
Windows 2000 (x86, x64)	Y	Y	Y
Windows XP (x86, x64)	Y	N	Y
Windows 2003 (x86, x64)	Y	Y	Y
RHEL 3 (x86, x64)	Y	Y	Y
RHEL 4 (x86, x64)	Y	Y (only x86)	Y
SLES 9 (x86, x64)	Y	Y (only x86)	Y
SLES 10 (x86, x64)	Y	N	Y
Solaris 8 (SPARC)	Y	Ν	N
Solaris 9 (SPARC)	Y	N	N
Solaris 10 (x86, x64, SPARC)	Y	N	Y
AIX	Y	N	Y
VMware	Y	Ν	N

RELIABLE HIGH PERFORMANCE SANS



Reason #6: Management Tools

Efficient management depends on the availability of tools that fit the way SAN Administrators prefer to work and the specific task being performed. To provide for all types of management styles and functions, both a graphical and command line interface (CLI) are required to enable the highest level of manageability.

Software initiator tools support varies widely between operating systems, leaving SAN Administrators no consistent method to manage their iSCSI SANs. This makes NIC based and TCP offload engine based initiators difficult to manage in an easy and logical manner. With their consistent storage interface, full offload iSCSI HBAs can easily support both a graphical interface and CLI that are consistent between operating systems, enabling SAN Administrators to quickly perform any management task or automate repetitive tasks with ease.

QLogic's full offload iSCSI HBAs include the iSCSI HBA Manager, a graphical user interface for configuring iSCSI HBAs, as well as a full featured CLI. In addition, QLogic as the singular and consistent point of contact for support and management across a wide range of operating systems makes contacting and receiving support easy.



SN0130907-00 Rev A

WP06007

Reason #7: Single Step iSCSI Boot

iSCSI boot allows a host server to boot its operating system over an iSCSI connection. To enable iSCSI boot, an adapter must be able to attach to target LUNs and present those LUNs to the host server during the initial power-on stage of the server.

Software iSCSI initiators running on NICs or TOEs require a two step boot process using an additional server to perform an iSCSI boot. This process requires additional hardware for the second boot server and perhaps more problematic additional management time to administer the server and boot LUN assignments. Full offload iSCSI HBAs can be simply configured to directly boot from a LUN on the SAN in a single step. They can also be configured to learn which LUN to use while leveraging existing services like DHCP.





Reason #8: Network Tuning

Network configuration is a key factor in the performance of any networked application. Vendor differences in switch implementation, topology difference, and traffic patterns all contribute to varying levels of performance and reliability. No standard set of TCP parameters can possibly optimize network performance for every application and workload. Software initiators are subject to the underlying parameters provided by the NIC or TOE adapter they are using for connectivity. iSCSI initiators implemented on generic networking NIC or TOE adapters can only take advantage of limited tuning parameters offered by these products. Full offload iSCSI HBAs offer a wider range of tunable TCP parameters to account for iSCSI specific network traffic and network differences.

In addition to standard TCP/IP networking configuration parameters, QLogic's full offload iSCSI HBAs provide the following configurable parameters:

- TCP Window Size
- TCP Window Scale
- TCP Timestamps
- TCP Delayed ACKs
- TCP Selective ACK
- 802.1x Flow Control / Ethernet Pause
- Maximum Segment Size (MSS) Jumbo Frames

Reason #9: Investment Protection

Choosing an iSCSI HBA means choosing a vendor with a commitment and understanding of the storage system market and technology. NIC and TOE products, with their one-size-fits-all model that implement the iSCSI protocol in a software layer above the hardware, don't have as much experience as a vendor who is committed to the storage market.

Several iSCSI players have disappeared, where as QLogic has deployed three generations of 1Gb iSCSI products **with an estimated 90% market share in iSCSI HBAs**. By choosing QLogic's iSCSI HBA, a user gets the additional benefits of a strong iSCSI roadmap and commitment to emerging technologies such as 10Gb speeds and virtualization.

Reason #10: Support for SAN Infrastructure

No single technology can meet every need in the storage system market. QLogic supports SAN Administrators by supplying a wide variety of popular technologies used in storage networking including:



- iSCSI HBAs
- Fibre Channel HBAs
- Fibre Channel Switches
- Fibre Channel/iSCSI Routers
- InfiniBand Adapters

By implementing QLogic storage networking components, SAN Administrators can enjoy centralized management tools and support from a single vendor.

Reason #11: Single Vendor Global Support

Storage technologies can be complex, so QLogic's Global Support infrastructure is committed to support customers from the initial installation and systems planning stages to solving operational issues when they arise. NIC and TOE vendors do not have the broad experience in storage system implementations that is needed when implementing high availability systems required in today's business environment. In addition, when using software initiators for iSCSI connectivity, the customer has a different contact point for each operating system. By using a hardware based iSCSI HBA from QLogic, the user has a single contact point for ALL iSCSI initiator support issues, greatly simplifying administration and lowering the overall total cost of ownership of an iSCSI SAN.

In a continuing commitment to bring best-in-class pre- and post-sales support to customers, QLogic offers the broadest support, fastest response times, and the most extensive geographic coverage of any HBA supplier in the industry today. QLogic offers standardized global support via email and telephone support to contact personnel **in the customer's preferred language**. QLogic support personnel are focused on QLogic products, receive uniform training, and are committed to responding either by phone or email to all phone requests within 30 minutes during normal business hours.

QLogic support personnel are also authorized to provide advanced replacement products as needed. QLogic's comprehensive support services are offered at <u>no additional</u> <u>charge</u> for products that are covered by warranty.

	QLogic	
Comprehensive Global Support Services		
Products Supported	 FC HBAs iSCSI HBAs InfiniBand HCAs FC Switches FC Storage Solutions Multi-Protocol Bridges/Routers 	
Support Forums	http://solutions.qlogic.com/	
Advanced Product Replacement	√	
30-Minute Response Time	1	



SN0130907-00 Rev A



Summary and Conclusion

Many factors must be taken into consideration when making a choice between a software initiator with a generic NIC or TOE and full offload iSCSI HBA for an iSCSI initiator. The full offload iSCSI HBA is the right choice for IT environments where performance, reliability, and manageability are key factors.

QLogic's iSCSI HBAs are the perfect choice in a full offload iSCSI HBA due to their superior:

- **Predictable Performance:** Full iSCSI offload HBAs consistently have low CPU overhead, leaving additional cycles for application processing.
- **Reliability:** Only full iSCSI offload based HBA implementations can offer robust iSCSI digest reliability at line speeds without impacting the performance of host applications.
- **Manageability:** A single hardware and management platform offers the flexibility to support heterogeneous operating systems.

In addition to the technical reasons for choosing full iSCSI offload based HBAs, they are also the right choice for lowering total cost of ownership and raising overall return on investment (ROI). By offloading all of the iSCSI protocol (including hardware digests), they free up host cycles that can now be used by other applications, extending the life of systems or allowing IT Administrators to purchase fewer systems. Having a single solution for iSCSI SAN initiators significantly reduces training and operational costs, further improving ROI. Finally, when purchased from a vendor that is also a supplier of Fibre Channel SAN solutions, it is now possible to further save costs on support when there is a single vendor contact for all issues across all SAN implementations in an organization.

QLogic iSCSI Host Bus Adapters offer best in class performance, features, manageability, compatibility, and support. The QLA405*x* family is based on QLogic's second generation iSCSI technology and is the industry's de facto standard for iSCSI adapters.





Appendix A

References

- 1) QLogic's iSCSI drivers, management tools, and documentation at: <u>www.qlogic.com</u>
- 2) MicroSoft's iSCSI initiators can be downloaded at: www.microsoft.com
- 3) IOMeter download and documentation at: www.iometer.org
- 4) Stevens, Richard, TCP/IP Illustrated, Volume 1, Addison-Wesley Professional, ISBN 0201633469



Disclaimer

Reasonable efforts have been made to ensure the validity and accuracy of these comparative performance tests. QLogic Corporation is not liable for any error in this published white paper or the results thereof. Variation in results may be a result of change in configuration or in the environment. QLogic specifically disclaims any warranty, expressed or implied, relating to the test results and their accuracy, analysis, completeness or quality.

RELIABLE | HIGH PERFORMANCE | SANS

