




The Value of Full Hardware Offload in a Converged Ethernet Environment

QLogic Offload Conserves Critical CPU Resources

QLogic 8200 Series Offers Superior Performance Over Software Initiators



Executive Summary

Most SANs today are built using Fibre Channel technology, which offers a highly reliable, robust, and mature storage protocol. The protocol meets the data integrity and performance requirements of enterprise data center customers running critical applications and enterprise storage solutions. However, there is a fast-emerging new standard, FCoE (Fibre Channel over Ethernet), which promises to introduce the data center trend of consolidation to the network. FCoE provides a direct mapping of Fibre Channel onto Ethernet and enables the benefits of Fibre Channel traffic to be natively transported over the ubiquity of Ethernet networks. Migrating to FCoE brings additional benefits, such as I/O consolidation and lower management costs, while preserving investments by leveraging existing Fibre Channel knowledge and management tools, which can be applied directly to FCoE.

When deciding on implementing an FCoE solution, it is important to know that you have two choices: software initiators or offload engines. While software initiators are a low-cost way for an organization to explore the benefits of FCoE SANs using existing 10GbE NICs in servers, offload engines are specialized adapters (Converged Network Adapters) designed for concurrent I/O

support, conserving precious CPU cycles for applications, services, and virtual server environments. In addition, offload engines can be used to address emerging performance requirements.

Besides cost, three other factors—performance, reliability, and scalability—need to be balanced, along with other considerations, when determining the proper interconnect to use within an FCoE-enabled data center. The questions many administrators will be asking include the following:

- To achieve I/O consolidation and cost reduction, should I utilize low-cost FCoE software initiators or the more expensive Converged Network Adapters with built-in processors to offload from the CPU?
- What are the tradeoffs of saving money on a NIC versus a Converged Network Adapter?

This white paper will discuss the advantages and provide guidance for making an informed decision.

The Value of Full Hardware Offload in a Converged Ethernet Environment QLogic Offload Conserves Critical CPU Resources

Key Findings

QLogic® 8200 Series Converged Network Adapter Advantages
• On-chip processing reduces CPU workload
• Field-proven Fibre Channel driver stack ensures data integrity
• Fibre Channel storage market leader
• 8.5 million ports install base assures widest interoperability and compatibility
• Principal design for all major OEMs
• Largest set of OEM server, OS, and storage qualifications
• Concurrent support for NIC, FCoE, and iSCSI
• Enterprise-class reliability and error recovery
• Experienced storage and networking support staff
Open FCoE Initiative with a 10Gb Ethernet NIC Advantages
• Low initial cost

Introduction – The Situational Impacts of Open FCoE

Open FCoE, a Linux community open source project, was started by Intel® with the goal of encouraging the development of a native FCoE code. This code base provides for Fibre Channel protocol processing over Ethernet-based transport and acts as a low-level device driver to send and receive data packets. Open FCoE is now being released within Red Hat® Enterprise Linux® (RHEL®) distributions, which will help propel storage over converged Ethernet solutions. However, there are several caveats; not just any NIC can be used. Ethernet NICs must support new Ethernet standards, such as Priority-based Flow Control (802.1Qbb) and Enhanced Transmission Selection (802.1Qaz). In addition, an FCoE Switch with a Fibre Channel Forwarder (FCF) is required to log in to a Fibre Channel fabric.

Despite the fact that Open FCoE is a good indication of how FCoE is being accepted as part of the Linux infrastructure, there are many factors to consider when preparing for migration to FCoE. The need for I/O consolidation, cost reduction, scalability, application performance, and data integrity are all key factors in the decision-making process. Although Open FCoE solutions leveraging a software initiator and an inexpensive NIC can be great for many applications, when it comes to enterprise-class applications and storage an Open FCoE driver solution not only fails to meet basic requirements but can also be detrimental to data center virtualization objectives.

Scalability Within Server Virtualization Environments

Within the evolution of data centers, server virtualization is the key solution driving server hardware consolidation, as well as I/O consolidation. FCoE provides a platform for I/O consolidation by reducing the number of adapters required for transporting multiple standard I/O protocols. This follows the trend of consolidation in data centers, providing increased flexibility and cost savings. Given this, it stands to reason that FCoE would be used first in data center environments employing VMware®, Hyper-V®, or other hypervisors that are being used for consolidation of hardware. It's true today that we see the very first FCoE implementations taking place within these virtualized environments.

In a virtualized server environment, there are two important considerations that should be well thought out prior to implementing FCoE: the increased density of applications per physical server and the addition of a virtualization layer. Both will require increased I/O performance. Application density across an I/O adapter is increased within these environments due to virtualized servers running multiple CPU cores with up to 12 or more virtual machines (VMs) per core. Each of these virtual servers and their applications are often running over a single physical adapter (in a fault tolerant situation, more than one adapter is used to provide redundancy). This creates a very dense application environment and places increasing I/O performance demands on adapters.

Although virtualization is driving consolidation, it also imposes additional overhead on the CPU, network, and I/O. Virtualized environments use the hypervisor to perform software emulation to abstract IT resources away from the physical hardware. This method comes at a considerable cost, as CPU overhead is incurred because virtual resources need to be mapped to physical resources. Additional overhead in a VM scaling environment further compounds the problem. For example, consider a single physical processor core within a server that supports multiple virtual processors. When there are 12 or more VMs installed per physical processor core, a single physical processor is now hyper-threading across many virtual machines, hindering the application within the virtual servers from scaling linearly.

With limited CPU resources, it is not logical to further burden the CPU with processing I/O storage requests, which can be done with the addition of an offload adapter, such as the QLogic 8200 Series Converged Network Adapter. To put this into perspective, consider a video card in a personal computer used for gaming that offloads the video processing from the CPU. Gamers are well aware of the advantages of offloading the video processing and using graphic adapters to conserve all available CPU resources for the video game itself. Advantages of the video card include stutter-free video

The Value of Full Hardware Offload in a Converged Ethernet Environment QLogic Offload Conserves Critical CPU Resources

playback, vibrant high-definition images, and, most importantly, CPU acceleration to support the most performance-hungry gaming applications.

When it comes to an enterprise-class application, an Open FCoE driver solution not only fails to meet enterprise requirements but it also can be detrimental to data center virtualization goals, a targeted environment for FCoE. Data center consolidation through virtualization will require CPU horsepower to efficiently scale VMs and provide the I/O bandwidth needed for enterprise applications. This can only be provided by an adapter with offload capabilities.

Enterprise Requirements for FCoE	Converged Network Adapter	Software Initiator
Concurrent I/O Support for Consolidation	Yes	No
Scalability within Virtual Operating Environments	Yes	No
Support for I/O Intense Applications	Yes	No
Data Integrity Assurance	Yes	No
Efficient CPU Utilization	Yes	No
IOPS Scalability	Yes	No
Enterprise Reliability	Yes	No
Investment Protection	Yes	Yes
Broad OS Support	Yes	No
Interoperability	Yes	No

The Underlining Cost of Open FCoE

The first thing you may hear about software initiators and the Open FCoE initiative is that it's free. While there may be some truth to the initial savings from a NIC compared to a Converged Network Adapter, it's important to consider the hidden costs of a software initiator. For example, a software initiator will consume approximately 1GHz of processing power per instance managing FCoE transmissions for full duplex read/write operations. With today's multi-core processors this may not seem like a problem, but combined with server virtualization, where a single processor core is being shared by multiple VMs, each VM will require 1GHz each just to process storage and data networking requests. In enterprise environments, this leaves inadequate CPU capacity devoted to virtual servers and business-critical applications that run on the same server. QLogic 8200 Series Converged Network Adapters free these valuable CPU cycles, improving the computing system's performance without the cost of upgrading or adding additional CPU(s). This will allow you to

run more applications per server and extend the useful life of your servers. Additional cost factors should be considered as well, and have been summarized in the following information.

Operational Efficiencies

The aforementioned software initiator consumption of 1GHz of an embedded processor to process FCoE transmission for full duplex read/write operations can be reduced up to 88 percent by installing a QLogic 8200 Series Converged Network Adapter with offload capabilities. The adapter can increase operational efficiency because more VMs can be supported per CPU; an increase in throughput is delivered per adapter, which allows increased scalability for applications and VMs. The same QLogic adapter can be used to offload iSCSI and provide Ethernet connectivity, replacing three adapters with one, further increasing operational efficiencies.

CPU Efficiency

Moving to a converged infrastructure can provide significant advantages, but relying on server resources to drive a protocol processing request is a less than optimal solution. If users believe that some applications will require a higher I/O load, and if they want to share the infrastructure with the rest of the servers and applications, moving to a converged infrastructure can provide these advantages. However, through Open FCoE, software initiators will use server CPU resources to complete processing of all supported I/O protocol stacks. Many claim that with the low cost of processors today a valid option would be to throw cheap CPU resources at I/O requests. This does not necessarily make sense; IT shops need to evaluate which applications require the advantages of FCoE networks, and which operating systems and applications will be running on them. When doing this, server virtualization appears to come out on top.

Server virtualization is an ideal candidate for FCoE because the increased bandwidth of 10GbE can support high throughput from multiple VMs. In addition, mobility of virtual servers, load balancing, and failover will require similarly high throughput, especially when it is being shared by multiple VMs. Herein is the problem with Open FCoE; decreased scalability, increased bottlenecks, and hidden costs.

Decreased Scalability

One adapter running a software initiator could easily consume up to one-third of the CPU processors' capabilities. The addition of multiple adapters into a single system relying on the CPU to perform their multiple operations only compounds the problem, as does the addition of VMs. However, offloading onto a QLogic 8200 Series Converged Network Adapter leaves plenty of room for scaling to

The Value of Full Hardware Offload in a Converged Ethernet Environment QLogic Offload Conserves Critical CPU Resources

multiple adapters and multiple VMs without impacting the overall performance of the server.

Bottlenecks

Using a software initiator on a NIC requires that every incoming TCP/IP, FCoE, and iSCSI packet traverse the PCI bus in the server. Sending packets back and forth increases the PCI bus' busy state, and can cause bottlenecks with other hardware on the PCI bus. The QLogic 8200 Series Converged Network Adapter offloads all protocol processing (FCoE, TCP, iSCSI, and SCSI digest for data integrity) onto the adapter. Therefore, using a QLogic 8200 Series Converged Network Adapter will result in reduced bottlenecks and increased throughput of application data across the PCI bus.

Key Differentiator

Offloading functions from the CPU is not a novel or risky concept. Look at the popularity of dedicated graphics controllers in the gaming world. Offload engines are a better alternative to software initiators in servers. Businesses can maximize CPU cycles available for application or I/O services and address emerging performance and scalability requirements of enterprise data centers.

Hidden Cost

To simplify this, let's start by assuming you spend \$3,000 on a typical server with one CPU core. With a software initiator you'd be relying on one-third of the CPU to process FCoE and/or iSCSI requests for one NIC. Simple math tells us that you're actually paying \$1,000 to process your FCoE traffic. Now, let's complicate things by looking at a similar solution with next-generation servers running a dual socket, quad core CPU with server virtualization. Today, typical virtual operation environments (VOE) are estimated at running approximately four VMs per CPU core (IDC reports indicate that the industry is quickly moving toward 10 – 12 VMs per core). Even in this light VOE, the server is supporting 32 VMs with eight cores or one-fourth of a core per VM. In this environment, there just is not enough CPU processing power to support protocol processing for all VMs, let alone headroom for future scaling requirements. This work can be done at a fraction of the cost with the hardware acceleration offered by the offload engine of a QLogic 8200 Series Converged Network Adapter, leaving the CPU to process business applications as intended and leaving plenty of room to scale.

When it comes to cost, the initial offering of Open FCoE and an inexpensive NIC may sound like you are getting something for free, but do not forget about the opportunity costs you'll be missing out on by not having offload functions. If you still are unsure about an offload solution, be advised that offloading functions from the CPU is not a novel or risky concept. The popularity of dedicated graphics controllers in the gaming world as well as dedicated disk controllers in the server world has become the norm. Offload engines are a better alternative to software initiators in enterprise servers because organizations can maximize CPU cycle availability for application or I/O services, as well as address emerging and future performance and scalability requirements within enterprise data centers.

What's the Cost of Data Integrity?

Along with the server efficiencies that the QLogic 8200 Series Converged Network Adapter can provide, there are also efficiencies and advantages within the I/O processing of a Converged Network Adapter over an Open FCoE solution. These advanced features will ensure that your data is being delivered accurately, as well as the validity of the data when it reaches the disk.

Reliable Data Delivery

The FCoE protocol adopts a data processing mechanism similar to Fibre Channel to maintain the same level of data integrity, while sending storage data over Ethernet. These mechanisms analyze storage packet headers and the data transmitted to ensure the integrity of the data. This integrity checking is a compute-intensive process that is either performed by the CPU in solutions using a software initiator or performed by the offload engine of a Converged Network Adapter, such as a QLogic solution. An offload engine ensures the highest performance of your Converged Network Adapter and maintains the highest level of data integrity as native Fibre Channel.

Low Impact Error Recoverability

As traffic is increased across an Ethernet network, dropped and out-of-order data frames will result. With a software initiator, recovering from both these issues will become a significant burden on the CPU and may cripple overall performance on a 10Gb Ethernet network. A QLogic offload engine, on the other hand, can reassemble out-of-order frames and complete the process of resending those that are dropped while requiring minimal assistance from the CPU and ensuring ongoing performance.

The CPU cycle stealing approach, whether it's for processing the driver stack, traversing data packets over the PCI bus, or for verifying the integrity of the data, is inefficient. By saving a few pennies on

The Value of Full Hardware Offload in a Converged Ethernet Environment QLogic Offload Conserves Critical CPU Resources

a NIC rather than purchasing a Converged Network Adapter, hidden costs will begin to materialize in the form of degrading performance of the data center.

The QLogic Approach

While a wide range of data centers will benefit from FCoE implementations, expect to see a majority of applications that move a significant amount of data, those with compute-intense applications, and environments with server virtualization to be running FCoE over Converged Network Adapters with offload engines. In addition to the two implementation choices, software initiators or offload engines, there are other engine solutions available in the industry. While the below information is relevant in deciding whether to implement a software initiator, it is also helpful in deciding which offload engine to implement.

Battle-Hardened Enterprise Fibre Channel Driver Stack

When it comes to a mature and robust storage driver stack, there is no comparison to QLogic's market share leadership install base of more than 8.5 million Fibre Channel ports. The QLogic 8200 Series Converged Network Adapters leverage these existing Fibre Channel drivers, making the company's Converged Network Adapter the most dependable and battle-hardened adapter on the market. Leveraging existing drivers enables QLogic to offer 12 years of proven field experience, combined with full hardware offload for FCoE and iSCSI protocol processing. The QLogic 8200 Series Converged Network Adapter has an extensive set of features to enhance system virtualization while ensuring industry-leading LAN networking performance. Unproven FCoE implementations from NIC vendors or storage start ups present a high risk for enterprise-class data centers.

OEM Acceptance and Interoperability

QLogic boasts the broadest set of driver certifications in the industry. This is a result of collaborations with leading server, OS, and storage manufacturers, including EMC®, Dell®, HP®, IBM®, Microsoft®, Oracle®, Red Hat, Novell®, and VMware to name a few. To support this effort, QLogic has expanded relationships with each company to fully engage and conduct the industry's most rigorous interoperability testing, spanning every major platform and operating system. This highly comprehensive qualification process tests configurations that replicate real-world environments and help ensure compatibility in the field.

In addition, QLogic has secured the most tier-one OEM design wins for FCoE Adapters in the industry. This achievement indicates two things about QLogic FCoE Adapters: OEMs are backing an FCoE

offload engine from a storage vendor for their FCoE solutions and QLogic products have passed intense scrutiny. QLogic is the leading Fibre Channel vendor.

The same Fibre Channel drivers that have passed countless OEM product and interoperability tests are repurposed by QLogic's Converged Network Adapters and have been put to additional scrutiny to verify storage and Ethernet capabilities within intricate testing trials. Using the same Host Bust Adapter software stack that has been qualified by all of the leading server and storage vendors brings the QLogic Converged Network Adapter an immense advantage that the OEMs recognize.

From an infrastructure perspective, industry-standard technology solutions from companies such as QLogic cannot be touched by Open Source or first-release solutions because their software stacks and hardware will need to go through qualification and testing cycles that will take years to complete. OEMs are not trading low-cost solutions with community collaboration testing because it will result in less interoperability with existing products. Early adopters of the FCoE technology will follow OEMs because they value enterprise-quality interoperability.

Enterprise-quality Reliability and Support

Reliability and support should be top considerations when deciding on an FCoE purchase. QLogic Fibre Channel Host Bust Adapters are unmatched in the category of reliability; they provide the highest Mean Time Between Failure (MTBF). This level of hardware and software reliability will take years for FCoE software initiators to attain. With this in mind, where are you going to turn for support while software initiators are playing catch up? QLogic has years of experience supporting storage and networking solutions while NIC companies attempting to enter the storage market will have no previous knowledge base to reference when they receive storage support calls.

If you want accountability when implementing an FCoE solution, QLogic's enterprise-quality and reliable adapters will emerge as key differentiators over software initiators. Without a history of reliable products and storage support, it is not surprising that early adopters are not accepting less than the enterprise-quality reliability and support that can be offered by a veteran company, such as QLogic.

Summary and Conclusion

When deciding on an FCoE solution, there are two implementation choices, software initiators or offload engines as used in the QLogic 8200 Series Converged Network Adapter. Both deliver desired functionality, but only the Converged Network Adapter approach

The Value of Full Hardware Offload in a Converged Ethernet Environment QLogic Offload Conserves Critical CPU Resources

conserves critical CPU resources. A hardware offload engine strategy provides a much more CPU-efficient approach versus software initiators, which only offer cost savings. Despite the fact that cost is an important factor in the adoption of FCoE, a successful implementation is not possible if performance and reliability suffer materially as a consequence. Outside of price, new FCoE software initiators and partial offload solutions will have a difficult time competing with full offload technology.

With the company's latest hardware release, QLogic enters its third generation of Converged Network Adapters, further demonstrating its commitment to continuously delivering value and innovation. As enterprise application customers strive to achieve density and resource utilization objectives, high I/O performance will emerge as a requirement. Virtualization, a target application for FCoE, will require CPU horsepower to efficiently scale VMs. Furthermore, the addition of a virtualization layer can add overhead, which will degrade I/O performance and eliminate additional virtualization benefits. With the QLogic 8200 Series Converged Network Adapter, the Fibre Channel driver stack is processed by the adapter. Offloading Fibre Channel processing means that more CPU resources can be used to support more VMs, providing greater cost savings and ROI for virtualized environments. The QLogic 8200 Series Converged Network Adapter is uniquely positioned to remove performance barriers for the virtualization of transactional enterprise applications.

Flexibility is another important factor in an FCoE solution. The QLogic 8200 Series Converged Network Adapter supports both iSCSI with full TCP offload and 10Gb Ethernet NIC features with stateless offload, further allowing the host CPU to be used for application and virtualization scaling. In addition, QLogic is currently the only vendor to support NIC, FCoE, and iSCSI protocols concurrently, giving you the opportunity to choose strategic, long-term, and flexible solutions based on your IT needs.

Along with scalability, performance, and offload feature benefits provided over software initiators, QLogic Converged Network Adapters offer additional operational efficiencies, such as eliminating PCI bus bottlenecks, ensuring data integrity with enhanced data delivery and error recoverability capabilities, a mature battle-hardened driver, and the largest interoperability base in the industry.

FCoE applications that move significant amounts of data or information systems that process large files will require solutions that free up CPU resources and increase bandwidth and performance. Software initiators lack maturity for enterprise applications and it will take years of qualification testing to establish the elevated levels set by QLogic. Meanwhile, the QLogic 8200 Series Converged Network Adapter is uniquely positioned to remove the existing barriers of a software initiator.

The Value of Full Hardware Offload in a Converged Ethernet Environment QLogic Offload Conserves Critical CPU Resources

Disclaimer

Reasonable efforts have been made to ensure the validity and accuracy of these 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.



Corporate Headquarters QLogic Corporation 26650 Aliso Viejo Parkway Aliso Viejo, CA 92656 949.389.6000

www.qlogic.com

International Offices UK | Ireland | Germany | India | Japan | China | Hong Kong | Singapore | Taiwan

© 2010 QLogic Corporation. Specifications are subject to change without notice. All rights reserved worldwide. QLogic and the QLogic logo are registered trademarks of QLogic Corporation. HP is a registered trademark of Hewlett-Packard Company. Hyper-V is a registered trademark of Microsoft, Inc. Dell is a registered trademark of Dell Inc. EMC is a registered trademark of EMC Corporation. IBM is a registered trademark of International Business Machines Corporation. Intel is a registered trademark of Intel Corporation. Linux is a registered trademark of Linus Torvalds. Microsoft is a registered trademark of Microsoft Corporation. Novell is a registered trademark of Novell, Inc. Oracle is a registered trademark of Oracle Corporation. Red Hat and RHEL are registered trademarks of Red Hat, Inc. VMware is a registered trademark of VMware, Inc. All other brand and product names are trademarks or registered trademarks of their respective owners. Information supplied by QLogic Corporation is believed to be accurate and reliable. QLogic Corporation assumes no responsibility for any errors in this brochure. QLogic Corporation reserves the right, without notice, to make changes in product design or specifications.