



An Oracle White Paper
May 2012

The Value of Full Hardware Offload in a Converged Ethernet Environment

Executive Overview	3
Introduction — The Situational Impacts of Open FCoE	5
Scalability Within Server Virtualization Environments.....	5
The Underlining Cost of Open FCoE	7
Operational Efficiencies	7
Decreased Scalability	7
Bottlenecks	8
Hidden Cost	8
What’s the Cost of Data Integrity?	8
Reliable Data Delivery	9
Low Impact Error Recoverability	9
The Oracle Approach.....	9
Battle-Hardened Enterprise Fibre Channel Driver Stack	9
Enterprise-quality Reliability and Support	10
Summary and Conclusion.....	10

Executive Overview

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, Fibre Channel over Ethernet (FCoE), 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 understand the two choices of software initiators and offload engines. 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, which can conserve 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 discusses the advantages of each approach and provides guidance for making an informed decision.

TALBE 1. KEY ADVANTAGES OF ORACLE FCOE OFFERINGS

ORACLE'S SUN STORAGE 10GBE FCOE CONVERGED NETWORK ADAPTER ADVANTAGES	OPEN FCOE INITIATIVE WITH A 10GB ETHERNET NIC ADVANTAGES
<ul style="list-style-type: none">• On-chip processing reduces CPU workload• Solaris field-proven Fibre Channel driver stack helps ensure data integrity with built-in failover, load balancing, and persistent binding• High performance fully tested FCoE solution supported by Oracle Service• Complete heterogeneous operating support including Oracle Solaris, which is available on both SPARC and x86 platforms, Oracle Linux, RedHat Linux, SUSE Linux, Windows, VMware, and Oracle VM• Concurrent support for FCoE and NIC• Enterprise-class reliability and error recovery	<ul style="list-style-type: none">• 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 login 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 that leverage 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 Oracle VM Server Virtualization, VMware®, Hyper-V®, or other hypervisors that are being used for consolidation of hardware. It's true today that the very first FCoE implementations are 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 first is increased density of applications per physical server, and the second is 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 hyper-threading across many virtual machines, hindering the applications 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. Instead, these requests can be processed with an offload adapter such as the Oracle Sun Storage 10GbE FCoE 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 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 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 that has offload capabilities.

TABLE 2. ENTERPRISE REQUIREMENTS FOR FCOE

ENTERPRISE REQUIREMENT	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 people often 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 for managing FCoE transmissions of 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 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.

Oracle 10GbE FCoE 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 more applications per server and extend the useful life of servers. Additional cost factors should be considered as well, and have been summarized in the following sections.

Operational Efficiencies

Moving to a converged infrastructure can provide significant advantages, but relying on server resources to drive a protocol processing requests is a less than optimal solution. If users believe that some applications will require a higher I/O load, and if these applications are to share the IT infrastructure with the rest of applications, moving to a converged infrastructure can provide definite advantages. However, even 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 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 utilizing Oracle's Sun Storage 10GbE FCoE Converged Network Adapter leaves plenty of room for scaling to 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 and FCoE 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. Oracle's Sun Storage 10GbE FCoE Converged Network Adapter offloads all protocol processing (FCoE, TCP, and SCSI digest for data integrity) onto the adapter. Therefore, using an Oracle's Sun Storage 10GbE FCoE Converged Network Adapter will result in reduced bottlenecks and increased throughput of application data across the PCI bus.

Hidden Cost

If an organization spends \$3,000 on a typical server with one CPU core, how much of that cost is used up by processing I/O requests with a software initiator? As discussed above, one-third of the CPU is used to process FCoE requests for one NIC, meaning that the organization is actually paying \$1,000 to process FCoE traffic.

A slightly more complex example would be to assume the use of next-generation servers running a dual socket, quad core CPU with server virtualization. Today, typical virtual operation environments (VOE) are estimated to run approximately four VMs per CPU core (IDC reports indicate that the industry is quickly moving toward 10 – 12 VMs per core). Even with only four VMs per CPU, 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 an Oracle 10GbE FCoE 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 be attractive. However, it is important not to forget about the opportunity costs that come from not having offload functions. 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 Oracle's Sun Storage 10GbE FCoE 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 help ensure that data is being delivered accurately and that data integrity is maintained when the data 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 help ensure data integrity. 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. An offload engine helps ensure the highest performance of the Converged Network Adapter and maintains the highest level of data integrity as in 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 of these issues will become a significant burden on the CPU and may cripple overall performance on a 10Gb Ethernet network. Oracle's FCoE 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, helping to ensure 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 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 Oracle Approach

While a wide range of data centers will benefit from FCoE implementations, the majority of applications that move a significant amount of data will be running FCoE over Converged Network Adapters with offload engines. In addition to the choices of software initiators and offload engines, there are other engine solutions available in the industry. The information in the sections that follow is relevant both for deciding whether to implement a software initiator and for deciding which offload engine to implement.

Battle-Hardened Enterprise Fibre Channel Driver Stack

Oracle's Sun Storage 10GbE FCoE Converged Network Adapters leverage existing Fibre Channel drivers, making the Sun Storage Converged Network Adapter the most dependable and battle-hardened adapter on the market. Leveraging existing drivers enables Oracle to offer years of proven field experience, combined with full hardware offload for FCoE protocol processing. The Oracle's Sun Storage 10GbE FCoE 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.

Enterprise-quality Reliability and Support

Reliability and support should be top considerations when deciding on an FCoE purchase. Oracle 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, it is wise to consider going to turn for support while software initiators are playing catch up? Oracle 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.

For accountability when implementing an FCoE solution, Oracle's enterprise-quality and reliable converged network adapters are vastly superior to software initiators. Without a history of reliable products and storage support, it is not surprising that early customers are not accepting anything less than the enterprise-quality reliability and support that can be offered by a veteran company such as Oracle.

Summary and Conclusion

When deciding on an FCoE solution, there are two implementation choices, software initiators or offload engines as used in the Oracle's Sun Storage 10GbE FCoE Converged Network Adapter. Both deliver the desired functionality, but only the Converged Network Adapter approach 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.

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 Oracle's Sun Storage 10GbE FCoE Converged Network Adapter, the Fibre Channel driver stack is processed by the adapter. By offloading Fibre Channel processing results, more CPU resources can be made available to support more VMs. This ultimately provides greater cost savings and ROI for virtualized environments. The Sun Storage 10GbE FCoE 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. Oracle's Sun Storage 10GbE FCoE Converged Network Adapter supports 10Gb Ethernet NIC features with stateless offload, further allowing the host CPU to be used for application and virtualization scaling.

Along with scalability, performance, and offload feature advantages over software initiators, Oracle Converged Network Adapters also offer operational efficiencies. They eliminate PCI bus bottlenecks

and offer robust data integrity through 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 and information systems that process large files will require solutions that free up CPU resources to 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 Oracle. Meanwhile, the Oracle's Sun Storage 10GbE FCoE Converged Network Adapter is uniquely positioned to overcome the existing barriers that can be experienced with a software initiator.



The Value of Full Hardware Offload in a
Converged Ethernet Environment

May 2012

Author: Michael Wojnar

Oracle Corporation
World Headquarters
500 Oracle Parkway
Redwood Shores, CA 94065
U.S.A.

Worldwide Inquiries:
Phone: +1.650.506.7000
Fax: +1.650.506.7200

oracle.com



Oracle is committed to developing practices and products that help protect the environment

Copyright © 2011, Oracle and/or its affiliates. All rights reserved. This document is provided for information purposes only and the contents hereof are subject to change without notice. This document is not warranted to be error-free, nor subject to any other warranties or conditions, whether expressed orally or implied in law, including implied warranties and conditions of merchantability or fitness for a particular purpose. We specifically disclaim any liability with respect to this document and no contractual obligations are formed either directly or indirectly by this document. This document may not be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose, without our prior written permission.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. UNIX is a registered trademark licensed through X/Open Company, Ltd. 0512

Hardware and Software, Engineered to Work Together