

Intel and VMware: Enabling Open FCoE in VMware vSphere™ 5

EXECUTIVE SUMMARY

Unified networking over 10 Gigabit Ethernet (10GbE) in the data center offers compelling benefits, including a simplified infrastructure, lower equipment and power costs, and the flexibility to meet the needs of the evolving, virtualized data center. In recent years, growth in Ethernet-based storage has surpassed that of storage-specific fabrics, driven in large part by the increase in server virtualization. The ubiquity of storage area network (SAN) access will be critical as virtualization deployments continue to grow and on-demand, cloud-based data center models emerge.

This paper explores the many benefits of unified networking, the approaches to enabling it, and the pivotal roles Intel and VMware are playing in helping to bring important, consolidation-driving technologies to enterprise data center customers.

White Paper provided by:

vmware®



SIMPLIFYING THE NETWORK WITH 10GbE

As IT departments look to reduce costs and improve server efficiency, they are turning increasingly to server virtualization and consolidation. The benefits of virtualization are widely acknowledged:

- Lower capital and operational expenditures through physical server to virtual machine (VM) consolidation
- Increased availability and responsiveness
- A more dynamic, flexible service model at lower costs

Today's servers are based on powerful new processors, including the Intel® Xeon® processor 5600 and E7 families, that support more VMs per physical host than ever before, helping IT realize greater consolidation ratios.

- The latest generation of Intel® Xeon® processors enables IT to consolidate servers at a 15:1 ratio, delivering power savings of up to 90 percent and a five-month return on investment.
- New four-socket processors are delivering 20 times the performance of previous-generation processors.
- Nearly 50 percent of the four-socket servers shipped today are being used for virtualization.¹

With VMware vSphere™ 5 organizations can create and run "super-VMs" for even the most resource-intensive applications. vSphere 5 VMs support up to 1 terabyte of memory and can enable up to 32 virtual CPUs. As for network throughput, vSphere 5 can support in excess of 1M I/O operations per second. These capabilities surpass the requirements of even the most resource-intensive applications available today, so there are many technical reasons to virtualize any workload on vSphere.

As VM density and use of advanced virtualization features increase, a physical server's networking needs also increase, adding to both cost and complexity. A typical virtualized server uses eight to 10 GbE local area network (LAN) ports and two dedicated SAN ports.

10GbE and unified networking allow IT departments to simplify server connectivity. Consolidating the traffic of multiple GbE connections onto a single 10GbE adapter significantly reduces cable and infrastructure complexity and overall TCO. Recent enhancements to the Ethernet standard enable 10GbE support for both LAN and SAN traffic, allowing IT to realize further benefits by converging data and storage infrastructures. Thanks to its ubiquity, cost effectiveness, flexibility, and ease of use, Ethernet has emerged as the unified data center fabric.

EVOLVING WITH THE DATA CENTER

The growth in server virtualization has helped data center networks evolve from discrete, siloed infrastructures to more flexible fabrics with the scalability and agility necessary to address the needs of new usage models and provide an excellent foundation for enterprise cloud computing.

Over 2.5 billion users will connect to the Internet in the next five years² with over 10 billion devices.³ This usage will require eight times the amount of storage capacity, 16 times the network capacity, and over 20 times the compute capacity by 2015.⁴ A new infrastructure must emerge to power this growth and enable the most efficient use of resources; this infrastructure is cloud computing. The cloud is an evolution of computing that delivers services over the Internet to consumers and enterprises. Services scale—as needed and only when

needed—without user intervention. Highly scalable and efficient cloud architecture is needed to provide both the technical attributes and the extreme resource utilization and efficiency cloud computing promises.

With its reduced hardware requirements, fewer points of management, and broad ecosystem support, 10GbE delivers the flexible, simplified network infrastructure needed to support cloud computing. These characteristics make 10GbE the ideal fabric for cloud infrastructures:

- **Ubiquity.** Ethernet connectivity ships standard on nearly every server today, and Ethernet infrastructures are a universal data center component. 10GbE products have been available for several years, but when 10GbE LAN on motherboard (LOM) connections are integrated in the next generation of servers, unified LAN and SAN connectivity will be available by default.
- **Advanced virtualization support.** Advanced server virtualization enables dynamic resource allocation and is required for most cloud computing infrastructures. Technologies from companies such as Intel and VMware are delivering near-native 10GbE throughput in virtualized environments.
- **Unified networking.** A 10GbE unified fabric simplifies the network infrastructure by consolidating LAN and SAN traffic. Internet small computer system interface (iSCSI) and network file system (NFS) are storage protocols that use Ethernet, and the recent ratification of the Fibre Channel over Ethernet (FCoE) standard extends this capability. In addition, recent Ethernet enhancements ensure quality of service (QoS) for critical traffic.

Intel and VMware are two companies helping to advance the shift to 10GbE-based unified networking in the data center. The latest Intel® Ethernet 10 Gigabit controller and server adapters offer virtualization optimizations and advanced unified networking features. These include optimizations for lossless Ethernet; intelligent, hardware-based accelerations for FCoE and iSCSI; and support for Open FCoE, which is discussed later in the paper.

VMware vSphere 5 is compatible with the widest range of industry-standard x86 servers, storage solutions, and network equipment so you can virtualize your business critical applications with confidence.

These solutions help to simplify network connectivity for today's virtualized servers and lay the foundation for the next-generation data center.

The Promise of Ethernet Storage

New usage models and the explosive growth of data in their organizations have forced IT administrators to deal with complicated technical and business challenges. Today, most IT departments deploy separate LANs and SANs, with storage often divided between network-attached storage (NAS) and SAN, which require multiple data recovery solutions, a variety of data management models, and, potentially, different support teams. The goal of unified networking is to allow a single fabric—Ethernet—to carry these disparate traffic types.

Ethernet has served as a unified data center fabric for years, supporting LAN, NAS (NFS, common Internet file system (CIFS)) and iSCSI SAN traffic. With recent Ethernet enhancements and the ratification of the FCoE specification, standard Ethernet adapters can now connect servers to Fibre Channel (FC) SANs. Extending Ethernet's inherent

advantages, including proven reliability, ubiquity, and wide familiarity, to FC SAN traffic will help accelerate the move to 10GbE-based I/O consolidation in virtualized data centers, reduce costs, and improve simplification and agility.

Given its flexibility and long history, it is not surprising that Ethernet storage is the fastest growing segment of the storage systems market. According to the industry research firm IDC, Ethernet-based storage unit shipments surpassed FC storage shipments in 2008,⁵ with growth driven largely by broad iSCSI adoption in Microsoft Windows*, virtual server, and blade server environments. IDC projects that in 2011, Ethernet-based storage will represent the majority of all new external storage capacity with 53 percent⁶ of total worldwide capacity.

As shown in figure 1, growth in Ethernet-based storage capacity will continue to increase, widening the gap between it and FC storage. This growth will be due to increasing deployment of "Ethernet only" data centers (which use a unified

10 Gigabit Ethernet infrastructure for all data and storage traffic), the emergence of cloud computing, and FCoE solutions entering the mainstream.

Ethernet Enhancements for Storage

Data Center Bridging for Lossless Ethernet

To strengthen 10GbE as a unified data center fabric, the IEEE has developed and ratified standards for Ethernet enhancements to support storage traffic. These enhancements strengthen 10GbE as a unified data center fabric for running FC over Ethernet and iSCSI. Known collectively as "Data Center Bridging" (DCB), these extensions enable better traffic prioritization over a single interface and an advanced means for shaping traffic on the network to decrease congestion. In short, DCB provides the QoS that delivers a lossless Ethernet fabric for storage traffic. For more information, see the [DCB white paper](http://ethernetalliance.org/files/static_page_files/83AD0BBC-C299-B906-8F5985957E3327AA/Data%20Center%20Bridging.pdf) from the Ethernet Alliance. (http://ethernetalliance.org/files/static_page_files/83AD0BBC-C299-B906-8F5985957E3327AA/Data%20Center%20Bridging.pdf)




Figure 1. Growth in Worldwide Storage Capacity of Ethernet-based Storage vs. Fibre Channel

Fibre Channel over Ethernet: Extending Consolidation

FCoE is a logical extension of Ethernet that uses FC's Network, Service, and Protocol layers to carry data packets over Ethernet's physical and data link layers. Using FC's upper layers smoothes the transition to FCoE because existing SAN-based applications do not need to change to benefit from the performance and cost benefits of FCoE.

Many enterprises have extensive FC installations, and FCoE provides easy FC SAN access for any server with a capable 10GbE port. By using standard Ethernet fabrics, FCoE eliminates the need for dedicated FC host bus adapters (HBAs), reducing cabling and switch-port requirements, while coexisting with existing FC hardware and software infrastructures. The result is a simplified data center infrastructure, lower equipment and power costs, and universal SAN connectivity across the data center over the trusted Ethernet fabric.

Introducing Open FCoE

The Open FCoE approach consists of standard 10 Gigabit Ethernet adapters and native OS- or hypervisor-based FCoE initiators, which together provide a robust, scalable, and high-performance server connectivity option without expensive, proprietary hardware. As shown in Figure 2, Open FCoE implements the complete FC protocol in the hypervisor system kernel. It provides libraries for different system-level implementations, allowing vendors to implement data plane functions of the FCoE stack in hardware to deliver optimum performance.

Open FCoE: Momentum in the Linux® Community

Initiated by Intel, the Open FCoE project was accepted by the Linux community in November 2007 with the goal of accelerating development of a native FCoE initiator in the Linux kernel. The industry responded enthusiastically, and today there are over 190 active participants in the community who are contributing code, providing review comments, and testing the Open FCoE stack. To date, the Open FCoE Source Web site (www.open-fcoe.org) has received over 20,000 hits. Open industry standards and Open Source play a significant role in the modern data center, as they lower research and development costs and enable access to a multi-vendor supply chain that is designed for heterogeneous interoperability, ultimately resulting in greater choice and lower equipment costs.

The Open FCoE approach offers a number of advantages in terms of accessibility and ease of use.

- **Accessibility.** The Open FCoE approach makes FC SAN access available to any server with a standard, compatible 10GbE adapter installed. Rather than having to purchase proprietary hardware, IT can continue using its current Ethernet adapters to gain this additional functionality. With FCoE support integrated into the hypervisor, IT can select compatible Ethernet adapters from multiple vendors, helping to avoid vendor lock-in. With Open FCoE, FC SAN access will become universal once 10GbE LOM technology becomes pervasive.




Figure 2. Overview of Open Fibre Channel over Ethernet initiator solution.

- **Ease of use.** Because the Open FCoE approach uses standard 10GbE adapters, IT can leverage existing knowledge to configure and manage these adapters for FCoE deployments. In fact, IT can standardize on a single product or product family for all LAN and SAN connectivity. FCoE initiator integration into the hypervisor also means common OS-based tool support across a product family or even adapters from multiple vendors. Using converged network adapters (CNAs) requires learning new tools and installing additional, proprietary software. Use of CNAs may also require using different products for LAN and SAN connectivity because CNAs use the Ethernet fabric but are optimized for storage functions.

Native Initiator Success: iSCSI

iSCSI provides an excellent example of the success of native storage initiators integrated into the OS.

In the early days of iSCSI, proponents of iSCSI HBAs claimed that these dedicated adapters were necessary to deliver acceptable performance. Much like CNAs, iSCSI HBAs offload storage protocol processing to a separate processor on the adapter, rather than allowing the host processor and OS to handle these tasks.

While early HBAs did offer some benefits in terms of throughput, they did so at levels beyond what was needed in real-world performance. They also suffered from many of the same issues as today's CNAs, including high costs, vendor lock-in, reliability, and stability.

Today, all major server OSs and hypervisors include native iSCSI support, delivering the same benefits as described above for Open FCoE. The maturity of native iSCSI initiators along with advanced adapter and platform features now provides enterprise-level performance with standard Ethernet products.

IT departments standardizing on Intel® Ethernet Server Adapters for iSCSI connectivity are able to use a single initiator, TCP/IP stack, and set of management tools and IT policies. This standardization can deliver easier server provisioning, lower the likelihood of human error, and simplify management while reducing capital and operational expenditures. With FCoE initiators integrated into leading hypervisors and OSs and 10GbE LOM implementations coming soon, Intel expects Open FCoE-based solutions to deliver these same benefits.

INTEL ETHERNET UNIFIED NETWORKING

As discussed earlier, Ethernet-based storage has been growing at a faster rate than traditional FC networks, particularly because Ethernet is a trusted, widely deployed, and well-understood technology. Ethernet is the foundation of unified networking, and with 30 years of experience delivering quality Ethernet products, Intel is uniquely positioned to drive the transition to 10GbE unified networking.

Reliability

Intel is the volume leader in Ethernet adapter shipments and has led the industry through speed transitions and enhancements to the Ethernet standard, including iSCSI, DCB, and FCoE. Intel Ethernet products have the broadest OS support in the industry, and Intel's long-term product roadmaps align future products with new server platform capabilities and upcoming data center trends, including cloud computing. Intel has expanded the capabilities of its trusted Ethernet product line to incorporate FCoE for true unified networking. By contrast, CNA vendors are acquiring Ethernet technology and adding it to their existing FC-centric products.

Cost Effectiveness

Standardizing on Intel® Ethernet 10 Gigabit Server Adapters takes advantage of Ethernet economics to deliver cost-effective, broad-based unified networking deployment. A single Intel Ethernet 10 Gigabit Server Adapter supports FCoE, iSCSI, NAS, and LAN traffic for true unified networking without the need for additional, expensive hardware or upgrades.

Scalable Performance

In terms of performance, Intel Ethernet 10 Gigabit Server Adapters combine high throughput, intelligent hardware-based offloads with native OS initiators, and stable hardware, while leaving processor cycles available for application processing. These adapters offload the main data paths to improve FCoE throughput. The net result is comparable FCoE performance for real-world workloads compared to that of CNAs.

In real-world implementations, SAN performance is determined by a number of factors, including application threading, storage target processor speed, and disk speeds. At typical input/output operations per second (IOPS) levels, processor utilization is typically quite low, especially in systems with the latest multi-core processors and memory controllers.

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

Ease of Use

By using the native storage initiators integrated into the OS, Intel Ethernet 10 Gigabit Server Adapters make it easy to connect any server to the SAN without the need for complicated proprietary solutions. Standardizing on Intel Ethernet 10 Gigabit Server Adapters delivers a number of benefits to enterprise IT:

- Trusted and validated OS-based and hypervisor-based initiators
- Single interface to configure an adapter for LAN or SAN traffic
- Plug-and-Play with existing management tools
- No proprietary software to install

FCoE CNAs, on the other hand, require vendor lock-in. The management interfaces, APIs, and drivers of each FCoE CNA brand typically differ, which adds complexity.

CONCLUSION

10GbE unified networking provides a simple, flexible, and well-understood fabric for today's virtualized data centers and lays the groundwork for new computing models, including cloud computing, which will deliver more intelligent, responsive data centers and greater business agility. With support for unified networking, including Open FCoE, Intel and VMware are two companies that are helping IT organizations transition to a common data center fabric on 10GbE that will simplify the infrastructure while reducing costs.

FOR MORE INFORMATION

www.intel.com/go/unifiednetworking

www.vmware.com/technical-resources/virtual-networking/index.html

¹ Source: Intel estimates as of January 2010. Performance comparison using SPECjbb*2005 bops (business operations per second). Results have been estimated based on internal Intel analysis and are provided for informational purposes only. Any difference in system hardware or software design or configuration may affect actual performance.

² IDC, "Server Workloads Forecast," 2009.

³ IDC, "The Internet Reaches Late Adolescence," December 2009.

⁴ 8x Network: 800 terabytes per second of IP traffic estimated on internal Intel analysis "Network Supply/Demand 2010-2020" forecast; 16x Storage: 60 exabytes of data stored from Barclays Capital "Storage Bits" September 2009, extrapolation by Intel for 2015; 20x Compute: Intel internal long-range planning forecast. Extrapolated to one billion virtual servers using one virtual machine per core.

⁵ IDC WW Storage Systems Tracker, December 2009.

⁶ Worldwide Enterprise Storage Systems 2011-2015

Forecast: "Emerging" Once Again is a Keyword in the Storage Market, IDC #228255, May 2011

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 www.intel.com/performance/resources/limits.htm

© 2011 Intel Corporation. All rights reserved. Intel, the Intel logo, and Xeon are trademarks of Intel Corporation in the U.S. and other countries.

*Other names and brands may be claimed as the property of others.