

Simplifying the Network with 10 Gigabit Ethernet Unified Networking

An Intel and EMC Perspective

Executive Summary

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. Unified networking over 10 Gigabit Ethernet (10GbE) in the data center offers compelling benefits for the evolving virtualized data center, including a simplified infrastructure, lower equipment and power costs, and greater flexibility. Ethernet is the foundation of Cloud Computing, as easy and simple access to the storage area network (SAN) will be critical as virtualization deployments continue to grow and new, on-demand data center models emerge.

Intel and EMC are two leaders in the transition to 10GbE unified networking. This paper explores the benefits of unified networking, approaches to enabling it, and how Intel and EMC are helping to bring important, consolidation-driving technologies to enterprise data center customers.

Table of Contents

Simplifying the Network with 10GbE	3
10GbE: The Fabric of Cloud Computing	3
The Promise of Ethernet Storage	4
Ethernet Enhancements for Storage	5
Data Center Bridging for Lossless Ethernet.....	5
Fibre Channel over Ethernet (FCoE):	
Enabling End-to-End Unified I/O	5
Introducing the Open FCoE Architecture	5
Open FCoE: Momentum in the Linux* Community	5
Native Initiator Success: iSCSI	6
Intel® Ethernet Unified Networking.....	6
Broad Support Commitment.....	6
Cost-Effective Efficiency.....	7
Ease of Use	7
EMC's Unified Storage Solutions.....	7
Conclusion	7

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 understood: less server hardware to purchase, lower power and cooling needs, and centralized management. Today's servers are based on powerful new processors, including the Intel® Xeon® processor 5600 and 7500 series, that support more virtual machines (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,¹ delivering power savings of up to 90 percent² and a five-month return on investment.³
- New four-socket processors are delivering 20 times the performance of previous-generation processors.⁴
- Nearly 50 percent of the four-socket servers shipped today are being used for virtualization.⁵

Unfortunately, the success achieved by many organizations in attaining these benefits has been hindered by past practices for networking virtualized servers. As VM density increases, a physical server's networking needs also increase, adding both cost and complexity. A typical virtualized server contains eight to ten Gigabit Ethernet (GbE) LAN ports and two dedicated Storage Area Network (SAN) ports.

As server virtualization continues to grow, 10GbE unified networking is simplifying server connectivity. Consolidating the traffic of multiple GbE connections onto a single 10GbE adapter significantly reduces cable and infrastructure complexity and overall TCO. Enhancements to the Ethernet standard also 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.



Figure 1. Multiple GbE connections in a typical virtualized server

10GbE: The Fabric of Cloud Computing

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 years⁶ with over 10 billion devices.⁷ This usage will require eight times the amount of storage capacity, 16 times the network capacity, and over 20 times the compute capacity by 2015.⁸ A new infrastructure must emerge to power this growth and enable the most efficient use of resources; this 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. A 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 key 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. 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 any Cloud Computing infrastructure. Technologies from companies such as Intel, VMware, and Microsoft are delivering line-rate 10GbE throughput and support for platform virtualization enhancements.



Figure 2. Simplified server connectivity using 10GbE

Unified networking: A 10GbE unified fabric simplifies the network infrastructure by consolidating LAN and SAN traffic. Recent Ethernet enhancements ensure quality of service (QoS) for critical traffic.

Intel and EMC are companies leading the shift to 10GbE-based unified networking in the data center. The latest Intel® Ethernet 10 Gigabit controllers and server adapters include virtualization optimizations and advanced unified networking features, including optimizations for lossless Ethernet, intelligent, hardware-based acceleration for Fibre Channel over Ethernet (FCoE) and iSCSI, and support for Open FCoE, which is discussed later in this paper.

EMC provides services to plan, design, implement, and maintain Ethernet-based solutions. In addition, EMC Select* includes a complete line of 10GbE switching products from Brocade and Cisco.

These solutions 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 LAN and storage networks, with storage often divided between network attached storage (NAS) for file-based applications and SAN (Fibre Channel and iSCSI) for

block-based applications. The goal of unified networking is to allow a single-fabric infrastructure, based on 10GbE, to carry all of 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, Data Center Bridging (DCB)-based Ethernet adapters can now facilitate connecting servers to Fibre Channel (FC) SANs. Extending Ethernet's ubiquity and wide familiarity to FC SAN traffic will help to 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. The industry research firm IDC estimates that the worldwide Ethernet-based storage systems (NAS and iSCSI SAN) market grew at a compounded annual growth rate (CAGR) of approximately 23 percent between 2005 and 2009.⁹ iSCSI storage shipments, in particular, experienced the highest growth rates (70 percent) during 2005-2009, driven by broad iSCSI adoption in Windows,* virtual server, and blade server environments.

Industry analysts project continued gains in the Ethernet storage market share due to increasing deployment of "Ethernet only" data centers (which use a unified 10GbE infrastructure for all data and storage traffic), the emergence of Cloud Computing, and as FCoE solutions enter the mainstream.

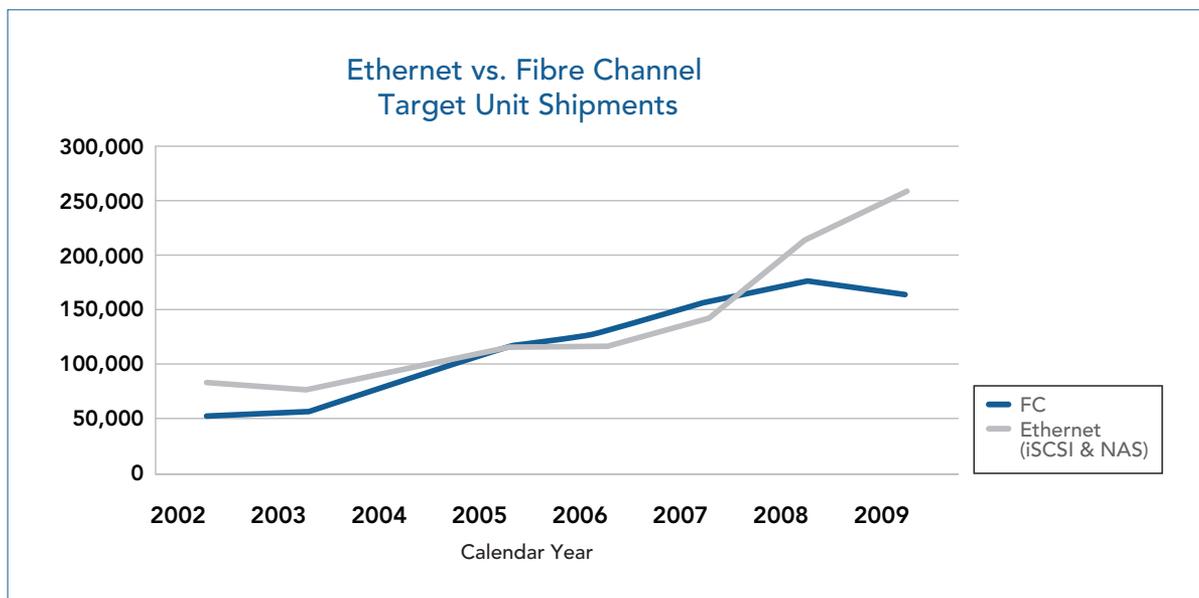


Figure 3. Actual unit shipments for Ethernet and Fibre Channel storage (IDC, 2009)

Ethernet Enhancements for Storage

Data Center Bridging for Lossless Ethernet

To strengthen 10GbE as a unified data center fabric, the IEEE has standardized on Ethernet enhancements to support storage traffic, including FCoE and iSCSI. These extensions, known collectively as Data Center Bridging (DCB), enable better traffic prioritization over a single interface, as well as advanced means for shaping traffic on the network to decrease congestion. In short, DCB delivers a lossless Ethernet fabric for storage traffic.

Fibre Channel over Ethernet (FCoE): Enabling End-to-End Unified I/O

Fibre Channel over Ethernet (FCoE) is a logical extension of Ethernet that uses Fibre Channel's Network, Service, and Protocol layers to carry data packets over Ethernet's physical and data link layers. Fibre Channel's unique network-centric management model has proven administrative capacities to scale to thousands of end nodes in a data center. Using Fibre Channel'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. The provisioning responsibilities, now split between the server and storage administrators, can be maintained in the transition from Fibre Channel SANs to Ethernet-based FCoE-powered SANs.

Many Enterprises have extensive FC installations, and the availability of software-based FCoE initiators makes the FC SAN easily accessible for any server with a 10GbE Ethernet port. By allowing FC to use DCB-capable Ethernet, FCoE eliminates the need for dedicated Fibre Channel host bus adapters (HBAs), reducing cabling and switch-port requirements, while coexisting with existing Fibre Channel 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 the Open FCoE Architecture

The Open FCoE approach consists of standard 10 GbE adapters and native operating system-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 operating 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

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

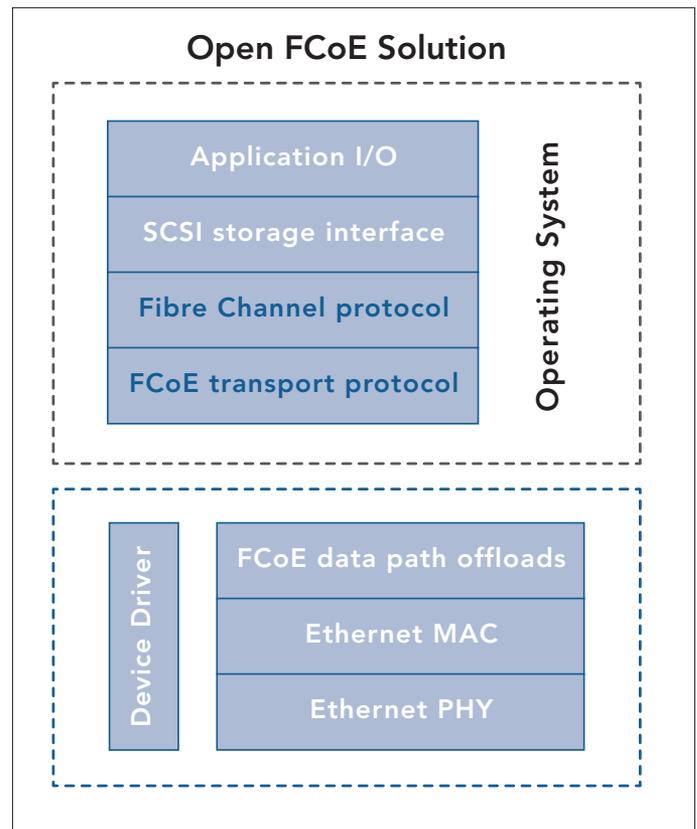


Figure 4. Overview of the Open FCoE initiator solutions

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 R&D investment 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 DCB-capable 10GbE adapter installed. IT departments can standardize on these adapters to simplify server connectivity across multiple use cases. With FCoE support integrated into the operating system, FC SAN access will become readily available across these converged networks and even more accessible once 10GbE LOM technology grows to pervasive status.

Ease of use: Because the Open FCoE approach uses DCB-capable 10GbE adapters that comply with broad, well-known Ethernet standards, 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 OS also means common OS-based tool support across a product family or even adapters from multiple vendors, as they become available.

Native Initiator Success: iSCSI

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

In the early days of iSCSI, proponents of iSCSI HBAs claimed that these dedicated adapters were necessary to deliver acceptable performance. iSCSI HBAs offload iSCSI processing to a separate processor on the adapter, rather than allowing the host CPU and operating system to handle these tasks.

Today, all major server operating systems include native iSCSI support, delivering the same benefits as detailed above. Native iSCSI initiators have continued to mature and now support advanced adapter features and platform advancements that help deliver excellent performance. In fact, Intel recently demonstrated a standard Intel® Ethernet 10 Gigabit Server Adapter driving 1.25 million IOPS using the native iSCSI initiator in Windows Server* 2008 R2.¹⁰

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. Easier server provisioning, lower likelihood of human error, and simpler management enable lower capital and operational expenditures. That standardization also allows IT to enjoy the benefits of simplified management and integrated tools. Intel expects the same benefits for Open FCoE-based solutions.

Intel® Ethernet Unified Networking

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. Qualification by EMC's E-Lab* means customers can be confident that servers with the Intel® Ethernet Server Adapter X520 installed will interoperate with EMC systems for iSCSI and now FCoE.

Broad Support Commitment

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 operating system 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 is expanding the capabilities of its trusted Ethernet product line to incorporate FCoE for true unified networking.

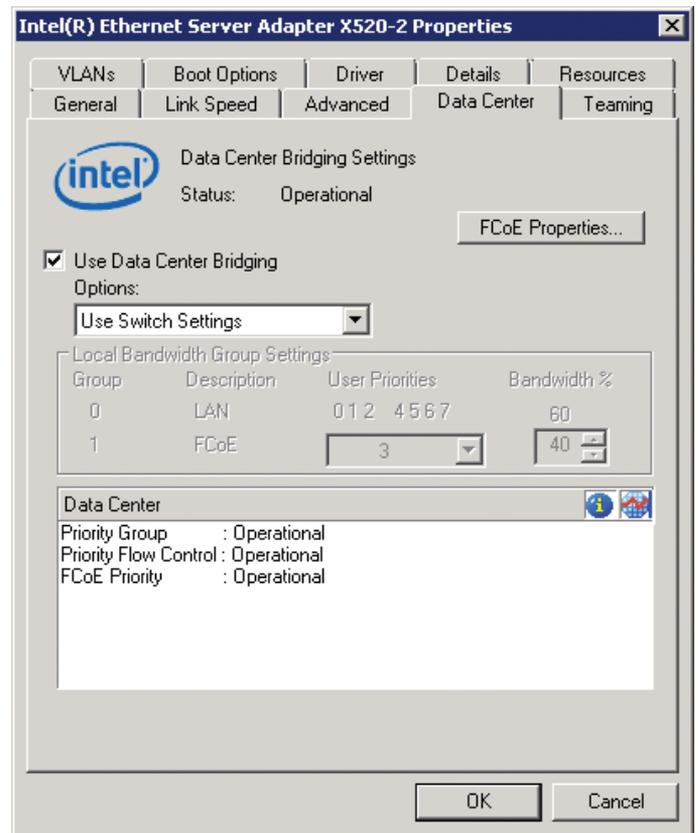
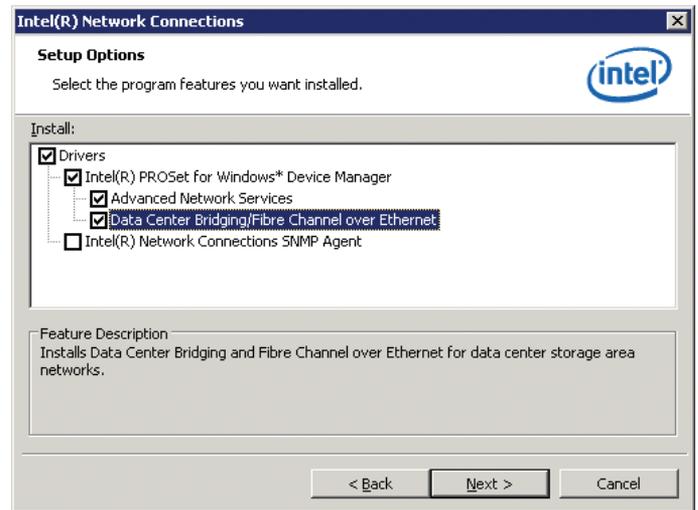


Figure 5. IT administrators use the same interface for configuring LAN and FC SAN traffic

Cost-Effective Efficiency

Standardizing on Intel Ethernet 10 Gigabit Server Adapters leverages Ethernet economics for cost-effective broad-based unified networking deployment. A single Intel Ethernet Server Adapter X520 supports FCoE, iSCSI, NAS, and LAN traffic for true unified networking without the need for expensive upgrades.

In terms of performance, the Intel Ethernet Server Adapter X520 combines 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 excellent FCoE performance for real-world workloads.

In real-world implementations, SAN performance is determined by a number of factors, including application threading, as well as storage target processor and disk speeds. At typical I/O per second (IOPS) levels, processor utilization is typically quite low, thanks in part to recent increases in processor performance.

Ease of Use

By using the native storage initiators integrated into the operating system, the Intel Ethernet Server Adapter X520 makes it easy to connect any server to the SAN. Standardizing on Intel Ethernet Server Adapter X520 family delivers a number of benefits to Enterprise IT:

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

EMC's Storage Solutions

As the leader in networked storage, EMC offers a comprehensive line of storage systems to satisfy a diverse set of customer needs. The EMC VNX* series of products delivers affordable and easy to use mid-tier storage solutions that provide high efficiency for multi-protocol environments. EMC Symmetrix* VMAX delivers unmatched scalability, high availability, and scale-out performance levels for the most demanding needs of the virtual data center. These EMC storage systems all offer simplified management with advanced functionality including virtual provisioning, fully automated storage tiering (FAST), remote and local replication, and data compression. EMC's storage systems support native FCoE. These high-speed FCoE I/O modules can easily be installed into EMC arrays non-disruptively.

Conclusion

10 Gigabit Ethernet's simplicity, flexibility, and familiarity make it ideal for a unified data center network fabric and for laying the groundwork for new models, including Cloud Computing, which will deliver more intelligent, responsive data centers and greater business agility. Through their product and industry leadership, Intel and EMC are helping IT organizations transition to 10GbE and build next-generation data center infrastructures.



¹ Intel consolidation based on replacing 15 four-year-old single-core Intel® Xeon® processor 3.8 GHz with 2M cache-based servers with one new Intel Xeon processor X5670-based server while maintaining performance as measured by SPECjbb2005* business operations per second. 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, visit www.intel.com/performance/server.

Baseline platform: Intel server platform with two 64-bit Intel Xeon processor 3.80 GHz with 2M L2 Cache, 800 FSB, 8x1 GB DDR2-400 memory, 1 hard drive, 1 power supply, Microsoft Windows® Server 2003 Ent. SP1, Oracle JRockit® build P27.4.0-windows-x86_64 run with 2 JVM instances.

New platform: Intel server platform with two six-core Intel® Xeon® processor X5670, 2.93 GHz, 12MB L3 cache, 6.4QPI, 12 GB memory (6x2 GB DDR3-1333), 1 hard drive, 1 power supply, Microsoft Windows Server 2008 64 bit SP2, Oracle JRockit® build P28.0.0-29 run with 2 JVM instances.

² Intel comparison replacing 20 5-year-old single-core Intel® Xeon® processor 3.33 GHz-based servers with one new Intel Xeon processor X7560-based server. Costs have been estimated based on internal Intel analysis and are provided for informational purposes only.

³ Based on comparison between 2S Single Core Intel® Xeon® 3.80 with 2M L2 Cache and 2S Intel® Xeon® X5670-based servers. Calculation includes analysis based on performance, power, cooling, electricity rates, operating system annual license costs and estimated server costs. This assumes 8kW racks, \$0.10 per kWh, cooling costs are 2x the server power consumption costs, operating system license cost of \$900/year per server, per server cost of \$7200 based on estimated list prices and estimated server utilization rates. All dollar figures are approximate. Performance and power comparisons are based on estimated server side java benchmark results (Intel Corporation Jan 2010). Platform power was estimated during the steady state window of the benchmark run and at idle. Performance gain compared to baseline was 15x.

Baseline platform: Intel server platform with two 64-bit Intel Xeon processor 3.80 GHz with 2M L2 Cache, 800 FSB, 8x1GB DDR2-400 memory, 1 hard drive, 1 power supply, Microsoft Windows® Server 2003 Ent. SP1, Oracle® JRockit® build P27.4.0-windows-x86_64 run with 2 JVM instances.

New platform: Intel server platform with two six-core Intel® Xeon® processor X5670, 2.93 GHz, 12MB L3 cache, 6.4QPI, 12 GB memory (6x2 GB DDR3-1333), 1 hard drive, 1 power supply, Microsoft Windows Server 2008 64 bit SP2, Oracle JRockit® build P28.0.0-29 run with 2 JVM instances.

⁴ Intel performance comparison using SPECjbb2005* business operations per second between 5-year-old single-core Intel® Xeon® processor 3.33 GHz-based servers and one new Intel Xeon processor X7560-based server. 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, visit www.intel.com/performance/server.

⁵ 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. For detailed calculations, configurations and assumptions refer to the legal information slide in backup.

⁶ IDC, "Server Workloads Forecast," 2009.

⁷ IDC, "The Internet Reaches Late Adolescence," December 2009.

⁸ 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.

⁹ IDC WW Storage Systems Tracker, December 2009.

¹⁰ Test Configuration: Iometer v. 2006.7.27, number of managers = 1, number of workers/manager = 30, (total number of workers) = 30, number of LUNs = 30, number of outstanding I/Os = 50, IO size = 512B. 10 iSCSI targets with 3 LUNs per target. Target is StarWind Enterprise configured with RAM disk. SUT: Supermicro 6026T-NTR+, Intel® Xeon® processor W 5680 (12 M Cache, 3.33 GHz, 6.40 GT/s Intel® QPI), 24 GB DDR3, Windows Server® 2008 R2 x64 Network Configuration: Cisco Nexus® 5020, Intel 82599EB 10 Gigabit Ethernet Controller connected @ 10 Gbps.

INTEL MAKES NO WARRANTIES, EXPRESS OR IMPLIED, IN THIS DOCUMENT. This document is for informational purposes only.

Copyright © 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.