WHITE PAPER

Driving 10 Gigabit Ethernet Adoption in the Data Center

INTRODUCTION

Data volumes and network bandwidth consumed by data centers double every 18 months while devices accessing networks double every 2.5 years. To support this thirst for data, a corresponding advance in network bandwidth needs to occur. Given that Ethernet has become the networking infrastructure of not only the Internet, but also information service delivery to corporations of all sizes, the logical answer is migration to the next advance in Ethernet—10 Gigabit Ethernet.

What's driving the evolution to 10 Gigabit Ethernet? The drive comes from advances in other information technologies—multi-core processors, virtualized environments, advances in storage architectures, server clustering and consolidation, new forms of information delivery using the Internet, and the next wave of digital media content. Associated with these technological advances is the advance in networking itself that is making 10 Gigabit Ethernet affordable across the network from the core to the edge.

Intel and Extreme Networks® are two companies leading the evolution to 10 Gigabit Ethernet adoption. Intel® 10 Gigabit Ethernet Server Adapters provide fast and reliable connections for high-performance servers and workstations. These adapters have been designed for multi-core processors, optimized for virtualization, and provide support for unified networking over Ethernet. Extreme Networks provides 10 Gigabit Ethernet networking solutions for layer 2 and layer 3 switching in both modular and fixed configurations. Extreme Networks versatile design for these configurations offers customers a choice while providing an excellent migration path from existing Gigabit Ethernet network infrastructures. By combining 10 Gigabit Ethernet technology from Intel and Extreme Networks, customers can realize the advantages of 10 Gigabit Ethernet today at affordable costs in seamless migration and coexistence with their existing Ethernet data center network.

Motivation Driving 10 Gigabit Ethernet

Today, most server, desktop, and laptop systems come standard with Gigabit Ethernet connectivity. It has become natural for us to just "plug into" a Gigabit Ethernet network. Although 10 Gigabit Ethernet has been around since 2002, its adoption has been limited to the core of data center networks. That is, until now.

Advances in technology delivering 10 Gigabit Ethernet and falling prices are driving the adoption of 10 Gigabit Ethernet from the core to the edge of the network. Let's look at the motivators driving the adoption of 10 Gigabit Ethernet in modern data centers.

MULTI-CORE PROCESSORS

Keeping pace with the processing demands of multi-core and multi-threaded processors requires robust I/O interfaces and devices. Constraining the I/O paths results in the "hour glass" effect, or the metering of I/O requests through a single, narrow interface. As in an hour glass, I/O requests can be constrained by slow interfaces and devices resulting in ineffective utilization of processing resources as they wait for I/O requests to complete.

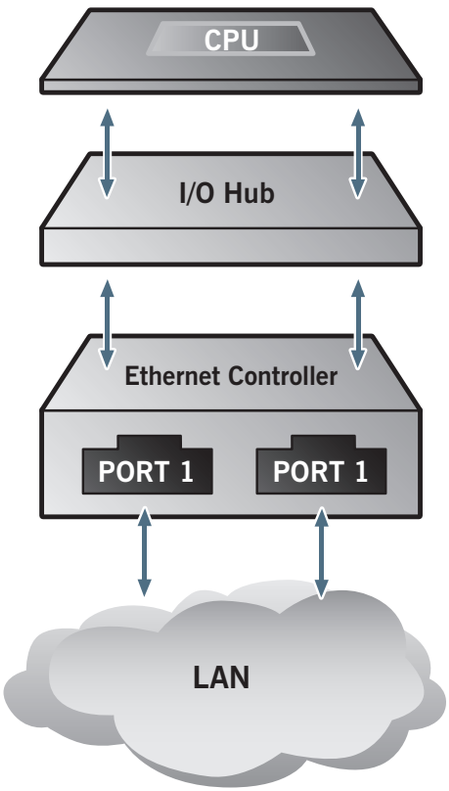

Figure 1: Network data flow in a previous-generation platform

The interface to the network and even the network itself can be one of the leading contributors to limiting processing efficiency in multi-core processor systems. Ten years ago it was typical to find single-core, single-threaded processors being supported by a Gigabit Ethernet network. Today, trying to support eight to sixteen concurrent execution streams supported by multi-core, multi-threaded processors on yesterday's Gigabit Ethernet is not sufficient. To achieve processing efficiency with these advanced multi-core, multi-threaded processors requires network interfaces and the network itself to be exponentially greater—10 Gigabit Ethernet.

VIRTUALIZED ENVIRONMENTS

In today's information technology vernacular, "server virtualization" means the ability to run more than one operating system (OS) image on a single physical server. Server virtualization uses a software hypervisor running as the kernel on a physical server to present multiple virtual machine (VM) images to the guest OSs. In this model, the OSs can be disparate; for example, Microsoft Windows* Server and Linux* running on a VMware hypervisor on an Intel® Xeon® processor-based server. Key benefits to server virtualization are:

- Better utilization of physical server resources (specifically with multi-core processors)
- Improved deployment of applications on virtual servers (you don't need to order another physical server to deploy an application)
- Ability to balance physical server workloads across many virtualized servers

However, along with these benefits come challenges. It is easy to overtax the physical I/O resources in a virtualized server environment. Each virtualized OS thinks it has exclusive use of the physical resources when in reality these resources are shared across the virtualized OSs running concurrently on a physical server. The result is an oversubscription of resources and, in a networking environment, the physical network server adapter. The problem with the demand on the network by a multi-core processor is exacerbated in a virtualized environment. One logical solution is to increase network bandwidth to satisfy this increased virtualized server network demand by deploying 10 Gigabit Ethernet at the network edge.

ETHERNET-BASED STORAGE

Ethernet-based storage offerings provide huge volumes of connected data at attractive prices. As these storage offerings continue to evolve and improve, organizations can adopt strategies to consolidate their storage infrastructure in the data center on an Ethernet network achieving improved economies of scale and simplified management of the infrastructure all delivered at an affordable cost.

Today, Ethernet-based network attached storage (NAS) and storage area network (SAN) systems are being deployed by large and small organizations to realize these benefits. The ability to scale storage delivery performance up to 10 gigabits per second gives these organizations a simple and cohesive migration path for their NAS and SAN storage systems.

SERVER CLUSTERING

Another major trend is data centers clustering together servers to improve performance and application availability. It is not uncommon to hear of physical server deployments in the hundreds consolidated in a single data center. Applications are being rewritten to take advantage of the parallel processing paradigm made available through server clustering to improve not only application performance, but also application availability. Ethernet is becoming the de facto choice for clustering servers in a network. Using inter-processor communication (IPC) over an Ethernet network, applications can deliver higher degrees of parallelization as long as the network throughput delivers the required performance.

The ability to migrate to 10 Gigabit Ethernet clustered interconnect solutions provides customers the ability to scale parallelized applications without the turmoil of moving to a different network technology to realize increased bandwidth, lower latencies, and better performance.

WEB 2.0

It has been estimated that the volume of web traffic will experience a fifty-fold increase by the year 2015. What's driving this web traffic increase? Today, we have the delivery of video, image, and audio content. Eight hours of video content is loaded onto YouTube* every minute. In the near future, online video will transition to high-definition (HD) video, which is seven to ten times more bandwidth-hungry than today's video streams. Corporate delivery of information both externally in describing products and services and internally for training and information delivery will be delivered through HD video and audio streams. To handle these vast bandwidth requirements traffic management appliances will be required to spread out the workload to clustered servers in a server load-balancing configuration.

To support this performance demand requires network performance that is equal to the task. No longer will Gigabit Ethernet be sufficient to respond to the demands of advances in the Web 2.0 Internet. Addressing this challenge will be 10 Gigabit Ethernet solutions that drive through the traffic management appliances directly to the web-facing servers handling Web 2.0's data-intensive requests.

DIGITAL IMAGING AND EDITING

One of the fastest advances in technology has been in the area of digital imaging. By 2010, it is estimated that medical imaging alone will consume 30 percent of the world's data storage. Human interaction with digital images will be through a global interconnected network as productivity gains are realized by effectively using this information form. To move hundreds of gigabytes of digital content to computers in a timely fashion for display and revision will require rethinking the infrastructure and cabling requirements for computers.

COST

When the 10 Gigabit Ethernet specification was released in 2002, the cost of early 10 Gigabit Ethernet offerings exceeded USD 10,000 per port. The only rational use at this price was in the core of the network. Extreme Networks can now deliver 10 Gigabit Ethernet at a cost of USD 500 per port.

With this cost advantage, 10 Gigabit Ethernet will move out of the core and access layers of the network to the network edge. 10 Gigabit Ethernet adoption will further accelerate the market for copper-based adapters and switches making network infrastructure costs even more reasonable.

The Solution

Moving to a 10 Gigabit Ethernet environment has never been easier. With solutions available from Intel and Extreme Networks, deploying 10 Gigabit Ethernet in both copper and fiber infrastructures in the data center is both affordable and efficient. By using Intel 10 Gigabit Ethernet Server Adapters and Extreme Networks' 10 Gigabit Ethernet switching solutions, the reality of achieving 10 Gigabit performance at the edge of the network is obtainable today.

INTEL® ETHERNET 10 GIGABIT SERVER ADAPTERS

Intel has been the leading supplier and innovator of Ethernet adapters for more than 25 years. Intel's new 10 Gigabit Ethernet adapters are designed for multi-core Intel Xeon processor-based systems, optimized for virtualization, and include support for unified networking to ensure fast, power-efficient solutions for a broad range of applications, including virtualized server environments, blade servers, and copper-based infrastructures for volume deployment.

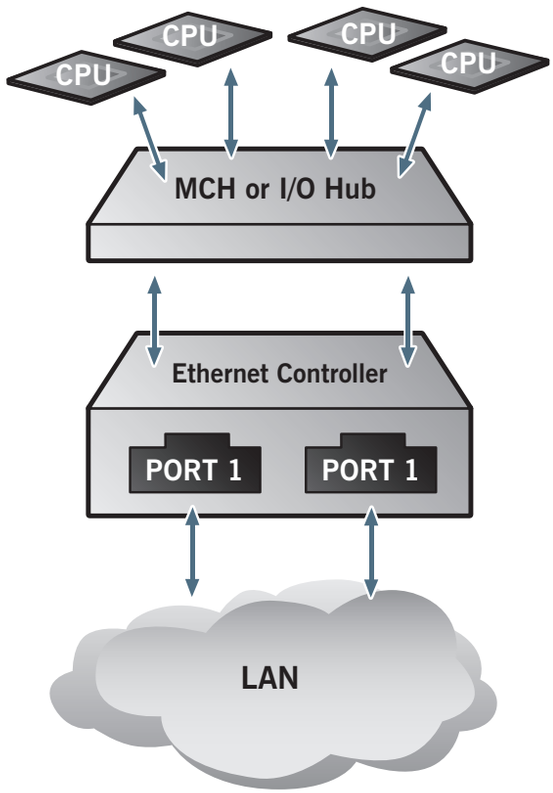

Figure 2: Multiple queues, RSS, MSI-X working together in a multi-core system.

Optimized for multi-core processors

Intel® Ethernet 10 Gigabit adapters include a number of latency-lowering features that are optimized for multi-core Intel Xeon processor-based servers:

- ***Extended Message-Signaled Interrupts (MSI-X)*** provide multiple interrupt vectors, which allow multiple interrupts to be handled simultaneously and load-balanced across multiple processor cores, helping to improve CPU utilization and lower latency. Previous-generation MSI passed interrupts to a single processor core, leading to less efficient performance.
- ***Receive-side Scaling (RSS)*** improves server load-balancing by segregating incoming packets into flows and directing those flows to separate hardware queues in Intel® Ethernet adapters, allowing them to be processed simultaneously. (On Linux systems, this technology is known as Scalable I/O.) RSS can also be used to direct multiple TCP/IP streams to specific processor cores for handling.

Optimized for virtualization

Intel® Virtualization Technology for Connectivity (Intel® VT-c) enhances server I/O solutions by integrating extensive hardware assists into the Ethernet controller. This collection of technologies addresses I/O bottlenecks by either offloading data-packet processing to the Intel Ethernet adapter from the hypervisor or providing direct I/O connectivity to the VMs for faster application responsiveness and improved processor utilization.

Two technologies comprise Intel VT-c: Virtual Machine Device Queues (VDMq), which supports virtual switch emulation, and Virtual Machine Direct Connect (VMDc), which provides direct assignment between a network connection and a VM.

- ***VMDq (emulation)*** improves overall CPU utilization and throughput levels by offloading the network I/O management burden from the virtual switch in the hypervisor to the Ethernet controller. Multiple queues and sorting intelligence in the silicon support enhanced network traffic flow in the virtual environment, freeing processor cycles for application work.
- ***VMDc (direct assignment)*** provides near-native I/O performance by facilitating dedicated I/O and data isolation among VMs and also enables VM migration. VMDc uses PCI-SIG SR-IOV and Intel® Virtualization Technology for Directed I/O (Intel® VT-d) to support this functionality.

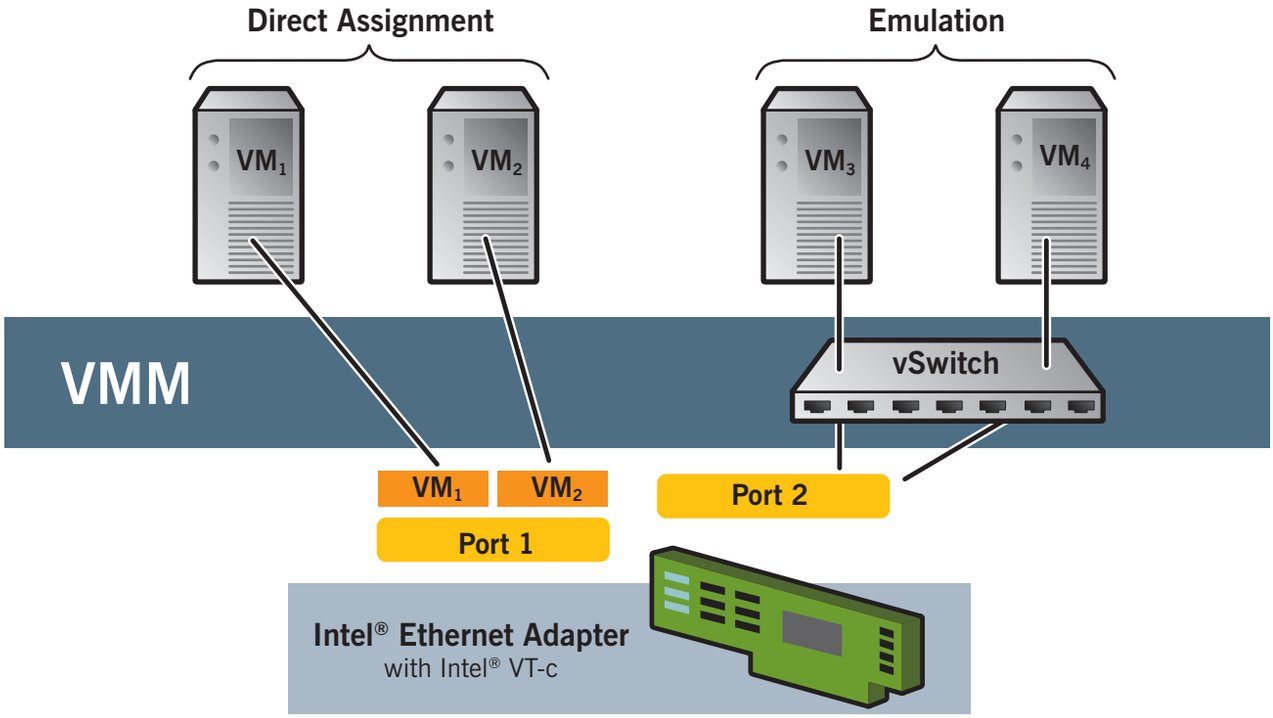

Figure 3: Intel® Virtualization Technology for Connectivity provides direct assignment and switch emulation modes.

Driving storage over Ethernet

The fast growth in storage capacity coupled with server virtualization has brought the need for the SAN to the forefront, and technologies such as iSCSI are leading the charge for storage over Ethernet. Limited bandwidth and TCP processing overhead are two of the key factors that have limited the adoption of iSCSI. However, today's multi-core processor-based platforms have more than enough power to handle these workloads, and 10 Gigabit Ethernet bandwidth provides sufficient throughput.

Intel Ethernet 10 Gigabit server adapters include support for iSCSI acceleration and advanced features for unified storage connectivity. The controller enables fast and reliable networked storage with native iSCSI initiator support with Microsoft, Linux, and VMware OSs as well as support for iSCSI remote boot.

The growth in 10 Gigabit Ethernet will enable greater deployments of virtualized servers and Ethernet SANs, providing unparalleled throughput. These performance improvements come at a cost, however, as servers will be forced to handle more network data than ever before.

Optimized for multi-core platforms, virtualized environments, and storage over Ethernet, Intel 10 Gigabit Ethernet products provide the throughput and traffic processing optimizations for the next generation of data centers.

Intel 10 Gigabit adapters offer a broad range of OS support, including Windows, Linux, and ESX, and its broad Ethernet portfolio includes single-port and multi-port adapters for fiber and copper networks.

EXTREME NETWORKS 10 GIGABIT ETHERNET SWITCHING SOLUTIONS

Extreme Networks is known for its leadership and innovation in Ethernet switching. With more than 10 years of experience developing innovative products to advance network technology, it is not surprising that Extreme is a leader in providing 10 Gigabit Ethernet switching solutions. With Extreme Networks award-winning BlackDiamond® 8800 series switches the customer can tailor the chassis through purpose-built core, aggregation, and edge modules to fit the exact requirements of data center deployments. Extreme Networks new Summit® X650 fixed configuration switch provides 24 non-blocking 10 Gigabit Ethernet ports. The

Summit X650 is ideal as a top-of-rack switch for data center server connectivity. Correspondingly, the BlackDiamond 8800 is well suited as an end-of-row switch, aggregating edge and core connectivity in the data center. Complementing the BlackDiamond 8800 and the Summit X650 is Extreme Networks ExtremeXOS® modular OS resulting in a highly reliable, yet simple to configure, solution that meets the demanding networking needs found in modern data centers.

BlackDiamond 8800 series switch

The design of the BlackDiamond 8800 series switch provides optimal flexibility to data center network configurations with its modular design. Customers can select purpose-built modules or blades to tailor the BlackDiamond 8800 port types and speeds. With its non-blocking interconnect design and a total switching capacity of 1.45 terabits per second and 816 gigabits per second in the backplanes of the BlackDiamond 8810 and BlackDiamond 8806, respectively, there is plenty of capacity to scale these switches when future network needs increase.

The latest BlackDiamond 8800's "c" series blade selection offers an eight port 10 Gigabit Ethernet blade. With nine of these blades configured in a BlackDiamond 8810, the result is seventy-two 10 Gigabit Ethernet ports all within a 14RU chassis. Similarly, the BlackDiamond 8806 can be configured with up to five 10 Gigabit Ethernet blades for a maximum of forty 10 Gigabit Ethernet ports contained in a 10RU chassis.

Both BlackDiamond 8800 models include a passive backplane complemented by high availability design elements, such as isolated control and data planes, redundant controller boards for power distribution, and physical monitoring, to identify potential problems before they affect network availability. To ensure continuous operation the BlackDiamond 8800 series offers redundant power supplies and redundant Management Switch Modules (MSMs) with active failover.

Extreme Networks BlackDiamond 8800 series switches deliver both Gigabit Ethernet and 10 Gigabit Ethernet wherever it is needed—from the network edge to the network core and the ability to aggregate high-speed connections to eliminate bottlenecks between the edge and the core. The BlackDiamond 8800 is an ideal end-of-row solution for dense rack server environments in the ever-expanding data center.

Summit X650 series switch

The Summit X650 series switch offers remarkable performance in a very small package. This switch is a purpose-built top-of-rack design offering 24 non-blocking 10 Gigabit Ethernet ports in a 1RU form factor. It is ideal for deployment in data centers where connectivity to blade servers and Ethernet-based storage systems is contained in a single data center rack.

The Summit X650 is offered in two models: a 10GBASE-T version to accommodate the need for copper twisted-pair cable connections and an SFP+ version (10GBASE-X) to accommodate optical fiber connections. Both models contain the most advanced 10 Gigabit Ethernet technology to support exceptionally high-density switching with very low latencies and highly scalable routing. The result is an Ethernet switch that can meet the critical demands of the data center.

One of the key features of the Summit X650 series switch is its ability to stack with other Summit switches through Extreme's switch stacking technology, SummitStack™. The Summit X650 offers flexible stacking options supporting existing stacking configurations Summit X250 and Summit X450 switches using a 40 Gigabit SummitStack interconnection, a new ultra high speed Summit Stack256 that supports the stacking of eight Summit X650s for a massive 192-port 10 Gigabit Ethernet, or a SummitStack512 allowing for two Summit X650s to be stacked together for non-blocking performance across all forty-eight 10 Gigabit Ethernet ports. These stacking options are supported through Extreme's unique Summit X650 Versatile Interface Module (VIM) giving customers the ability to configure the Summit X650 to meet their specific data center needs.

The Summit X650 is the answer for migrating existing servers and storage systems from Gigabit Ethernet environments to 10 Gigabit Ethernet in modern, virtualized data centers. With its versatile design, the Summit X650 has the flexibility to accommodate massive 10 Gigabit Ethernet edge configurations along with the ability to aggregate traffic onto a data center's core backbone network.

ExtremeXOS

Common to both the BlackDiamond 8800 and the Summit X650 is a modular OS: ExtremeXOS. With one OS supporting Extreme Networks entire product switching portfolio, configuring and maintaining the network environment has never been simpler. ExtremeXOS extensive XML APIs and CLI scripting facility allows network administrators to manage a complex data center network through a single interface eliminating the need to learn multiple management interfaces.

Combined with ExtremeXOS ease of management is its rich-feature set. Some of these features relevant to the data center include seamless support for mixed IPv4 and IPv6 network environments, security support through access control lists (ACLs) to provide protection to the network, and its high availability architecture and Ethernet Automated Protection Switching (EAPS) protocol to reduce network downtime for business continuity of mission-critical applications.

ExtremeXOS common modular OS image spanning the entire network switching environment affords a data center the necessary command-and-control infrastructure while, at the same time, removing inherent network complexities.

Deploying Extreme Networks 10 Gigabit Ethernet switching solutions at the data center network edge enables servers, storage systems, and applications to achieve greater levels of performance and productivity. These performance and productivity improvements can dramatically change the landscape of information processing as we know it today—delivering vast amounts of information to applications and users at data rates previously not attainable. The result is new efficiencies of scale and the corresponding realized cost savings as productivity increases.

Conclusion

One thing is certain: The demand for information is growing exponentially. Along with this growth is a corresponding demand for delivering this information to the applications and end-users that process it. Advances in information-processing technology, such as multi-core processors, server virtualization, dramatically increased capacities of storage systems, and a wealth of new information forms, all are driving the demand on networks to deliver more information faster and more efficiently. To achieve these ends, Intel and Extreme Networks are providing 10 Gigabit Ethernet solutions for existing and new data center networking deployments, ensuring that increased network demand can be realized at affordable costs.

FOR MORE INFORMATION

Intel® Ethernet: **www.intel.com/go/ethernet**

Extreme Networks®: **www.extremenetworks.com**

Intel® 10 Gigabit Ethernet Server adapters:
www.intel.com/go/10Gbe

Intel® Virtualization Technology for Connectivity:
www.intel.com/go/vtc

Printed in USA 0409/BY/MESH/PDF Please Recycle 321816-001US