

A large, vertical, grey, curved shape on the left side of the page, resembling a stylized 'S' or a drop shadow, extending from the top to the bottom of the page.

10 GbE RISING FACTORS DRIVING THE EMERGENCE OF 10 GIGABIT ETHERNET

White Paper
February 2007

Table of Contents

Executive Summary	1
Introduction/Ethernet Overview	2
The Arrival of 10 Gigabit Ethernet	4
The Emergence of 10 Gigabit Ethernet Over Copper	5
10 Gigabit Ethernet Unleashed	6
Datacenter Design Reloaded	7
Server Consolidation	7
Server Virtualization	8
Sun Virtualizes the 10 Gigabit Ethernet NIC	9
The Rise of Bandwidth	11
The Telco Market	12
Advanced Telecommunications Computing Architecture (ATCA)	12
The Enterprise Network	13
Client-driven Application Environment	13
Business Continuity	14
Solaris OS	15
Project FireEngine	15
Project Crossbow	16
Summary	17

Executive Summary

IT capacity and demand are on a collision course for the datacenter. Industry trends such as Chip Multithreading (CMT), multicore CPUs such as CoolThreads™, Intel Xeon, and AMD Opteron®, blade computing, and virtualization are just a few of the significant advances that have increased datacenter capacity. But the demand for IT compute and storage capabilities is expected to increase at an even faster rate. Ethernet, the staple of datacenter networking infrastructure, has kept pace with the introduction of 10 Gigabit Ethernet trunking on commodity Layer 2 switches, but until recently server Network Interface Cards (NICs) were stalled at Gigabit Ethernet speeds.

Demanding workloads combined with advances in CPU design challenge conventional Gigabit Ethernet plumbing. Trends such as 64-bit computing, terabyte desktop storage, and multicore CPUs are also driving demand for LAN throughput, and ultimately, access capacity to the datacenter. But the imbalance between server throughput and Ethernet attachment is about to be righted. The standardization of 10 Gigabit Ethernet over Category 6 UTP copper, combined with more cost-effective silicon, is changing the dynamics of local-area networking.

Maximum performance means eliminating bottlenecks between the CPU, end-user applications, and network access, a solution characterized simply as end-to-end throughput computing. Sun is an undisputed leader and innovator in high-throughput computing. From the advantageous position of a systems vendor, Sun is well qualified to approach the impedance mismatch between server performance and network throughput with complementary investments across the solution stack in both hardware and software. The results have pioneered a volume high-throughput server architecture. Sun has made many advancements, including CMT and cutting-edge contributions to the Solaris™ Operating System (OS), taking full advantage of unprecedented parallelism for servers and network interfaces. From its groundbreaking CPU design, Sun is now applying CMT to network processing. In addition to its market-leading performance, Sun has designed a wealth of innovative features that help enable new networking functionality from the edge to the core—from the desktop to the datacenter.

This paper will examine the forces that are converging to push 10 Gigabit Ethernet from its niche status to pervasive deployment across both enterprise and telco markets. The paper will showcase Sun's emergence as a leader in 10 Gigabit Ethernet technology with a CMT-optimized networking architecture designed for NICs, LAN on the motherboard (LOM), and even embedded in the CPU die. This novel implementation helps deliver unprecedented performance at price points that accelerate wide deployment.

Introduction/Ethernet Overview




Figure 1: While Ethernet speeds have increased by several orders of magnitude over the years, the economics underlying recent gains were poorly aligned with advancements in CPU and storage capabilities. 10 GbE has kept pace with multicore CPUs and massive storage capabilities, but was too expensive to implement.

The origins of Ethernet stem from work done by Robert Metcalfe at Xerox PARC in 1972, and the technology was subsequently patented in 1977. Ethernet began as an experimental coaxial cable network by Xerox with a shared data rate of 3 Mb/sec. Early deployment success led to the 1980 joint development of the shared 10 Mb/sec Ethernet Version 1.0 specification by the three-company consortium of DEC, Intel, and Xerox. The draft standard was approved by the 802.3 working group in 1983 and was subsequently published as an official standard in 1985 (ANSI/IEEE Std. 802.3-1985).

Ethernet got its next speed increase to 100BaseT (802.3u) in the mid-1990s, followed by the introduction of Gigabit Ethernet (802.3ab). Known as 1000BaseT, it was ratified in 1999 and defined transmission over Unshielded Twisted-Pair (UTP) Category 5, 5e, or 6 cabling. As the transmission speeds increased, the deployment technology progressed from a bandwidth-sharing model (built on collision-detection multiple access) to dedicated switched LAN segments. Switching LAN segments (also known as bridging) meant that each connected host had access to 100 percent of the bandwidth.

As terrestrial switched Ethernet grew in speed, the wireless standards emerged and reversed the trend toward dedicated bandwidth, providing wireless clients originally with a technology that shared 11 Mb/sec (802.11b), then 54 Mb/sec (802.11g and 802.11a). The wireless industry will soon announce even faster throughput – perhaps as fast as 500 Mb/sec (802.11n), though still shared.

Gigabit Ethernet has enjoyed wide-scale deployments in switching infrastructures, computer NICs, and LOMs. The versatility of Ethernet makes it the ideal interconnect, with useful features such as speed auto-negotiation (10/100/1000 Mb/sec), hardware flow control, ability to bond multiple links together to make fatter pipes (also known as link aggregation, or 802.3ad), mature loop detection protocols, and ultra-high-volume deployment.

While Ethernet has made dramatic advances in overall bandwidth performance, the economics of the most recent speed increases (10 Gb/sec, up from 1 Gb/sec) have not aligned well with the pace of CPU and storage improvements. Consider that multiple cores per chip and multiple threads per core have pushed chip processing from barely 1 MIPS in the first generation 8088 CPUs to billions of instructions per second in modern CPUs, with performance doubling approximately every 20 months. Storage access with Fibre Channel interfaces has grown to 4 Gb/sec throughput, and InfiniBand has increased to 20 Gb/sec with Double Data Rate (DDR) running over copper. These breakthroughs in server and storage throughput have blown past the utility of embedded Gigabit Ethernet interfaces. Since its initial ratification in 2002, 10 Gigabit Ethernet has, until now, been a poor fit for server and storage network access.

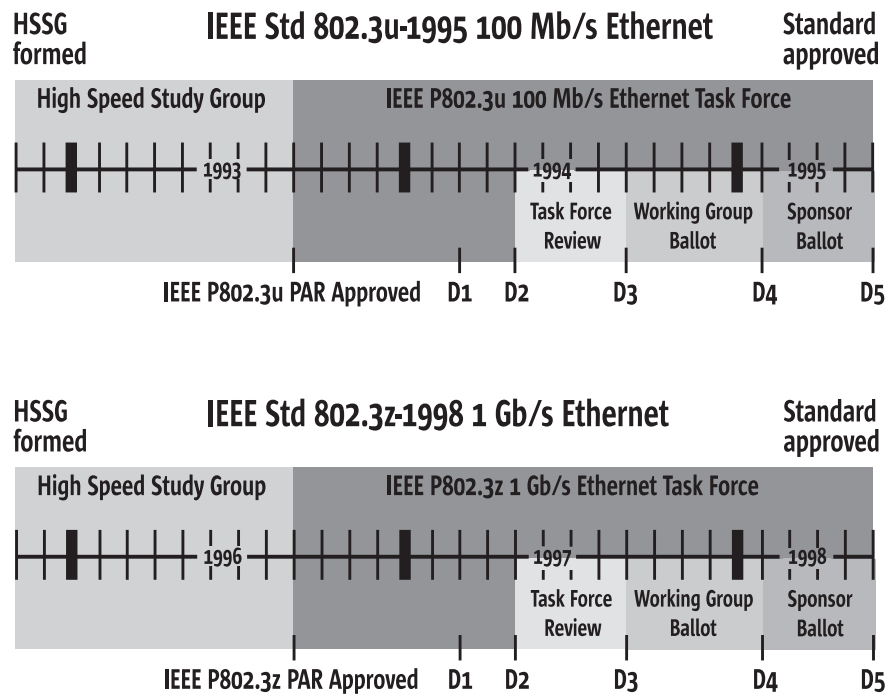


Figure 2: 100 Mb/sec and 1 Gb/sec ratification timelines (Source: David Law 3Com)

The Arrival of 10 Gigabit Ethernet

In 2002, 10 Gigabit Ethernet was ratified (802.3ae), but unlike its older siblings, wide-scale server-network access deployment has been hampered by numerous factors. Initially, the two most prominent impediments to its adoption were chip costs and absence of support for UTP cabling (of any category). The use of 10 Gigabit Ethernet as a trunking technology for commodity Layer 2 switches and routers is flourishing, thanks to advances in optical interconnects¹.

While Sun's CMT-enabled 10 Gigabit Ethernet technology has addressed the chip cost issue, the absence of copper cabling was more problematic. In fact, the first IEEE standard for 10 Gigabit Ethernet transmission over copper defined an Ethernet physical layer (PHY) that borrows from the InfiniBand cabling standard, and is composed of dual Twinax coax cables (as defined by the 10GBaseCX4 initiative, designated as 802.3ak). Unfortunately, 10GBaseCX4 has limited appeal for most Ethernet deployment opportunities because it is relatively unknown to all but those who have invested in InfiniBand cabling plants, and cable length is limited to 15 meters. A 10GBaseCX4 copper solution meant that 10 Gigabit Ethernet was still not ready for widespread deployment.

In summary, 10 Gigabit Ethernet has had a slow start. This is illustrated in Figure 3, which shows the uptake of Ethernet technologies in their first six years of standardization. Note that Fast Ethernet and Gigabit Ethernet both had strong starts (100 percent in the first year) and had good growth trajectories after that. 10 Gigabit Ethernet shipped almost nothing until four years after the standard was ratified.

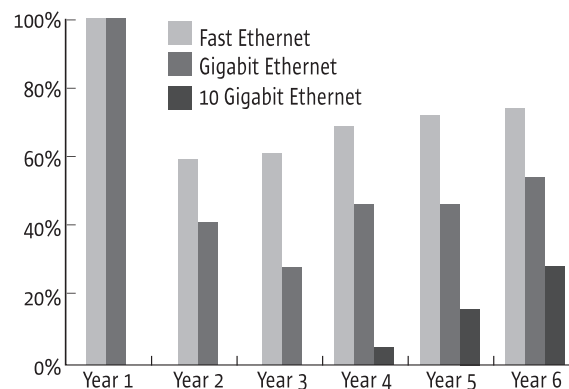


Figure 3: Ethernet switch port configuration shipments in first six years after ratification. (Source: Dell'Oro Group)

¹Optical interconnect form factors have rapidly evolved from XENPACK to XPACK to X2 to XFP, and now SFP+. With each generation of optical interconnect the form factor has become smaller, allowing for better port density, lower costs, and reduced power consumption.

The Emergence of 10 Gigabit Ethernet Over Copper

In June 2006, the IEEE ratified 10GBaseT (802.3an – 10 Gigabit Ethernet over twisted pair), removing one of the major hurdles to 10 Gigabit Ethernet adoption. At the same time, low-cost 10GBaseT chips started sampling in 2006, and a new generation of switch and NIC products is expected to arrive in 2007. These achievements will help enable 10 Gigabit performance with the simplicity and cost-effectiveness of twisted-pair copper. With costs coming down and its applicability to existing networking cable plants, 10GBaseT is expected to rapidly become the optimal high-speed networking solution for highly consolidated enterprise computing and high-performance computing (HPC) datacenters.

The 802.3an industry standard specifies running 10 Gb/sec of Ethernet data over Unshielded Twisted-Pair (UTP) cabling, with a maximum length of 100 meters. Standards crafters say 10GBaseT will reach the maximum distance easily over Category 7 and Category 6a UTP cabling. Users can expect only around 55 meters of reach using Category 6 cabling, and only around or 45 meters of reach over Category 5e, which is one of the more widely installed network cabling types in corporate nets.




Figure 4: 10 Gb/sec ratification timeline (Source: David Law, 3Com)

10 Gigabit Ethernet Unleashed

Anticipating the trends in 10 Gigabit Ethernet networking, Sun has developed breakthrough silicon that maximizes end-to-end throughput from the CPU through the Ethernet port. Closely aligned with the proven performance architecture of CMT, Sun's 10 Gigabit Ethernet solution harnesses the bandwidth of an 8x PCI-E bus (32 Gb/sec full duplex) commonly available in today's server and workstation environments. Now users and applications can access the enormous power of high-performance, multicore and CMT CPUs over the network.

Sun's CMT-optimized networking technology contains two 10 Gigabit and two Gigabit Ethernet MACs. Unlike competitive offerings, which supply a modest number of DMA channels, the CMT-optimized networking ASIC provides 24 DMA channels on the TX (transmit) side and 16 DMA channels on the RX (receive) side. More DMA channels are needed to align processor threads with DMA access to the network interface. Instead of multiple threads competing for a few—or even one—DMA channel, each thread has full use of an entire DMA channel. This maximizes concurrency for demanding transactional loads originating from millions of client/subscriber requests, and dramatically increases throughput.

Datacenter Design Reloaded

Now that issues relating to cost and transmission medium are resolved, 10 Gigabit Ethernet is poised for large-scale deployment into enterprise and telco environments. This technology will significantly change existing datacenter LAN architectures.

Server Consolidation

Client/server LANs are going through a period of unprecedented growth, while simultaneously experiencing tremendous consolidation. The number of applications in a modern enterprise network has been growing for decades—often there are hundreds, and sometimes thousands of applications. The original model of one server per application for a single client constituency has led to datacenters with far too many underutilized servers and poorly allocated storage resources. Today's underutilized compute farms are comprised of overwhelmingly complex LANs, which are necessary to accommodate the proliferation of server NICs necessary to meet client demands. The energy requirements alone justify a move toward consolidation and virtualization.

How does server consolidation affect LANs? Fewer servers doesn't necessarily equate to fewer network interfaces. On the contrary: most high-performance servers require multiple Gigabit Ethernet interfaces to meet the service level demands of networked applications, along with other requirements such as NAS, online backup, and management. A typical server motherboard may provide up to four Gigabit Ethernet ports and carry additional Gigabit Ethernet interfaces as PCI NICs. To simplify the burdens of assigning IP addresses, network administrators often group multiple Gigabit Ethernet interfaces into 802.3ad trunks, which look like a single IP interface to the server. Link aggregation eases IP addressing but does nothing to simplify datacenter cable management, as a Layer 2 switch must still accommodate every individual connection and, barring oversubscription, the switch must carry a large proportion of the server bandwidth upstream to the next hop. As indicated earlier, 10 Gigabit Ethernet over optical fiber is well suited to the tasks of aggregating multiple Gigabit Ethernet connections and delivering downstream bandwidth to an upstream entity.

When deployed within a NIC or as a LOM, Sun's CMT-optimized networking technology utilizes 8x PCI-E bus, providing 32 Gb/sec full-duplex bandwidth from the CPU. Sun has also implemented CMT-enabled networking technology directly onto the Niagara 2 CPU die to unleash its full potential, unconstrained by PCI-E bandwidth. When combined with low-cost copper PHYs in a 10GBaseT configuration, it's clear that a significant reduction in Ethernet NICs, and the associated cabling, is achievable. The potential for a 10:1 reduction in cabling represents a breakthrough in LAN design and deployment optimization.

Layer 2 aggregation, typically executed in commodity Layer 2 switches (see Figure 5), can now be pushed into the server. Using Sun’s CMT-enabled 10 Gigabit Ethernet technology, the NIC has the potential to emulate a wire-speed Layer 2 switch, further consolidating datacenter LAN infrastructure while improving overall functionality.

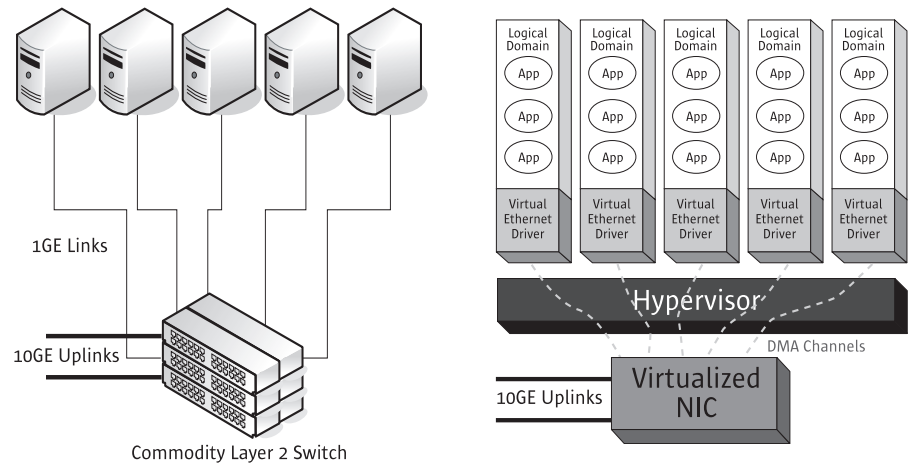


Figure 5: Consolidation and cabling reduction with virtualized NICs

Server Virtualization

The benefits of a consolidated 10 Gigabit Ethernet interface are substantially extended when server virtualization is added to the mix. Server virtualization enables a datacenter to be viewed and managed as a set of compute resources rather than a room of individual systems, allowing businesses to reduce costs through improved efficiency and better resource utilization, while improving service levels. Server virtualization permits datacenter operators to eliminate the wasteful single-server, single-application model by leveraging Solaris Containers, a server hypervisor, or both.

Solaris Containers enable a single machine to be securely partitioned into virtual OS domains, or containers. Each container runs independently and provides a complete virtual instance of the entire Solaris runtime environment, including root and user file systems, virtual interfaces, and user accounts. The Solaris 10 OS can support up to 8,000 containers, depending on machine size and performance. There is a global container that has secure access to all user-defined containers for global server administration. Each Solaris Container is provided with a virtual network interface that can optionally leverage VLAN tagging when consolidated over a single 10 Gigabit Ethernet interface.

Sun also supports hardware-level virtualization with its implementation of the Sun hypervisor. A hypervisor (also known as a virtual machine monitor) is a virtualization platform that helps enable multiple operating systems to run simultaneously on a single host computer, each contained within a specialized partition, collectively called logical domains (LDOMs). Each LDOM is provided its private network stack, including IP addressing as well as Ethernet MAC address². The hypervisor is often described as a software shim that runs directly on a given hardware platform. A guest operating system then runs on top of the hypervisor, one level removed from the hardware.

Multiple host operating systems have been modified to run as guest operating systems on the hypervisor. As of late 2006, the Solaris OS, Linux, and FreeBSD have been ported to run on top of the Sun hypervisor (each runs simultaneously on the same processor, as a fully virtualized independent guest OS). Wind River also has announced plans to run its version of Carrier Grade Linux on the Sun hypervisor. Full virtualization on UltraSPARC® processors was not difficult because the SPARC® architecture, since its inception in the mid-1980s, was deliberately kept clean of artifacts that would have adversely affected virtualization.

Sun Virtualizes the 10 Gigabit Ethernet NIC

The role of the Layer 2 switch is changing. The server now acts as the aggregation point for multiple virtual NICs when Sun 10 Gigabit multichannel technology is coupled to the virtualized application plane. The benefits of virtualized NICs are numerous:

- Multiple networked applications can run with private network stacks and dedicated Ethernet I/O.
- The 10 Gigabit Ethernet uplink acts like a Layer 2 switch-aggregation point for numerous virtual server instances, which reduces the number of Layer 2 switches required for Gigabit Ethernet aggregation.
- Individual application stacks can utilize a variety of addressing models, including dedicated IP addresses, dedicated Ethernet MAC addresses, VLAN tags, or combinations to suit the requirements imposed by customer best practices.
 - Fewer Gigabit Ethernet aggregation switches translates into lower operational expense and less power consumption.
 - 10 Gigabit Ethernet consolidated, virtualized NICs offer a 10:1 reduction in cabling or better, depending on tolerance for oversubscription.
 - When server uplink redundancy is factored into cable reduction, the benefits of a 10 Gigabit Ethernet consolidated NIC is compounded (often yielding 8:2 reduction per server).

² The number of dedicated Ethernet MAC addresses is a function of the virtualized 10 Gigabit Ethernet NIC. Sun's CMT-enabled 10 Gigabit Ethernet technology supplies 16 LDOMs, each with a dedicated Ethernet MAC address.

Sun's CMT enabled 10 Gigabit Ethernet technology has been optimized to virtualize the NIC: it's designed for LDOMs where each domain can be assigned unique network interface resources.

Sun's roadmap for end-to-end server virtualization includes many innovations, including the plumbing of LDOMs directly to DMA channels (designed into Sun's CMT-optimized networking technology). This approach effectively virtualizes the server's NIC while eliminating bottlenecks in Ethernet I/O. A true virtualized 10 Gigabit Ethernet NIC narrows the behavioral distinction between the Sun virtual server and legacy discrete servers with dedicated Gigabit Ethernet interfaces. Virtualized services host applications and follow administrative constructs as though they are served from a unique, dedicated system.

The Rise of Bandwidth

The Ethernet switch market has adapted rapidly to meet increased demand for enterprise and telco datacenters. Trends such as Web 2.0, high-performance computing grids, Internet Protocol television (IPTV), iSCSI, and business continuity are fueling bandwidth and processing requirements. 10 Gigabit Ethernet is expected to meet these new requirements.

There are many choices for high-density Gigabit Ethernet and a variety of 10 Gigabit Ethernet products, ranging from aggregated trunking to direct data connectivity to core datacenter routers using OC-48 (2.48 Gb/sec) and OC-192 (10 Gb/sec) Internet access. But the supply side of the equation is driven from the rise of network applications, network-based storage solutions, and network-attached computers (of which Sun was an early innovator).

The relationship between the server NIC bandwidth and the switch's trunking bandwidth is the key to understanding the drive to increasingly higher Ethernet bandwidth adoption. Computers are strung together using high-density commodity Ethernet switching, and current practices allow for varying levels of server-to-network oversubscription³. But too much oversubscription can lead to network congestion and service-level degradation, resulting in poor customer satisfaction. As servers increase in throughput, either the number of NICs and/or the throughput of each NIC must be increased. Faster NICs drive the need for faster switch ports and, in turn, for faster trunking.

³ The average Ethernet link utilization is typically a small fraction of available bandwidth, with periodic short bursts of higher utilization.

The Telco Market

Opportunities for increasing revenues at telco operators and service providers are primarily focused on emerging Internet Protocol (IP) services. Services such as push-to-talk, multiplayer gaming, video, broadband, VoIP, and SMS (texting) represent significant revenue opportunities over conventional voice services. As telcos move broadly to deploy new services such as IPTV and metroscale wireless mesh, the demand for large back-end pipes will only increase. These new services rely on IP protocols and packet-based technology, such as IP Multimedia Subsystem (IMS). Data-driven services will drive bandwidth requirements through telco networks, from regional datacenters to local central offices. Telcos are gearing up to deliver IP-based solutions on an enormous scale. 10 Gigabit Ethernet can help meet this demand cost effectively, and in some situations is expected to replace or augment SONET links.

Advanced Telecommunications Computing Architecture (ATCA)

To help standardize future IT platforms, the telco market is adopting the Advanced Telecommunications Computing Architecture (ATCA), helping telcos focus on creating applications instead of systems, and leveraging the cost advantages of a high-volume, general-purpose platform. Sun is reenergizing the ATCA platform with next-generation, CMT-powered networking. The availability of 10 Gigabit Ethernet creates a wealth of opportunities for telcos by providing the cost-effective bandwidth needed to open existing throughput bottlenecks and fuel a variety of emerging applications.

Sun Netra™ ATCA products are fully compliant with industry-standard ATCA specifications and are interoperable with third-party offerings. Sun's decades of experience in the telco industry and its solid relationship with the PCI Industrial Computer Manufacturers Group (PICMG) have also ideally positioned it for its role in the development of the ATCA standard. As a result, Sun has designed hardware and software systems that are well qualified to meet this new standard's requirements. Sun systems offer unexcelled choice, a complete ATCA system offering, superior system management software, and world-class services.

Sun's ATCA offering is unique. Sun offers high-performance UltraSPARC and AMD Opteron™ blades, as well as a choice of the two most popular carrier-grade operating systems, Solaris and Carrier Grade Linux, in the same system. Additionally, Sun brings its Netra High Availability Suite (NHAS) products to ATCA and implements a set of Service Availability Forum (SA Forum) interfaces. The Sun ATCA product family offers OS choice – the proven, carrier-grade Solaris OS on UltraSPARC and AMD Opteron blades, as well as OSDL-compliant, carrier-grade Linux OS on AMD Opteron blades.

The Enterprise Network

Network bandwidth requirements are increasing throughout the enterprise. Network-based applications, diskless computing, voice and video over IP, Web conferencing and training, and other services are driving the need for more bandwidth throughout the network.

Client-driven Application Environment

The adoption of SAN emulation (iSCSI and FCoIP) has been hampered by the obvious trade-offs in bandwidth and server TCP processing. iSCSI and FCoIP are typically used for different purposes:

- With iSCSI, SCSI commands and data frames are encapsulated in IP to support I/O disk access over an IP network.
- With FCoIP, Fibre Channel frames are encapsulated in IP so that both SCSI and non-SCSI frames can be transported over TCP/IP.

The bandwidth demands for carrying block-mode protocols over IP over Ethernet, combined with the latency incurred by typical server TCP stacks, have hindered the wide-scale deployment of these disruptive technologies. There are many point solutions to accelerate the TCP stack, including TCP Offload Engines (TOEs) built into the NIC or embedded in a network appliance. The TOE approach to TCP acceleration has been in existence for a number of years and owes its design center for augmenting single-threaded general-purpose processors as well as for application-aware network appliances. As mentioned earlier, server throughput is doubling every 20 months and when combined with CMT, the server processor can often outperform the TOE, rendering it obsolete. Sun's CMT, combined with CMT-optimized networking technology, will fundamentally change the dynamics of running FCoIP and iSCSI alongside traditional client-server application networking stacks. CMT eliminates processor stalls incurred during memory access by launching the next hardware thread without a context switch, effectively queuing traffic to the TCP stack with minimal latency. Processor threads utilize the abundant DMA channels running on the CMT-optimized networking ASIC to fill available Ethernet bandwidth.

The combination of Sun CMT-powered processors and the industry's only CMT-optimized networking ASIC enables enterprises to deploy FCoIP and iSCSI to interconnect SAN islands, provide SAN-level remote backup and replication, and run full production traffic—all over a single consolidated 10 Gigabit Ethernet interface.

Business Continuity

Secure and reliable data storage solutions with high capacity and availability have become a critical requirement for an ever-increasing number of enterprises. 10 Gigabit Ethernet provides ample bandwidth for production and maintenance traffic such as backup or provisioning.

Due to the huge increase in storage demand, precise, internationally recognized standards and storage-dedicated protocols had to be created. Until recently, the only valid protocol was ESCON, which enables operations up to 200 Mb/sec. Since higher speeds were needed, many IT professionals moved to the Fibre Channel protocol, which enables a rate of up to 4 Gb/sec for storage traffic. In addition, database storage for IP applications is also becoming mandatory for enterprise resource planning applications, security (video surveillance) and other requirements. This is driving the demand for more bandwidth through the IT environment, from the edge to the core. New regulations and mandates such as the Sarbanes-Oxley Act, the Health Insurance Portability and Accountability Act (HIPAA), and regulations from the European Community such as the International Convergence of Capital Measurement and Capital Standards—a Revised Framework, otherwise known as Basel II, were created to answer the high demand for crucial data safeguarding. Backup, whether to near-line drives or tape devices, must occur without taking systems offline—and 10 Gigabit Ethernet can provide that capacity.

Solaris OS

Project FireEngine

Project FireEngine, a new TCP/IP stack in the Solaris 10 OS, is a high-performance networking framework based on an IP classifier and vertical perimeters. This new architecture reduces the overhead of synchronization and cross-communication between CPUs, a necessary ingredient for scaling across a very large number of CPUs.

In addition to a number of overall performance and management improvements, Project FireEngine enables the Solaris 10 OS to drive 10 Gb/sec line rates—the only limitation is PCI-E bus bandwidth. Project FireEngine accommodates 10 Gb/sec NICs by supporting dynamic switching between interrupt and polling modes, allowing systems to pick up multiple packets. This gives system CPUs sufficient time to process packets, instead of continually bombarding them with the arrival of new packets. Project FireEngine boosts network performance without requiring any hardware upgrades or changes to existing 32- or 64-bit applications.

Project Crossbow

Project Crossbow is designed to provide the building blocks for utility computing by creating virtual stacks around services such as HTTP, HTTPS, FTP, and NFS; protocols such as TCP, UDP, and SCTP; and Solaris Containers technology. The virtual stacks are separated by means of a hardware classification engine such that traffic for one stack does not impact other virtual stacks. Each virtual stack can be assigned its own priority and bandwidth on a shared NIC without causing performance degradation to the system or the service/container. Project Crossbow helps address the I/O problems typically associated with virtualization. The architecture dynamically manages priority and bandwidth resources, and can provide better defense against denial-of-service attacks directed at a particular service or container by isolating the impact to just that service or container.

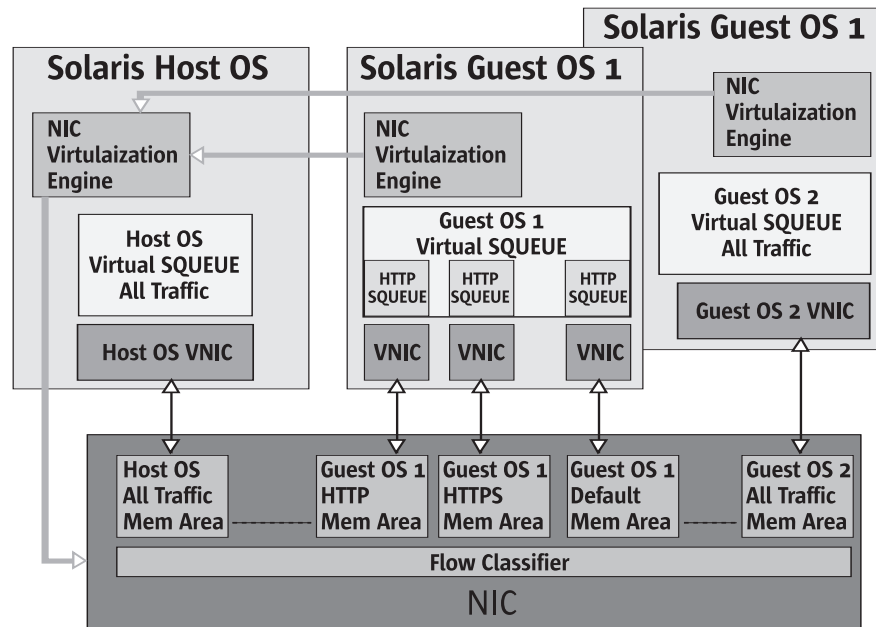


Figure 6: Crossbow combined with the Sun CMT-optimized networking ASIC

Summary

Sun is leading the campaign for datacenter consolidation in a number of important ways, including advances in chip design and multiple levels of server virtualization. Innovations such as CMT have resulted in pace-setting performance throughout its product line, including the Niagara 1 processor used in Sun Fire™ T1000 and Sun Fire T2000 servers, and the blade chassis with as many as 16 Opteron cores per server blade. And Sun's CoolThreads servers with CMT consume up to 80 percent less power per thread of computing.

CoolThreads technology-based servers include up to four Gigabit Ethernet ports on the motherboard, and for many common workloads they are able to drive full bandwidth across all interfaces. The roadmap for CMT, including Niagara 2 and Rock (the next-generation CPU core), promises improvement in computation as well as increased throughput capabilities. The design center for the CMT-optimized networking technology is throughput parity with all of Sun's next-generation CMT and multicore Opteron servers.

Sun is the only vendor supplying 10 g to meet the demands of virtualization by offering a true 10 Gigabit Ethernet over PCI-E, combined with 40 DMA channels, thereby aligning CMT or multicore CPUs with emerging networking requirements. In this way, Sun can bring unparalleled throughput, power, and efficiency to enterprise and telco environments.



Sun Microsystems, Inc. 4150 Network Circle, Santa Clara, CA 95054 USA **Phone** 1-650-960-1300 or 1-800-555-9SUN (9786) **Web** sun.com



© 2007 Sun Microsystems, Inc. All rights reserved. Sun, Sun Microsystems, the Sun logo, CoolThreads, Netra, Solaris, and Sun Fire are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States and other countries. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. in the U.S. and other countries. Products bearing SPARC trademarks are based upon an architecture developed by Sun Microsystems, Inc. AMD and Opteron are trademarks or registered trademarks of Advanced Micro Devices. Information subject to change without notice.
Printed in USA 02/07 SunWin 498481