

Sun CoolThreads™ Servers and Zeus Technology

Next-generation load balancing and application traffic management



Rapid and unpredictable growth in Web services applications is driving the need for highly flexible Web infrastructure solutions and new ways to accelerate application performance. Traditional load-balancing appliances based on dedicated hardware can restrict agility and add unnecessary complexity to the network topology. Recent advances in performance and virtualization with Sun CoolThreads™ servers have enabled Sun to team with Zeus Technology to provide a next-generation solution for load balancing and application traffic management — a solution that combines the high-performance throughput and virtualization capabilities of Sun CoolThreads servers with the flexibility and power of Zeus Extensible Traffic Manager (ZXTM) software.

Highlights

- Achieve the highest levels of observed throughput for ZXTM software by combining the flexibility and power of ZXTM software with the high performance of Chip Multithreading (CMT) in Sun CoolThreads servers
- Distribute and load balance incoming traffic to the fastest available servers, increasing service capacity and providing more reliable services
- Accelerate applications by boosting the number of transactions your Web infrastructure can handle
- Improve network application manageability, reducing the time, effort, and cost associated with deploying mission-critical application services
- Easily reallocate server resources for traffic management or other uses to most effectively utilize available server capacity

Service-level challenges in today's Web infrastructures

As today's businesses struggle to survive in a global online marketplace, they must serve millions of users with real-time application services that consistently deliver high service levels. Service-level breaches can be costly, so it's important that the IT infrastructure provide high reliability and the capacity to scale quickly in the face of unpredictable growth. Competitive pressures also require a cost-effective approach to application delivery, so IT infrastructures must be extremely efficient to keep overhead costs low.

These challenges are compounded by the narrow time to market windows in today's competitive markets and the explosive and unpredictable growth of application services. That means Web infrastructures must be flexible and adaptable so that business needs can be met without undue cost. Server and network resources must be able to be quickly and easily reallocated to where they're most needed. IT organizations have turned to server virtualization to consolidate application services to create a streamlined datacenter that's easier to manage and operates within their existing capacity for datacenter space, power, and cooling.

Whether consolidated onto virtualized servers or deployed on independent servers, today's Web applications utilize horizontal scaling, wherein multiple independent instances of

services are deployed to provide sufficient throughput and to offer built-in resilience.

To maintain acceptable service levels in this environment requires distributing incoming traffic across these service instances using a traffic-management solution that can intelligently monitor and manage the requests.

Most traffic-management solutions available on the market today are based on fixed-purpose hardware that can't easily be reconfigured or reprovisioned. In today's fast changing business conditions, these dedicated devices reduce IT flexibility, making it harder for the business to adapt quickly. If unexpected growth occurs, it can sometimes take several weeks to acquire additional load-balancing appliances. IT managers need a more flexible approach to deploying traffic management to maintain service levels during times of rapid growth.

A new approach to traffic management

A combined offering from Sun and Zeus Technology enables an entirely new approach to application traffic management and load balancing. The solution provides extremely high performance and reliable Web-service delivery and enables unsurpassed flexibility in deploying server resources for traffic management.

The traffic-management and load-balancing functions are delivered in software, enabling the solution to provide much greater flexibility than fixed-purpose hardware load-balancing appliances. Consisting of Zeus Technology's

ZXTM software running on Sun CoolThreads servers and the Solaris™ 10 Operating System, the solution offers high levels of sustained throughput—throughput that actually exceeds the capacity of many fixed-purpose hardware solutions.

The primary business benefits of the combined Sun and Zeus offering include:

- Increased agility through immediate capacity expansion or contraction when user demand changes
- Improved visibility and control by inspecting, transforming, and routing requests as they are load-balanced across the application infrastructure
- Increased capacity for the number of transactions that can be handled by the Web infrastructure, accelerating application performance
- High security with Solaris 10 security features and secure isolation within ZXTM that protects against direct attacks and invalid requests
- Reduced time, effort, and cost associated with deploying mission-critical application services
- Greater application throughput with content caching in ZXTM
- Improved user experience with fast response times and high service levels

Virtual switching technology

The first load-balancing solutions to come to market in the early days of Web infrastructures were software-based solutions. However, as Web services continued to expand, dedicated hardware for load balancing became necessary in order to meet the performance throughput and latency requirements. Now the trend has come full circle. Technology advancements in Chip Multithreading (CMT) with the Sun UltraSPARC® T1 and T2 processors provide the capacity for Sun CoolThreads servers to deliver real-time performance for packet processing and request routing, eliminating the need for dedicated hardware.

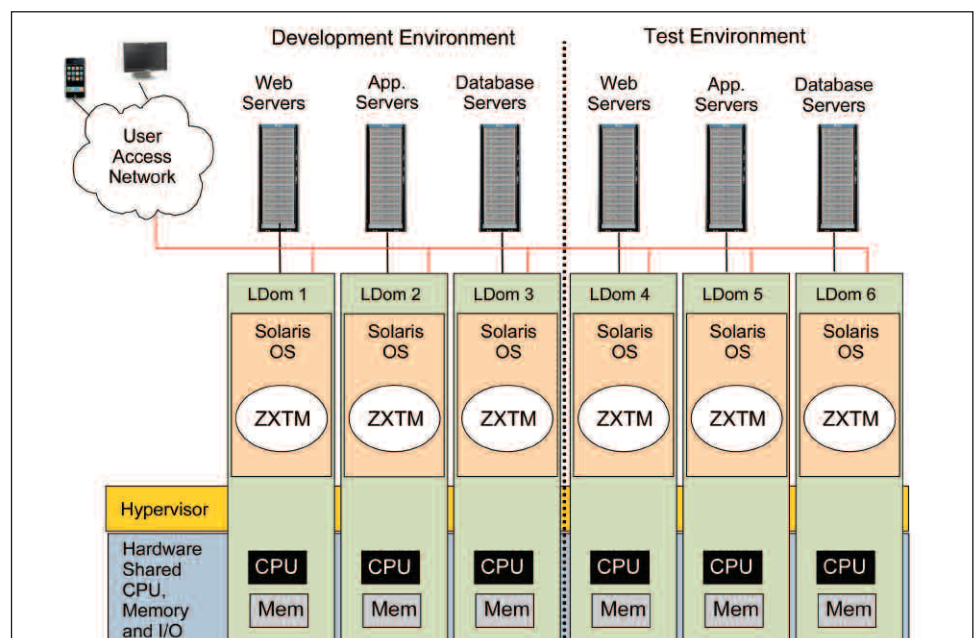
A significant advantage of being able to run load balancing and application traffic management functions on general purpose hardware such as Sun CoolThreads servers is the added flexibility that is gained through the virtual-switch approach: An administrator can quickly and easily change the network topology or reallocate hardware resources without having to reroute cables or move equipment.

The concept of a virtual switch is implemented on Sun CoolThreads servers by partitioning

server resources into secure isolated domains, each having its own set of resources. Each domain runs its own instances of the Solaris OS, an instance of ZXTM software and maintains its own separate IP network stack. As shown in Figure 1, Sun Logical Domains technology effectively divides the server into multiple isolated domains. The illustration shows each domain or virtual switch serving different network zones for Web, application, and database resources.

This virtual-switch capability can be used to separate customers or lines of business, support multiple tiers in Web applications, and implement various network architectures including Demilitarized Zones (DMZs) and extranets. Most importantly, with ZXTM software, configuration of the network topology and the business logic for traffic management is completely under software control. IT organizations can plug once, deploy, and redeploy server resources with no recabling or swapping of hardware.

The key technologies that enable this virtual-switch approach are Sun LDom and Solaris IP Instances.



Sun CoolThreads servers can be dynamically reconfigured to allocate resources across multiple virtual switches.

Logical Domains (LDoms) technology

Sun Logical Domains technology lets you consolidate multiple instances of software components onto a single Sun CoolThreads server while protecting each application instance from faults in other domains. You can allocate the server's resources across multiple LDoms to share those resources while dedicating specific resources to application components in a given domain.

Each logical server's OS, data, and applications are partitioned to run independently of each other while sharing the processing and storage resources of the Sun CoolThreads server.

You can even run different operating systems simultaneously, combining multiple UNIX® and Linux servers onto one Sun CoolThreads server.

Multiple LDoms are created by the addition of a software layer called a hypervisor, interposed between the operating system and the hardware platform. The hypervisor abstracts the hardware and can expose or hide various resources, allowing you to specify LDoms' resource partitions. Each domain is an independent, discrete system, separate from the others, with its own resource-level specification.

LDoms technology is standard on the UltraSPARC T1 and T2 processors, allowing you to deploy one of the industry's most open virtualization solutions with no additional licensing costs.

Virtualization with the Solaris 10 OS

The Solaris 10 OS also contains key virtualization technologies that make it possible to deploy multiple virtual switches in a single server. Virtualization at the network layer is made possible by Solaris IP Instances technology. Solaris IP Instances enable a discrete network stack for each virtual or physical OS instance. This gives you the option of assigning separate IP stacks to Solaris Containers or LDoms, thus providing much more flexibility. Security is maintained because the ability to view and control each instance of the IP network stack is limited by the privileges of the particular container or domain in which the IP stack resides.

If you want to deploy additional virtual zones within LDoms, Solaris Containers technology provides another facility for accomplishing this. Solaris Containers allow you to securely isolate application workloads within a single instance of the Solaris OS, helping to reduce administration overhead by reducing the number of OS instances to be managed.

Solaris 10 is the current version of Sun's tested, certified and supported enterprise operating system, available free for download. Future versions of the Solaris OS will be based on technology from the OpenSolaris™ project, a community-based open source initiative.

ZXTM Software from Zeus Technology

With today's racks of virtualized servers sometimes housing as many as several hundred applications in a single rack, network-level virtualization has become critically important. Virtual IP addresses can be used to automate the configuration of network IP addresses across application components. Routing of traffic across the application infrastructure can then be managed using flexible rules. ZXTM software running on Sun CoolThreads servers provides the ability to perform deep packet inspection and create highly flexible traffic-management rules that enable precise decisions about how to route and handle traffic.

Zeus Technology's ZXTM software can dramatically improve service levels by speeding applications, improving reliability and security, and easing management. The software enables customized business logic for directing incoming requests to the fastest currently available server and routing traffic away from busy or failed services. The ZXTM TrafficScript scripting language gives administrators the flexibility and power to implement traffic-management policies and customize the software by changing caching, access control, persistence, routing, error checking, health monitoring, and logic for rewriting requests and responses.

The traffic manager can also cache responses to common requests, reducing the number of requests that must be relayed to the back-end server. TCP requests and responses are buffered, so the server interacts with a fast local client rather than a slow and unreliable remote client—reducing the number of simultaneous connections and improving the capacity of Web servers and middleware applications.

Key features of ZXTM software include:

- Ability to create custom business logic for Layer 7 traffic management
- Full traffic inspection and service-level monitoring for intelligent routing
- High-availability configurations that utilize multiple ZXTM instances in clusters of active units, or a combination of active and standby units
- Application acceleration through TCP offload and buffering as well as HTTP optimization and full support for XML, XPath, and XSLT translation
- Enhanced security with high-performance SSL technology as well as optional encryption key management
- Simplified management and capacity planning through real-time reporting and self-diagnostics

Sun CoolThreads servers

By combining ZXTM software with Sun CoolThreads servers you can achieve the highest levels of sustained throughput ever observed with ZXTM software and also take advantage of the free and open source Solaris 10 OS and Sun's no-cost virtualization technologies. Sun's CoolThreads server offerings include both state-of-the-art rack servers and density-efficient blade servers:

- Sun SPARC® Enterprise T5120 and T5220 servers and the Netra™ T5220 server are based on the new UltraSPARC T2 processor, with up to 64 active execution threads and integrated dual 10 Gigabit Ethernet technology. These servers offer blistering network performance and provide a highly scalable platform for consolidating multiple instances of ZXTM software. The Netra T5220 server is optimized for carrier-grade networks and is NEBS Level 3 certified.
- The world's first eco responsible servers, the Sun Fire™ T1000 and T2000 servers (also called Sun SPARC Enterprise T1000 and T2000 servers) and Netra T2000 servers are based on the UltraSPARC T1 processor. These servers are a good fit for horizontally scaled Web infrastructures.
- Sun Blade™ Modular Systems, including the new Sun Blade T6320 server module and the NEBS-compliant Netra CP3260 ATCA blade servers, based on the UltraSPARC T2 processor, are designed to deliver increased performance and greater flexibility in a low-cost, high-density solution. These servers and blade servers can help relieve datacenter capacity constraints for the massive horizontal scaling of today's Web infrastructures.

- All systems come with integrated, no-cost, open virtualization technologies, providing the ability to quickly consolidate multiple ZXTM software instances onto rack and blade servers that are highly space and energy efficient.

Sun servers and blades powered by the UltraSPARC T2 "system on a chip" processor are designed for virtualization and eco efficiency. With up to 8 cores and 64 threads per system, they provide the flexibility and power of 64 virtual systems in a single server. These systems can help you realize dramatic space savings, huge reductions in electricity costs, and simplified server administration—all while delivering much faster processing for each consolidated server.

Wire-speed encryption

The integrated on-chip cryptographic acceleration of the UltraSPARC T1 and T2 processors provides high performance for Secure Sockets Layer (SSL) encryption. ZXTM software takes advantage of the Solaris Cryptographic Framework (SCF), which provides cryptographic services for kernel-level and user-level processes. SCF automatically offloads SSL processing for any PKCS11#-compliant application, enabling the application to transparently utilize the powerful hardware acceleration features of the UltraSPARC T1 and T2 processors.

The integrated cryptographic acceleration, coupled with security software features available in the Solaris 10 OS, practically eliminates the performance penalty of encryption and decryption services.

Performance of ZXTM software on Sun CoolThreads servers

ZXTM software is a highly multithreaded application that takes full advantage of the CMT technology in the UltraSPARC T2 processor and the multithreaded performance of the Solaris 10 OS.


The Solaris 10 OS enables ZXTM software instances to deliver excellent throughput on Sun CoolThreads servers by optimizing performance across the 64 threads available on the 8-core UltraSPARC T2 processor. Solaris has been scaling to 64 threads since the introduction of the 64-way Sun Enterprise™ 10000 server more than a decade ago. Performance gains over competitive operating systems are also made possible through innovative new features in the Solaris OS, such as enhanced multi-threaded networking and drivers as well as the DTrace tool for performance analysis.

ZXTM software can leverage up to 64 processor threads on Sun CoolThreads servers to process Web traffic requests in parallel and deliver extremely high overall throughput. CMT technology in the UltraSPARC processor enables ZXTM software to continue processing while a given thread is waiting for memory operations or I/O processing. Each thread within the UltraSPARC T2 processor can handle its own sequence of traffic requests due to the multithreading capabilities in ZXTM software. Fast switching between UltraSPARC processor threads then enables high utilization of the processor, providing extremely high levels of sustained throughput.

“Sun’s latest CoolThreads technology and the Solaris 10 OS form an extremely powerful solution for Zeus Technology’s best-of-breed application delivery controller, ZXTM. ZXTM’s highly scalable architecture enables it to use all of the cores and threads available in Sun’s CMT servers, providing market-leading performance that is nearly double that seen on other platforms.”

David Day

CTO, Zeus Technology



High I/O throughput is also made possible by the fully integrated, dual 10 Gigabit Ethernet network interfaces, which are built into the UltraSPARC T2 processor. The fully integrated dual 10 Gigabit Ethernet network interfaces eliminate PCI-E latencies, helping to accelerate multithreaded application performance and optimize I/O throughput for applications that utilize parallel threads. It provides the critical bandwidth and granular resource allocation required for parallel applications and high-speed networks.

According to Zeus Technology, the performance observed on the Sun SPARC Enterprise T5220 server represents nearly double the best complex application traffic management throughput ever seen with ZXTM software on other platforms. As shown in Figure 2, tests conducted in Sun’s labs showed 4.7 Gb/s sustained throughput for ZXTM running on the Sun SPARC Enterprise T5220 server. ZXTM software quickly reaches its peak throughput and then maintains this level of throughput with minimal variation for the duration of test.

Furthermore, a single Sun ATCA chassis with 12 blades running ZXTM provides 56 Gb/s layer 7 throughput, the most in the industry. And it uses only 200 watts per blade. Compared to the fastest proprietary blade-based traffic management hardware available from competitors, ZXTM on a T5220 server provides 55 percent more throughput per watt, and 10 percent more throughput per rack unit. Multiple ATCA chassis can work together as a single cluster, making ZXTM on Sun ATCA products the industry’s only solution for unlimited linear scaling of traffic management.

The test configuration consisted of a single Sun SPARC Enterprise T5220 server with 8 cores and 2 GB of RAM. The standard installation of ZXTM software was deployed and the operating system environment consisted of Solaris 10 07/08 with recommended patches, and soft-rings enabled. User requests were simulated using a hardware-based traffic-generation tool that emulated 100 clients across 20 Web servers. The client requests were 512 KB packets delivered using the HTTP 1.1 protocol and the

traffic was transferred to and from the SPARC Enterprise T5220 server using two 10 Gigabit Ethernet connections.

These test results clearly demonstrate that the Sun SPARC Enterprise T5220 and the Netra T5220 server can support high sustained throughput via concurrent processing with ZXTM software, eliminating the need for dedicated load-balancing hardware. IT managers now have the flexibility to quickly and easily reconfigure their Web infrastructure to meet changing business needs.

Get started now with a risk-free trial

Take the first step toward next-generation load balancing and application traffic management now. You can try the world's first eco responsible servers, the Sun Fire T1000 and T2000 servers or the new Sun SPARC Enterprise T5120 and T5220 servers, risk-free for 60 days. Also included in the risk-free trials are Netra T2000 and T5220 servers for telco customers. Visit sun.com/tryandbuy/ to apply for your free trial system.

Learn More.

Learn more about all of Sun's Web infrastructure solutions, go to sun.com/newheights and find out more about ZXTM software from Zeus Technology at zeus.com.