



SPARC T-SERIES

An Oracle White Paper
February 2011

Oracle's SPARC T3-1, SPARC T3-2, SPARC T3-4 and SPARC T3-1B Server Architecture

Introduction	1
The Evolution of Oracle's Multicore/Multithreaded Processor Design	2
Business Challenges for Enterprise Applications	2
Rule-Changing Multicore/Multithreading Technology	3
SPARC T3-1, T3-2, T3-4, and T3-1B Servers	6
SPARC T3 Processor	11
The World's First Sixteen-Core Massively Threaded System-on-a-Chip	11
Taking Oracle's Multicore/Multithreaded Design to the Next Level	13
SPARC T3 Processor Architecture	15
Server Architecture	19
System-Level Architecture	19
Chassis Design Innovations.....	21
Oracle's SPARC T3-1 Server Overview.....	22
Oracle's SPARC T3-2 Server Overview.....	24
Oracle's SPARC T3-4 Server Overview.....	27
Oracle's SPARC T3-1B Server Overview	31
Enterprise-Class Management and Software	33
System Management Technology	33
Scalability and Support for Multicore/Multithreading Technology .	35
Fault Management and Predictive Self-Healing.....	40
Multicore/Multithreading Tools: Performance and Rapid Time to Market	41
Conclusion	44

Introduction

Serving the dynamic and growing data center IT services space is challenging for data center operations. Services need to be able to scale rapidly, often doubling capacity in a short period even as they remain highly available. Infrastructure must keep up with these enormous scalability demands, without generating additional administrative burden. Unfortunately, most data centers are already severely constrained by both real estate and power—and energy costs are rising. There is also a new appreciation for the role that the data center plays in reducing energy consumption and pollution. Virtualization has emerged as an extremely important tool as organizations seek to consolidate redundant infrastructure, simplify administration, and leverage underutilized systems. Security too has never been more important, with the increasing price of data loss and corruption. In addressing these challenges, organizations can ill afford proprietary infrastructure that imposes arbitrary limitations.

Employing Oracle's SPARC T3 processor—taking the industry's first massively threaded system-on-a-chip (SoC) to the next level— Oracle SPARC T3-1, SPARC T3-2, SPARC T3-4 and SPARC T3-1B servers offer breakthrough performance and energy efficiency to drive data center infrastructure and address other demanding data center challenges. Fourth-generation multicore, multithreading technology supports up to 128 threads in as little as two rack units (2RU), providing increased computational density while staying within variously constrained envelopes of power and cooling. Very high levels of integration help reduce latency, lower costs, and improve security and reliability. Optimized system design provides support for a wide range of IT services application types. Uniformity of management interfaces and adoption of standards help reduce administrative costs, while an innovative chassis design shared across Oracle's volume servers provides density, efficiency, and economy for modern data centers. With both the processor and Oracle Solaris available under open source licensing, organizations are free to innovate with a worldwide technical community.

The Evolution of Oracle's Multicore/Multithreaded Processor Design

Oracle's UltraSPARC processors have led the industry for years—first, with the introduction of the multithreaded, multicore chip design in the first-generation UltraSPARC T1 processor in 2005 and now with the fourth-generation SPARC T3 processor. By any measure, the first generation multicore, multithreaded processors were an unprecedented success. Delivering up to five times the throughput in a quarter of the space and power, systems using these processors have rapidly been welcomed and accepted. Now fourth-generation multicore, multithreaded technology is evolving rapidly to meet the constantly changing demands of a wide range of enterprise data center applications.

Business Challenges for Enterprise Applications

Organizations across many industries hope to address larger markets, reduce costs, and gain better insights into their customers. At the same time, an increasingly broad array of wired and wireless client devices are bringing network computing into the everyday lives of millions of people. This strong demand has a “pull-through” effect on the IT services that must be satisfied in the data center. These trends are redefining data center scalability and capacity requirements, even as they collide with fundamental real estate, power, and cooling constraints.

Driving Data Center Virtualization and Eco-Efficiency

Coincident with the need to scale services, many data centers recognize the advantages of deploying fewer standard platforms to run a mixture of commercial and technical workloads. This process involves consolidating underused and sprawling server infrastructures with effective virtualization solutions that serve to enhance business agility, improve disaster recovery, and reduce operating costs. This focus can help reduce energy costs and break through data center capacity constraints by improving the amount of realized performance for each watt of power the data center consumes.

Eco-efficiency provides tangible benefits, improving ecology by reducing the carbon footprint to meet legislative and corporate social responsibility goals, even as it improves the economy of the organization paying the electric bill. As systems are consolidated onto more dense and capable computing infrastructure, demand for data center real estate is also reduced. With careful planning, this approach can also improve service uptime and reliability by reducing hardware failures resulting from excess heat load. Servers with high levels of standard reliability, availability, and serviceability (RAS) are now considered a requirement.

Building Out for Web-scale Applications

Web-scale applications engender a new pace and urgency to infrastructure deployment. Organizations must accelerate time to market and time to service, while delivering scalable high-quality and high-performance applications and services. Many need to be able to start small with the ability to scale very

quickly, with new customers and innovative new Web services often implying a doubling of capacity in months rather than years.

At the same time, organizations must reduce their environmental impact by working within the power, cooling, and space available in their current data centers. Operational costs too are receiving new scrutiny, along with system administrative costs that can account for up to 40 percent of an IT budget. Simplicity and speed are paramount, giving organizations the ability to respond quickly to dynamic business conditions. Organizations are also striving to eliminate vendor lock-in as they look to preserve previous, current, and future investments. Open platforms built around open standards help provide maximum flexibility while reducing costs of both entry and exit.

Securing the Enterprise at Speed

Organizations are increasingly interested in securing all communications with their customers and partners. Given the risks, end-to-end encryption is essential to inspire confidence in security and confidentiality. Encryption is also increasingly important for storage, helping to secure stored and archived data even as it provides a mechanism to detect tampering and data corruption.

Unfortunately, the computational costs of increased encryption can increase the burden on already overtaxed computational resources. Security also needs to take place at line speed, without introducing bottlenecks that can impact the customer experience or slow transactions. Solutions must help to ensure security and privacy for clients and bring business compliance for the organization, all without impacting performance or increasing costs.

Rule-Changing Multicore/Multithreading Technology

Addressing these challenges has outstripped the capabilities of traditional processors and systems, and required a fundamentally new approach.

Moore's Law and the Diminishing Returns of Traditional Processor Design

The oft-quoted tenet of Moore's law states that the number of transistors that will fit in a square inch of integrated circuitry will approximately double every two years. For more than three decades the pace of Moore's law has held, driving processor performance to new heights. Processor manufacturers have long exploited these gains in chip real estate to build increasingly complex processors, with instruction-level parallelism (ILP) as a goal. These traditional processors employ very high frequencies along with a variety of sophisticated tactics to accelerate a single instruction pipeline, including

- Large caches
- Superscalar designs
- Out-of-order execution
- Very high clock rates
- Sophisticated branching techniques

- Deep pipelines
- Speculative prefetches

Although these techniques have produced faster processors with impressive-sounding multiple-gigahertz frequencies, they have largely resulted in complex, hot, and power-hungry processors that are not well-suited to the types of workloads often found in modern data centers. In fact, many data center workloads are simply unable to take advantage of the hard-won ILP provided by these processors. Applications with high shared memory and high simultaneous user or transaction counts are typically more focused on processing a large number of simultaneous threads (thread-level parallelism, or TLP) rather than running a single thread as quickly as possible (ILP).

Making matters worse, the majority of ILP in existing applications has already been extracted, and further gains promise to be small. In addition, microprocessor frequency scaling itself has leveled off because of microprocessor power issues. With higher clock speeds, each successive processor generation has seemingly demanded more power than the last, and microprocessor frequency scaling has leveled off in the 2 GHz to 3 GHz range as a result. Deploying pipelined superscalar processors requires more power, limiting this approach by the fundamental ability to cool the processors.

Chip Multiprocessing with Multicore Processors

To address these issues, many in the microprocessor industry have used the transistor budget provided by Moore's law to group two or more conventional processor cores on a single physical die—creating multicore processors, or chip multiprocessors (CMPs). The individual processor cores introduced by many CMP designs have no greater performance than previous single-processor chips, and in fact, have been observed to run single-threaded applications more slowly than single-core processor versions. However, the aggregate chip performance increases since multiple programs (or multiple threads) can be accommodated in parallel (TLP).

Unfortunately, most currently available chip multiprocessors simply replicate cores from existing (single-threaded) processor designs. This approach typically yields only slight improvements in aggregate performance since it ignores key performance issues such as memory speed and hardware thread context switching. As a result, although these designs provide some additional throughput and scalability, they can consume considerable power and generate significant heat—without a commensurate increase in overall performance.

Chip Multithreading

Oracle engineers were early to recognize the disparity between processor speeds and memory access rates. While processor speeds continue to double every two years, memory speeds typically double only every six years. As a result, memory latency now dominates much application performance, erasing even very impressive gains in clock rates. This growing disconnect is the result of memory suppliers focusing on density and cost as their design center, rather than speed.

Unfortunately, this relative gap between processor and memory speeds leaves ultrafast processors idle as much as 85% of the time, waiting for memory transactions to complete. Ironically, as traditional

processor execution pipelines get faster and more complex, the effect of memory latency grows—fast, expensive processors spend more cycles doing nothing. Worse, idle processors continue to draw power and generate heat. It is easy to see that frequency (GHz) is a misleading indicator of real performance.

First introduced with the UltraSPARC T1 processor, multicore & multithreading takes advantage of chip multiprocessing advances, but adds a critical capability—the ability to scale with threads rather than frequency. Unlike traditional single-threaded processors and even most current multicore processors, hardware multithreaded processor cores allow rapid switching between active threads as other threads stall for memory. Figure 1 illustrates the difference between chip multiprocessing, fine-grained hardware multithreading (FG-MT), and multicore/multithreading. The key to this approach is that each core in an Oracle multicore/multithreaded processor is designed to switch between multiple threads on each clock cycle. As a result, the processor's execution pipeline remains active doing real useful work, even as memory operations for stalled threads continue in parallel.




Figure 1. Oracle's multicore/multithreaded approach combines CMP and fine-grained hardware multithreading.

Oracle's multicore/multithreading approach to processor design provides real value since it increases the ability of the execution pipeline to do actual work on any given clock cycle. Use of the processor pipeline is greatly enhanced because a number of execution threads now share its resources. The negative effects of memory latency are effectively masked, because the processor and memory subsystems remain active in parallel to the processor execution pipeline. Since these individual processor cores implement much-simpler pipelines that focus on scaling with threads rather than frequency (emphasizing TLP over ILP), they are also substantially cooler and require significantly less electrical energy to operate. This innovative approach results in a unique processor technology—multiple physical instruction execution pipelines (one for each core), with multiple active thread contexts per core. In addition, SPARC T3 processors feature two execution pipelines per core to further boost scalability.

The SPARC T3 Processor

Unlike complex single-threaded processors, multicore/multithreaded processors use the available transistor budget to implement multiple hardware multithreaded processor cores on a chip die. SPARC T3 processors take the multicore/multithreaded model to the next level, providing up to 16 cores per processor, with each core supporting up to eight threads via two independent pipelines—effectively doubling the throughput of UltraSPARC T2 and T2 Plus processors with minor increases in the clock frequency. In addition, these processors use the increased transistor budget resulting from the use of a 40 nm silicon technology to implement the industry's first massively threaded system-on-a-chip (SoC), with a single processor die hosting:

- Up to 128 threads per processor (up to sixteen cores supporting eight threads per core)
- On-chip Level 1 and Level 2 caches
- Newly designed floating point pipeline per core
- Per core cryptographic acceleration of 12 different ciphers
- Two on-chip 10 Gigabit Ethernet (GbE) interfaces
- Two on-chip PCI Express Generation 2 (PCIe Gen2) interfaces
- Six on-chip cache coherency links and logic

Through SoC design, the SPARC T3 processor significantly enhances the general-purpose nature of the CPU—building in 16 newly-designed floating-point units (one per core). Enhanced floating-point capabilities further open the SPARC T3 to the world of compute-intensive applications as well as the traditionally multicore/multithreaded-friendly data center throughput applications. No-cost security and cryptographic acceleration is provided by the on-chip, per-core streaming accelerators. In addition, the ability to move data in and out of the SPARC T3 processor is significantly aided by two integrated PCIe Generation 2 interfaces and dual 10 GbE interfaces. The SPARC T3 processor also implements cache coherency logic and links on the processor silicon that facilitate a multsocket, glueless system design.

SPARC T3-1, T3-2, T3-4, and T3-1B Servers

Oracle's SPARC T3-1, SPARC T3-2, SPARC T3-4 and SPARC T3-1B servers all are designed to leverage the considerable resources of the SPARC T3 processors in the form of cost-effective, general-purpose platforms (Figure 2). SPARC T3-based servers deliver up to twice the throughput of their predecessors, while leading competitors in terms of performance, performance per watt, and SWaP performance (as evaluated by the Space, Watts, and Performance metric detailed later in this section). SPARC T3-2 servers extend this scalability by adding dual sockets for SPARC T3 processors and considerably large memory support. Further extending this scalability is the quad socket SPARC T3-4 server. All these systems extend the benefits of multicore/multithreaded processor architectures from multithreaded commercial workloads into technical workloads oriented towards floating-point operations.



Figure 2. Oracle's SPARC T3-1, SPARC T3-2, SPARC T3-4, and SPARC T3-1B servers are designed to leverage the considerable resources of the SPARC T3 processor.

Overview

With support for up to 512 threads, large memory, cryptographic acceleration, and integrated on-chip I/O technology, these servers represent an even further departure from traditional system designs. SPARC T3-1, SPARC T3-2, SPARC T3-4 servers and the SPARC T3-1B blade server are ideal for providing high throughput within significant power, cooling, and space constraints. SPARC T3-1 servers and SPARC T3-1B blade servers support a single SPARC T3 processor. The SPARC T3-2 server supports two SPARC T3 processors, and the SPARC T3-4 server supports up to four SPARC T3 processors. All systems feature a compact and power-efficient package. As compute nodes within massive, horizontally scaled environments, SPARC T3-1, T3-2, T3-4, and T3-1B servers can help provide a substantial building block for application tier, Web services, or even high-performance computing (HPC) infrastructure. Network infrastructure applications such as portal, directory, network identity, file service, and backup are all a good fit for these servers.

The SPARC T3-1 server supports a single SPARC T3 processor and provides for throughput as well as expandability, with extra I/O and internal disk options afforded by a 2RU rack mount form factor. SPARC T3-2 servers support two SPARC T3 processors, and SPARC T3-4 servers support quad SPARC T3 processors. With greater I/O and internal disk, typical workloads include demanding mid-tier application server deployments or Web-tier and application-tier consolidation and virtualization projects requiring maximum uptime with future growth and integration into diverse environments. SPARC T3-2 and SPARC T3-4 servers are ideal for online transaction processing (OLTP) database deployments.

Designed to complement each other and the rest of Oracle's server product line, SPARC T3-1, SPARC T3-2, SPARC T3-4, and SPARC T3-1B servers address the dynamic needs of the modern data center.

- **Efficient and predictable scalability.** With support for 128 threads and large memories, SPARC T3-1 and T3-1B servers use the 10 Gigabit Ethernet, I/O, and cryptographic acceleration provided directly by the SPARC T3 processor. This approach provides leading levels of performance and scalability with extremely high levels of power, heat, and space efficiency.

SPARC T3-2 servers extend this breakthrough compute and memory density, delivering up to 256 threads in a single system, while typically consuming less power than an equivalently configured previous-generation system. SPARC T3-2 servers deliver twice the I/O bandwidth of Sun SPARC T5120 and T5220 servers by providing two PCIe root complexes associated with each SPARC T3 processor.

- **Accelerated time to market.** SPARC T3-4 servers running Oracle Solaris provide full binary compatibility with earlier SPARC systems, preserving investments and rapid time to market. The Cool Tools for SPARC help accelerate application selection, profiling, testing, tuning, debugging, and the deployment of key applications on multicore/multithreaded systems. This functionality has been integrated into the Oracle Solaris Studio 12 release.
- **Industry-leading tools for virtualization and consolidation.** Oracle's multicore/multithreaded technology is ideal for consolidation, providing low-level multithreading support for virtualization at every layer of the technology stack. Oracle's Virtual Machine Server for SPARC (OVMSS) technology exploits the SPARC T3 processor's up to 128 threads per socket, offering multiple guest operating system instances. In addition, Oracle Solaris Containers provide virtualization within a single Oracle Solaris instance. The advanced Oracle Solaris ZFS file system provides storage virtualization for storage and considerable scalability.
- **System and data center reliability.** Reliability is key to keeping applications available and costs down. With the greater levels of integration provided by an SoC design, SPARC T3-1, T3-2, T3-4, and T3-1B servers provide commensurately higher levels of reliability, availability, and serviceability (RAS). Lower power consumption and higher performance per watt greatly reduce generated heat loads and the associated issues they cause. Technologies such as Oracle's Solaris Predictive Self Healing are integrated with the hardware, and help keep systems available.
- **A tradition of leading eco-efficiency.** Oracle's Sun Fire and Sun SPARC Enterprise T1000 and Oracle's Sun SPARC Enterprise T2000 servers were the industry's first eco-responsible servers. Oracle's Sun SPARC Enterprise T5120, T5220, T5140, T5240, and T5440 servers continued this tradition by offering the best performance and performance per watt across a wide range of commercial and technical workloads. In addition, Oracle's UltraSPARC T2 and UltraSPARC T2 Plus processors were the first processors to incorporate unique power management features at both core and memory levels of the processor. Oracle's SPARC T3-1, T3-2, T3-4, and T3-1B servers utilizing the SPARC T3 processor take this functionality to an even higher level by achieving increased levels of integration at the processor level as compared to the UltraSPARC T2 and T2 Plus processors.
- **Zero-cost security.** Providing secure communications and data protection has never been more important, with attempted electronic intrusion and theft at an all-time high. With up to eight integrated cryptographic accelerators on each SPARC T3 processor, there simply is no need to send plain text on the network or store plain text in storage systems. SPARC T3-1, T3-2, T3-4, and T3-1B

servers support many more cryptographic operations per second than competitive systems with dedicated cryptographic accelerator cards—all with minimal impact to system overhead.

- **Simplified management.** Each SPARC T3-1, T3-2, T3-4, and T3-1B server provides an Integrated Lights Out Manager (ILOM 3.0) service processor, compatible with Oracle's x64 servers. Integrated Lights Out Manager provides a command-line interface (CLI), a Web-based graphical user interface (GUI), and Intelligent Platform Management Interface (IPMI) functionality to aid out-of-band monitoring and administration. Integrated Lights Out Manager will not provide an Advanced Lights Out Management (ALOM) backward-compatibility mode for administrators: the assignment of user `cli_mode=aloom` is no longer offered and at best, only provided a subset of full ALOM capabilities.

Innovative System Design

Beyond the capabilities of individual systems, Oracle understands that data centers have unique and pressing needs that require attention on the part of system designers. Density, performance, and scalability are all essential considerations, but systems must also be serviceable and fit in with modern data center strategies that consider power, cooling, and serviceability. SPARC T3-1, T3-2, T3-4, and T3-1B servers share an innovative design philosophy that extends across Oracle's volume x64 and SPARC server platforms. Principles of this philosophy include the following.

- **Maximum compute density.** Oracle's volume servers provide leading density in terms of CPU cores, memory, storage and I/O. This focus on density often lets Oracle's 2RU rack mount servers replace competitive 3RU servers, for a 33 percent space savings.
- **Continued investment protection.** Oracle designs for maximum investment protection. Even with breakthrough technology such as chip multithreaded processors, applications simply run without modification.
- **Leading storage capacity.** Oracle's volume servers provide leading density and flexible RAID options. Smaller disk drives and innovations in structure, airway, and carrier design allow more disk capacity in smaller spaces, while enhancing system airflow.
- **Common, shared management.** SPARC T3-1, T3-2, T3-4, and T3-1B servers are designed for ease of management and serviceability with service processors shared by other Oracle volume server platforms. Systems and components are designed for easy identification, and hot-swap components facilitate online replacement.

Table 1 compares the SPARC T3-1, T3-2, T3-4, and T3-1B servers.

TABLE 1. SPARC T3-1/T3-2/T3-4/T3-1B SERVER FEATURES

FEATURE	SPARC T3-1 SERVER	SPARC T3-2 SERVER	SPARC T3-4 SERVER	SPARC T3-1B BLADE SERVER
CPUs	• 16-core 1.65 GHz SPARC T3 processor	• 16-core 1.65 GHz SPARC T3 processor (Dual)	• 16-core 1.65 GHz SPARC T3 processor (Dual or Quad)	• 8- or 16-core 1.65 GHz SPARC T3 processor
Threads	• Up to 128	• Up to 256	• Up to 512	• Up to 128
Memory	• Up to 128 GB	• Up to 256 GB	• Up to 512 GB	• Up to 128 GB

Capacity	(8 GB DDR3 DIMMs)	(8 GB DDR3 DIMMs)	(8 GB DDR3 DIMMs)	(8 GB DDR3 DIMMs)
Maximum Internal Disk Drives	<ul style="list-style-type: none"> Up to 16 HDD (2.5-inch SAS2 300 GB disk drives) RAID 0/1, (5/6+BBWC) 	<ul style="list-style-type: none"> Up to 6 HDD (2.5-inch SAS2S 300 GB disk drives) RAID 0/1 	<ul style="list-style-type: none"> Up to 8 HDD (2.5-inch SAS2 300 GB disk drives) RAID 0/1 	<ul style="list-style-type: none"> Up to 4 HDD (2.5-inch SAS2 300 GB disk drives) Optional RAID Expansion Module
Video	<ul style="list-style-type: none"> One VGA port 	<ul style="list-style-type: none"> One VGA port 	<ul style="list-style-type: none"> One VGA port 	<ul style="list-style-type: none"> One VGA port (dongle)
Removable, Pluggable I/O	<ul style="list-style-type: none"> Slimline DVD-R/W Five USB 2.0 ports 	<ul style="list-style-type: none"> Slimline DVD-R/W Five USB 2.0 ports 	<ul style="list-style-type: none"> No DVD (Done via rKVMS) Four USB 2.0 ports 	<ul style="list-style-type: none"> No DVD (Done via rKVMS) Three USB 2.0 ports
PCI	<ul style="list-style-type: none"> Six x8 PCIe Gen2 slots 	<ul style="list-style-type: none"> Eight x8 PCIe Gen2 slots Two x4 PCIe Gen2 slots 	<ul style="list-style-type: none"> 16 EM x8 PCIe Gen2 slots 	<ul style="list-style-type: none"> Optional Fabric Expansion Module (RAID 0/1 or 5/6) Supports Gen2 PCIe
Ethernet	<ul style="list-style-type: none"> Four onboard Gigabit Ethernet ports (10/100/1000) Two 10 Gigabit Ethernet ports via XAUI combo slots (shared with PCIe) 	<ul style="list-style-type: none"> Four onboard Gigabit Ethernet ports (10/100/1000) Two 10 Gigabit Ethernet ports via XAUI combo slots (shared with PCIe) 	<ul style="list-style-type: none"> Four onboard Gigabit Ethernet ports (10/100/1000) Eight 10 Gigabit Ethernet ports via XAUI 2 QSFP Quad Connectors 	<ul style="list-style-type: none"> Two onboard Gigabit Ethernet ports (10/100/1000) Two optional 10GB XAUI Ethernet ports expansion modules
Power supplies	<ul style="list-style-type: none"> Two hot-swappable AC 1200 W power supplies (N+1 redundancy) 	<ul style="list-style-type: none"> Two hot-swappable AC 2000 W power supplies (N+1 redundancy) 	<ul style="list-style-type: none"> Four hot-swappable AC 2060 W power supplies (N+N redundancy) 	<ul style="list-style-type: none"> Contained within SunBlade 6000 Modular System Chassis
Fans	<ul style="list-style-type: none"> Six hot-swappable fan modules, with counter-rotating fans per module N+1 redundancy 	<ul style="list-style-type: none"> Six hot-swappable fan trays, with counter-rotating fans per module N+1 redundancy 	<ul style="list-style-type: none"> Five hot-swappable fan modules, with counter-rotating fans per module N+1 redundancy 	<ul style="list-style-type: none"> Contained within Blade 6000 Modular System Chassis
Operating System	<ul style="list-style-type: none"> Oracle Solaris 10 Update 8 + MU9, Solaris 10 Update 9 	<ul style="list-style-type: none"> Oracle Solaris 10 Update 8 + MU9, Solaris 10 Update 9 	<ul style="list-style-type: none"> Oracle Solaris 10 Update 8 + MU9, Solaris 10 Update 9 	<ul style="list-style-type: none"> Oracle Solaris 10 Update 8 + MU9, Solaris 10 Update 9

Leading Reliability, Availability, and Serviceability

SPARC T3-1, T3-2, T3-4, and T3-1B servers provide excellent reliability, availability, and serviceability (RAS) characteristics. Highly reliable parts and a relatively low total component count minimize the opportunity for system errors. Dual PCIe root complexes and the ability to configure multiple processors on SPARC T3-2 and T3-4 servers add to resiliency. In addition, these servers include core and thread off-lining capabilities, integrated disk RAID functions, and extensive ECC hardware protection—along with redundant hot-swap disks, power supplies, and fans. The following key design

elements in the SPARC T3-1, T3-2, T3-4, and T3-1B servers are key to improving the dependability of IT services.

- Processor thread and core off-lining and built-in RAID capabilities
- Redundancy and hot-swap components
- Parity protection and error correction capabilities
- System monitoring
- Integrated Lights Out Manager service processor
- Superior energy efficiency
- Robust virtualization technology
- Comprehensive fault management

Space, Watts, and Performance: The SWaP Metric

SPARC T3-1, T3-2, T3-4, and T3-1B servers deliver leading performance across a range of multithreaded workloads and benchmarks. However, with energy and real estate costs and pressures, it is not enough to measure performance in isolation. Delivering the required level of throughput in a fixed space and power envelope is critical. Traditional system-to-system benchmarks are valuable as a way of comparing one system to another, but are limited when it comes to understanding the power and density attributes of the systems being compared. For this reason, Oracle has developed the space, watts, and performance (SWaP) metric. Designed to provide a simple and transparent measure of overall server efficiency, SWaP is calculated using the following formula: $SWaP = \text{Performance} / (\text{Space} * \text{Power Consumption})$ where

- Performance is measured by industry-standard benchmarks
- Space refers to the height of the server in rack units
- Power is measured by watts used by the system, taken during actual benchmark runs or from vendor site planning guides

SPARC T3 Processor

The SPARC T3 processor is the industry's most highly integrated system-on-a-chip, supplying the most cores and threads of any general-purpose processor available, and integrating all key system functions.

The World's First Sixteen-Core Massively Threaded System-on-a-Chip

The SPARC T3 processor eliminates the need for expensive custom hardware and software development by integrating computing, security, and I/O onto a single chip. Binary compatible with earlier SPARC processors, no other processor delivers so much performance in so little space and with

such small power requirements—letting organizations rapidly scale the delivery of new network services with maximum efficiency and predictability. The SPARC T3 processor is shown in Figure 3.



Figure 3. The SPARC T3 processor allows organizations to rapidly scale the delivery of new network services as well as increasingly compute-intensive workloads with maximum efficiency and predictability.

Table 2 provides a comparison between the SPARC T3 and UltraSPARC T2 and T2 Plus processors.

TABLE 2. SPARC T3, ULTRASPARC T2, AND ULTRASPARC T2 PLUS PROCESSOR FEATURES

FEATURE	SPARC T3 PROCESSOR	ULTRASPARC T2 PROCESSOR	ULTRASPARC T2 PLUS PROCESSOR
Cores/Processor	• Up to 16	• Up to 8	• Up to 8
Threads/Core	• 8	• 8	• 8
Threads/Processor	• 128	• 64	• 64
Hypervisor	• Yes	• Yes	• Yes
Sockets Supported	• 1, 2, or 4	• 1	• 2 or 4
Memory	• Two memory controllers • Up to 16 DDR3 DIMMs	• Four memory controllers • Up to 16 FB-DIMMs	• Two memory controllers • Up to 16 or 32 FB-DIMMs
Caches	• 16 KB instruction cache • 8 KB data cache, 6 MB L2 cache (16 banks, 24-way associative)	• 16 KB instruction cache • 8 KB data cache, 4 MB L2 cache (8 banks, 16-way associative)	• 16 KB instruction cache • 8 KB data cache, 4 MB L2 cache (8 banks, 16-way associative)
Technology	• 40 nm technology	• 65 nm technology	• 65 nm technology
Floating Point	• 1 FPU with Mul/Add per core • 8 FPUs per chip	• 1 FPU per core • 8 FPUs per chip	• 1 FPU per core • 8 FPUs per chip
Integer Resources	• 2 integer execution units/core	• 2 integer execution units/core	• 2 integer execution units/core
Cryptography	• Stream processing unit/core • 12 most popular ciphers	• Stream processing unit/core • 10 most popular ciphers	• Stream processing unit/core • 10 most popular ciphers
Additional On-chip Resources	• Dual PCIe interface (x8) • Dual 10 GbE Interfaces • Coherency logic and links	• Dual 10 GbE interfaces • PCIe interface (x8)	• PCIe interface (x8) • Coherency logic and links (4.8 Gb/second)

(6 x 9.6 Gb/second)

* Two-socket implementation represents SPARC T3-2 server, whereas SPARC T3-4 server represents a four-socket implementation.

Taking Oracle's Multicore/Multithreaded Design to the Next Level

When designing the next-generation of Oracle's multicore/multithreaded processors, the in-house design team started with key goals in mind.

- Increasing computational capabilities to meet the growing demand from Web applications by providing twice the throughput of UltraSPARC T2 and T2 Plus processors
- Supporting larger and more diverse workloads with greater floating-point performance
- Powering faster networking to serve new network-intensive content
- Providing end-to-end data center encryption
- Increasing service levels and reducing downtime
- Improving data center capacities while reducing costs

Oracle's multicore/multithreaded architecture is ultimately very flexible, allowing different modular combinations of processors, cores, and integrated components. The considerations listed above drove an internal engineering effort that compared different approaches with regard to making improvements on the successful UltraSPARC T2 and T2 Plus architecture. For example, simply increasing the number of cores would have gained additional throughput, but would have not have addressed floating-point performance.

The final SPARC T3 processor design recognizes that memory latency is truly the bottleneck to improving performance. By increasing the number of cores supported by each processor, designing a new floating-point pipeline, and by further increasing network bandwidth, these processors are able to provide approximately double the throughput of the UltraSPARC T2 and T2 Plus processors.

Each SPARC T3 processor provides up to 16 cores, with each core able to switch between up to eight threads (128 threads per processor) using a modified LRU algorithm for thread choice. In addition, each core provides two integer execution pipelines, so that a single SPARC core is capable of executing two threads at a time. Figure 4 provides a simplified high-level illustration of the thread model supported by a 16 core SPARC T3 processor.

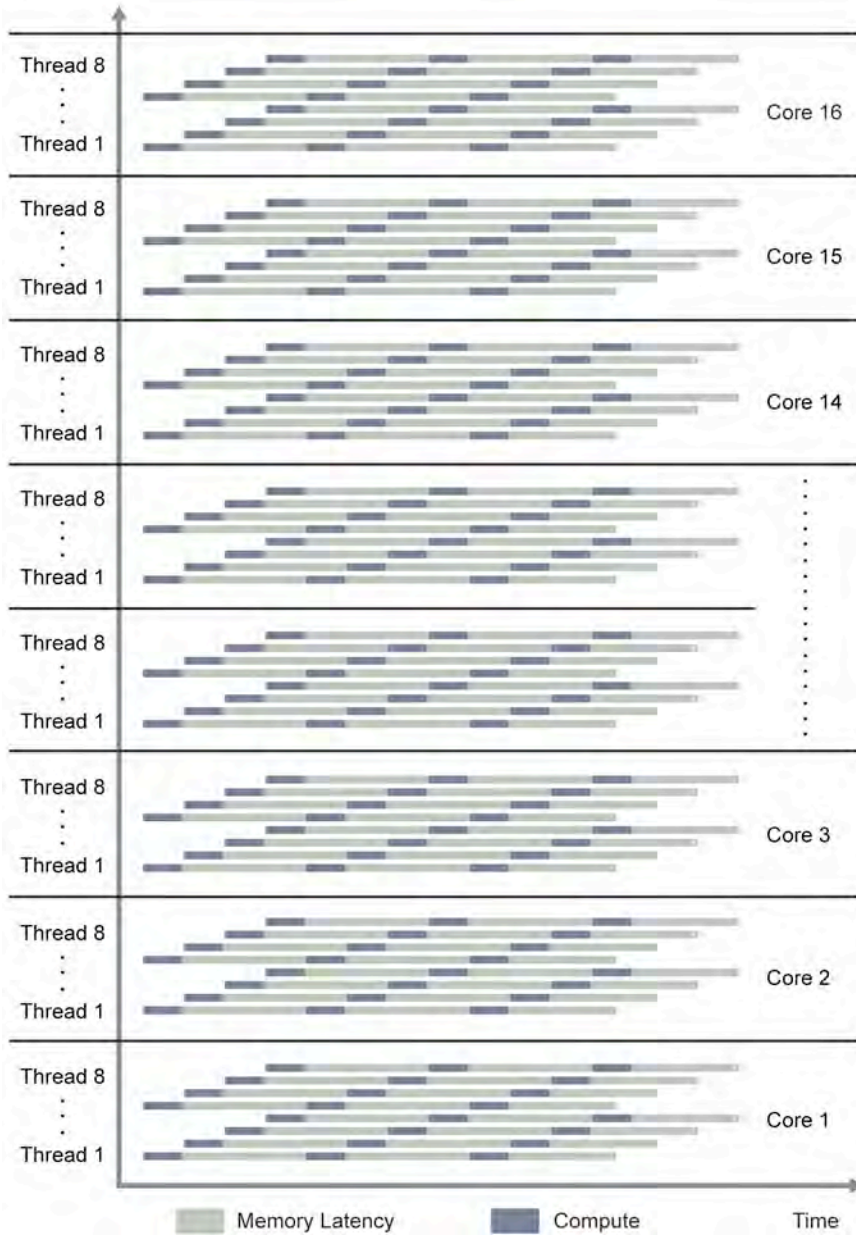


Figure 4. A single 16-core SPARC T3 processor supports up to 128 threads, with up to 2 threads running in each core simultaneously.

SPARC T3 Processor Architecture

The SPARC T3 processor extends Oracle's multicore/multithreaded initiative with an elegant and robust architecture that delivers real performance to applications. Figure 5 provides a block-level diagram of the SPARC T3 processor.

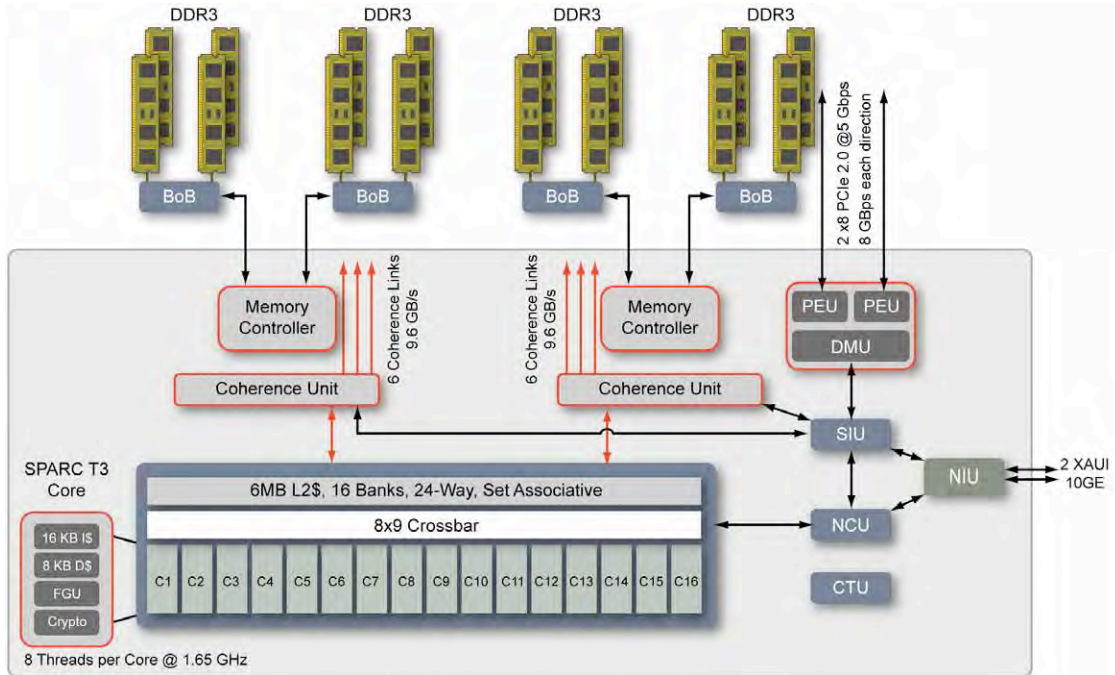


Figure 5. The SPARC T3 processor provides six coherence links to connect to up to four other processors.

The SPARC T3 has coherence link interfaces to allow communication between up to four SPARC T3 processors in a system without requiring any external hub chip. There are six coherence links, each with 14 bits in each direction running at 9.6 Gbps. Each frame has 168 bits, so maximum frame rate is 800M frames per second. The SPARC T3 has two coherence link controllers. Each includes two Coherence and Ordering Units (COU), three Link Framing Units (LFU) and a cross bar (CLX) between COUs and LFUs. Each COU interfaces to two L2 bank pairs. The coherence links run a cache coherence (snoopy) protocol over an FB-DIMM like physical interface. The memory link speed of the SPARC T3 was increased to 6.4 Gb/sec over the UltraSPARC T2 Plus processor's 4.8 Gb/sec, and 4.0 Gb/sec of the UltraSPARC T2 processor.

The SPARC T3 processor can support one-, two- and four-socket implementations. A typical two-socket implementation is shown in Figure 6. Dual-socket, as well quad-socket SPARC T3 implementations interconnect the processors' six coherence links; no additional circuitry is required.

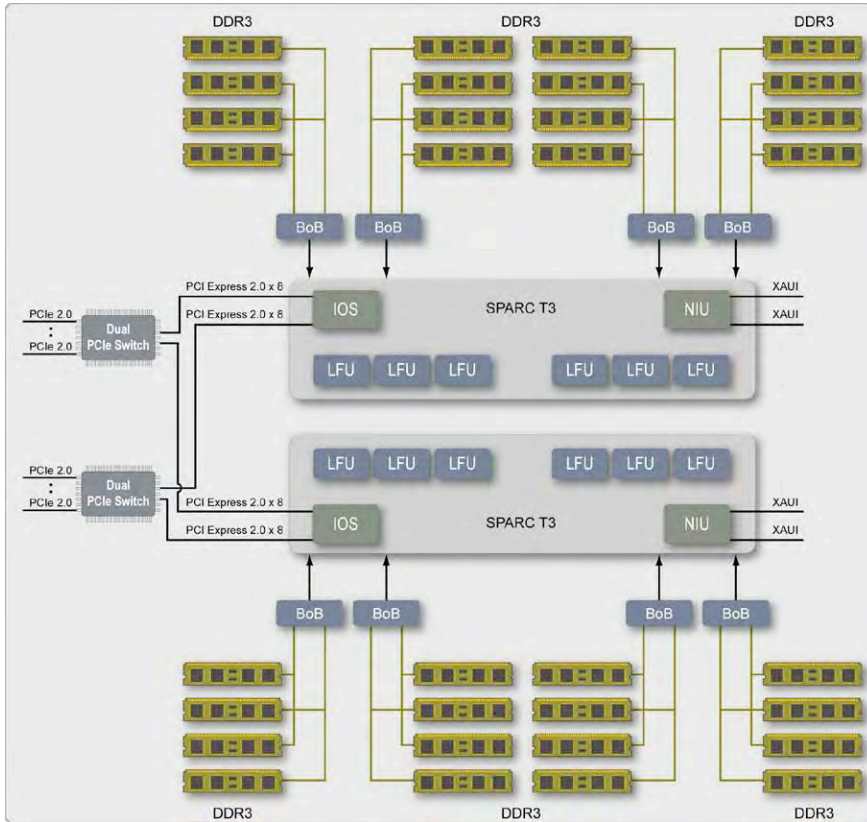


Figure 6. A typical dual-socket SPARC T3 configuration.

Figure 7 provides a block-level diagram representing a single SPARC core on the SPARC T3 processor. Up to 16 cores are supported per processor.

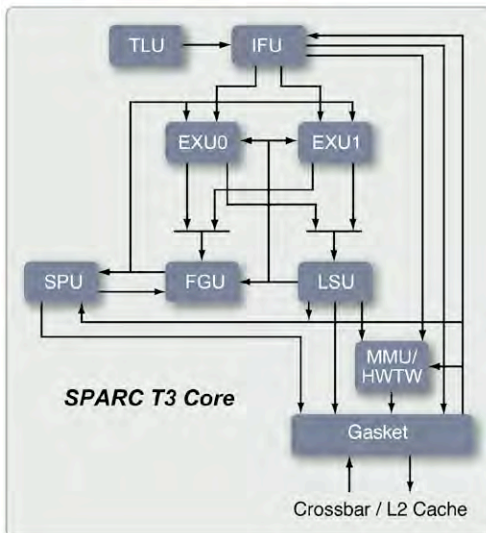


Figure 7. Block-level diagram of a core of the SPARC T3 processor. Components implemented in each core include the following.

- **Trap logic unit.** The trap logic unit (TLU) updates the machine state as well as handling exceptions and interrupts.
- **Instruction fetch unit.** The instruction fetch unit (IFU) includes a 16 KB instruction cache (32-byte lines, 8-way set associative) and a 64-entry fully associative instruction translation lookup buffer (ITLB).
- **Integer execution unit.** Dual integer execution units (EXUs) are provided per core with four threads sharing each unit. Eight register windows are provided per thread, with 160 integer register file (IRF) entries per thread.
- **Floating point/graphics unit.** A floating point/graphics unit (FGU) is provided within each core and it is shared by all eight threads assigned to the core. Thirty-two floating-point register file entries are provided per thread. A fused floating point Mul/Add instruction is implemented.
- **Stream processing unit.** Each core contains a stream processing unit (SPU) that provides cryptographic co-processing.
- **Memory management unit.** The memory management unit (MMU) provides a hardware table walk (HWTW) and supports 8 KB, 64 KB, 4 MB, and 256 MB pages.

An eight-stage integer pipeline and a new 9-stage floating-point pipeline is provided by the SPARC T3 processor core (Figure 8). A *pick* pipeline stage exists to choose two threads (out of the eight possible per core) to execute each cycle.




Figure 8. An 8-stage integer pipeline and a 12-stage floating-point pipeline are provided by the SPARC T3 processor core.

To illustrate how the dual integer pipelines function, Figure 9 depicts the integer pipeline with the load store unit (LSU). The instruction cache is shared by all eight threads within the core. A least-recently-fetched algorithm is used to select the next thread to fetch. Each thread is written into a thread-specific instruction buffer (IB) and each of the eight threads is statically assigned to one of two thread groups within the core.




Figure 9. Threads are interleaved between pipeline stages with very few restrictions (integer pipeline shown, letters depict pipeline stages, numbers depict different scheduled threads)

The pick stage chooses one thread each cycle within each thread group. Picking within each thread group is independent of the other, and a least-recently-picked algorithm is used to select the next thread to execute. The decode state resolves resource conflicts that are not handled during the pick stage. As shown in the illustration, threads are interleaved between pipeline stages with very few restrictions. Any thread can be at the fetch or cache stage, before being split into either of the two thread groups. Load/store and floating-point units are shared between all eight threads. Only one thread from either thread group can be scheduled on such a shared unit.

Integrated Networking

By providing integrated on-chip networking, the SPARC T3 processor is able to provide better networking performance. All network data is supplied directly from and to main memory. Placing networking so close to memory reduces latency, provides higher memory bandwidth, and eliminates inherent inefficiencies of I/O protocol translation. The SPARC T3 processor provides two 10 Gigabit Ethernet ports with integrated serializer/deserializer (SerDes), offering line-rate packet classification at up to 30 million packets/second (based on layer 14 of the protocol stack). Multiple DMA engines (16 transmit and 16 receive DMA channels) match DMAs to individual threads, providing binding flexibility between ports and threads. Virtualization support includes provisions for eight partitions, and interrupts may be bound to different hardware threads.

Stream Processing Unit

The SPU on each core runs in parallel with the core at the same frequency. The cipher/hash unit supports RC4, DES/3DES, AES-128/192/256, MD5, SHA-1, SHA-256 ciphers. Added to the SPARC T3 processor are SHA-384/SHA-512, Kasumi Bulk Cipher, and Galois Field Operations. The SPU is designed to achieve wire-speed encryption and decryption on the processor's 10 GbE ports.

Integral PCI Express Generation 2 Support

SPARC T3 processors provide dual on-chip PCIe Generation 2 interfaces. Each operates at 5 Gtps per x1 lane bi-directionally through a point-to-point dual-simplex chip interconnect, meaning that each x1 lane consists of two uni-directional bit-wide connections, one for northbound and the other for southbound traffic. An integral IOMMU supports I/O virtualization and process device isolation by using the PCIe BUS/Device/Function (BDF) number. The total theoretical I/O bandwidth (for a x8 lane) is 4 GB/sec, with a maximum payload size of 256 bytes per PCIe Gen 2 interface. The actual realizable bandwidth is more likely to be approximately 2.8 GB/sec. An x8 SerDes interface is provided for integration with off-chip PCIe switches.

Power Management

Beyond the inherent efficiencies of Oracle's multicore/multithreaded design, the SPARC T3 processor incorporates unique power management features at both the core and memory levels of the processor. These features include reduced instruction rates, parking of idle threads and cores, and ability to turn off clocks in both cores and memory to reduce power consumption. Substantial innovation is present in the areas of

- Limiting speculation, such as conditional branches not taken
- Extensive clock gating in the data path, control blocks, and arrays
- Power throttling that allows extra stall cycles to be injected into the decode stage

Server Architecture

Oracle's SPARC T3-1, T3-2, T3-4 and T3-1B servers have been designed to provide breakthrough performance while maximizing reliability and minimizing power consumption and complexity. This section details the physical and architectural aspects of these systems.

System-Level Architecture

The System-on-a-Chip (SoC) design of the SPARC T3 processor means that sophisticated system-level functionality can be accomplished with a minimum of high-quality components. The sections that follow describe the architecture of the various systems.

Oracle's SPARC T3-1, T3-2, T3-4, and T3-1B Servers

A separate motherboard design is used in each of Oracle's T3-1, T3-2, and T3-4 servers. Although each system deploys the same type of technologies (SPARC T3 processor, DDR3 Memory, PCIe Generation 2 switches, SAS2, and more) each Motherboard has its own unique design and physical form factor. Common features of the SPARC T3 motherboard (Processor Module for the SPARC T3-4 server) are:

- A minimum of one socket for a SPARC T3 processor

- Memory slots to supply memory for the SPARC T3 server
- A remote to local memory latency ratio of 1.47 (in the case of a 2 or 4 processor system running the `lmbench` read access test)

Memory Subsystem

In Oracle's SPARC T3-1, T3-2, T3-4, and T3-1B servers, the SPARC T3 processor provides on-chip memory controllers that communicate indirectly to DDR3 DIMMs via Buffer-on-Board (BoB) memory interfaces through four high-speed serial links. Two dual-channel FBDIMM memory controller units (MCUs) are provided on the SPARC T3 processor. Each MCU can transfer data at an aggregate rate of 6.4 Gb/sec. Sixteen motherboard memory socket locations on Oracle's SPARC T3-1, 32 DIMM slots broken up over four memory risers, each of which plug into the SPARC T3-2 motherboard, up to 64 motherboard memory socket locations on the SPARC T3-4 and 16 motherboard memory socket locations on the SPARC T3-1B provide sufficient board space for four 1066 MHz DDR3 DIMMs per channel.

I/O Subsystem

Each SPARC T3 processor interfaces through two (x8) PCIe Gen 2 ports capable of operating at 5 GB/sec bidirectionally. In each server, these ports natively interface through the I/O devices through PCIe Gen 2 switch chips, connecting either to PCIe card slots, Express Module slots or to bridge devices that interface with PCIe, such as those listed below.

- **Disk controller.** Disk control is managed by an LSI Logic SAS2008 SAS/SATA controller chip. RAID levels 0, 1, and 10 are supported. The LSI controller chip also drives the DVD (optical drive) in Oracle's SPARC T3-1 and T3-2 servers. (DVDs are not supported in Oracle's SPARC T3-4 or T3-1B servers.)
- **Modular disk backplanes.** Depending on the system, a 6-, 8-, or 16-disk backplane is attached to the LSI disk controller by one or more x4 SAS2 links. Oracle's SPARC T3-1 server supports eight and 16 SAS-2 capable disk backplanes. The 16-disk backplane contains a 6 Gb/sec SAS2 expander. The SPARC T3-2 server supports a SAS-2 capable 6 disk backplane, and the SPARC T3-4 server supports a SAS-2 capable 8-disk backplane.
- **Gigabit Ethernet.** Oracle's SPARC T3-1, T3-2, and T3-4 servers provide four 10/100/1000 Mb/sec Ethernet interfaces on the rear of each chassis, two 10/100/1000 Mb/sec Ethernet interfaces are provided on the SPARC T3-1B blade server.
- **Dual 10 Gigabit Ethernet.** Oracle's SPARC T3 processor provides dual 10 Gigabit Ethernet Attachment Unit Interfaces (XAUI) interfaces.
- **USB.** On all servers, a single-lane PCIe port connects to a PCI bridge device. A second bridge chip converts the 33-bit 66 MHz PCI bus into multiple USB 2.0 ports.

Chassis Design Innovations

Oracle's SPARC T3-1, T3-2, T3-4, and T3-1B servers share basic chassis design elements. This approach not only provides a consistent look and feel across the product line, but it simplifies administration through consistent component placement and shared components. Beyond mere consistency, this approach reflects a data center design focus that places key technology where it can make a difference for the data center.

Enhanced System and Component Serviceability

Finding and identifying servers and components in a modern data center can be challenging. Oracle's SPARC T3-1, T3-2, T3-4 and T3-1B servers are optimized for lights-out data center configurations with easy-to-identify servers and modules. Color-coded operator panels provide easy-to-understand diagnostics and systems are designed for deployment in hot-isle/cold-isle multiracked deployments with both front and rear diagnostic LEDs to pinpoint faulty components. Fault Remind features identify failed components.

Consistent connector layouts for power, networking, and management make moving between Oracle's systems straightforward. All hot-plug components are tool-less and easily available for serviceability. For instance, easy access provides access to fan modules so that fans can be serviced without exposing sensitive components or causing unnecessary downtime.

Robust Chassis, Component, and Subassembly Design

Oracle's volume servers share chassis that are carefully designed to provide reliability and cool operation. Even features such as the hexagonal chassis ventilation holes are designed to provide the best compromise for high strength, maximum airflow, and maximum electronic attenuation. Next-generation hard disk drive carriers leverage the hexagonal ventilation of the chassis and provide a small front plate for greater storage density while increasing airflow to the system.

In spite of their computational, I/O, and storage density, Oracle's servers are able to maintain adequate cooling using conventional technologies. Minimized DC-DC power conversions also contribute to overall system efficiency. This approach reduces generated heat, and introduces further efficiencies to the system.

Minimized Cabling for Maximized Airflow

To minimize cabling and increase reliability, a variety of smaller boards are employed, appropriate to each chassis. These infrastructure boards serve various functions in Oracle's SPARC T3-1, T3-2, and T3-4 servers.

- Power distribution boards distribute system power from multiple power supplies to the major components of the system.
- Fan boards provide connections for power and control for both the primary and secondary fans in the front or rear of the chassis. No cables are required because every fan module plugs directly into one of these PCBs.

- The disk backplane mounts to the disk cages in the chassis, delivering disk data through one or two four-channel, discrete mini-SAS cables from the motherboard.
- The Oracle SPARC T3-1, T3-2, and T3-4 servers all support USB 2.0 interfaces, two each on the front and back of the chassis. Oracle SPARC T3-1, T3-2, T3-4, and T3-1B servers support an internal USB thumb drive. In addition, Oracle SPARC T3-1B servers support two external USB 2.0 interfaces on the front of the blade server.

Oracle's SPARC T3-1 Server Overview

The compact Oracle SPARC T3-1 server provides significant computational power in a space-efficient, low-power 2RU rack mount package. With high levels of price/performance, a low acquisition cost, and tightly integrated high-performance 10 Gigabit Ethernet, this server is ideally suited to the delivery of horizontally scaled transaction and Web services that require extreme network performance, and can function as a very capable HPC compute node. The server is designed to address the challenges of modern data centers with greatly reduced power consumption and a small physical footprint.

Oracle's SPARC T3-1 server has a unique motherboard design (Figure 10).

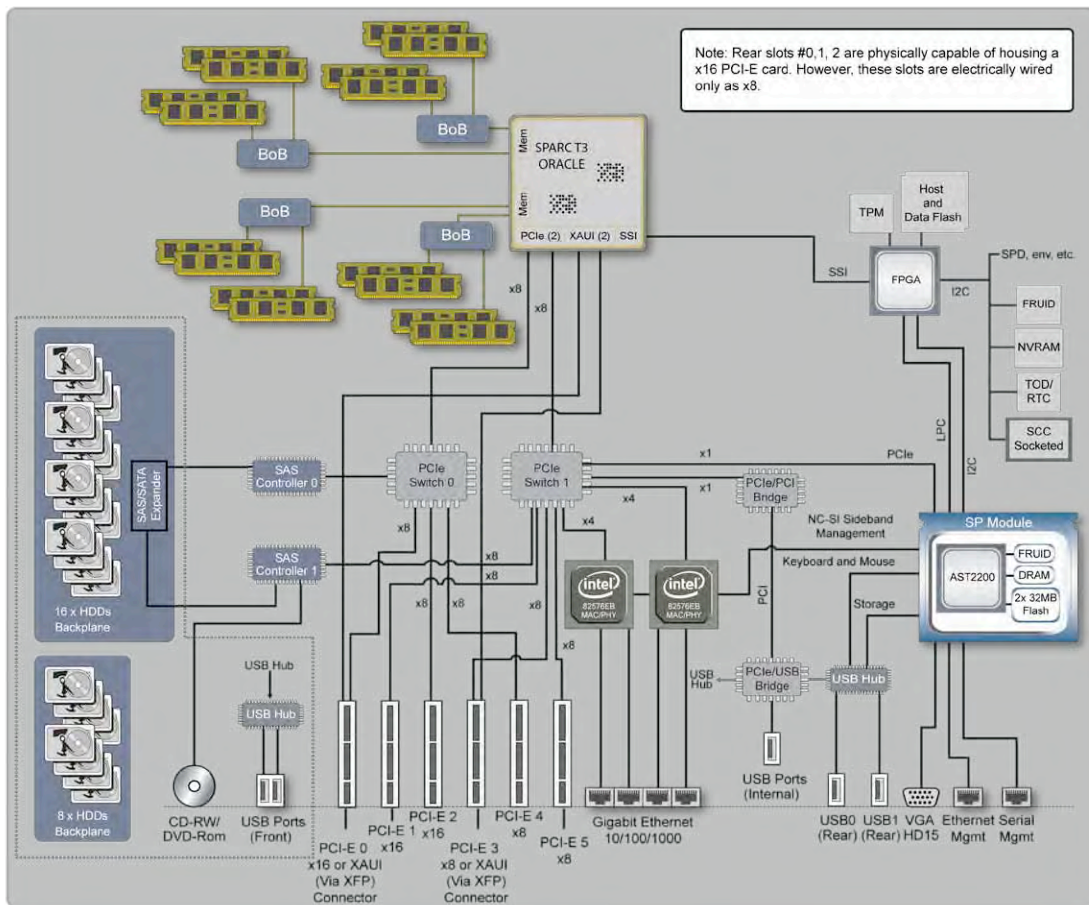


Figure 10. Oracle's SPARC T3-1 server motherboard design.

Enclosure

The 2RU Oracle SPARC T3-1 server enclosure is designed for use in a standard 19-inch rack (Table 3).

TABLE 3. DIMENSIONS AND WEIGHT OF ORACLE'S SPARC T3-1 SERVER

DIMENSION	U.S.	INTERNATIONAL
Height	3.49 inches (2RU)	88.65 millimeters
Width	17.6 inches	447 millimeters
Depth	26.5 inches	673.1 millimeters
Weight (Approximate, without PCIe cards or rack mounts)	60 pounds	27.2 kilograms

The Oracle SPARC T3-1 server includes the following major components.

- An SPARC T3 processor with sixteen cores operating at 1.65 GHz.
- Up to 128 GB of memory in 16 Dual Data Rate slots (2 GB, 4 GB, and 8 GB DDR3 DIMMs supported)
- Four onboard 10/100/1000 Mb/sec Ethernet ports
- Dedicated low-profile PCIe slots (x8)
- Two combination XAUI or low-profile PCIe x4 slots
- Five USB 2.0 ports (two forward, two rear facing, 1 internal)
- Eight or sixteen available disk drive slots support SAS2 disk drives
- ILOM 3.0 system controller
- Two (N+1) hot-swappable, high-efficiency 1200 watt AC power supplies
- Six fan assemblies (each with two fans), under environmental monitoring and control, N+1 redundancy. Fans are accessed through a dedicated top panel door.

Front and Rear Perspectives

Figure 11 illustrates the front and rear panels of Oracle's SPARC T3-1 server, showing both the 16-disk backplanes.




Figure 11. The front and rear panels of Oracle's SPARC T3-1 server.

External features of Oracle's SPARC T3-1 server include the following.

- Front and rear system and component status indicator lights that provide locator (white), service required (amber), and activity status (green) for the system
- Up to eight or 16 hot-plug SAS2 disk drives insertable through the front panel
- One slimline, slot-accessible DVD+RW accessible through the front panel
- Four USB 2.0 ports, two on the front panel, and two on the rear
- Two hot-plug/hot-swap (N+1) power supplies with integral fans insertable from the rear
- Rear power-supply indicator lights that convey the status of each power supply
- A single power plug on each hot-plug/hot-swap power supply
- Four 10/100/1000Base-T autosensing Ethernet ports
- A VGA video port
- A total of six PCIe card slots, two of which can alternately support XAUI cards connected to the SPARC T3 server 10 Gigabit Ethernet interfaces
- Two management ports for use with the ILOM 3.0 system controller; RJ-45 serial management port provides default connection to the ILOM 3.0 controller (network management port supports an optional RJ-45 10/100Base-T connection to the ILOM 3.0 system controller)

Oracle's SPARC T3-2 Server Overview

The expandable SPARC T3-2 server is optimized to deliver transaction and Web services, including Java™ 2 Platform and Enterprise Edition (J2EE platform) technology application services, Enterprise application services such as enterprise resource planning (ERP), customer relationship management (CRM), supply chain management (SCM) and distributed databases. With considerable expansion capabilities and integrated virtualization technologies, Oracle's SPARC T3-2 server is also an ideal platform for consolidated Tier 1 and Tier 2 workloads.

Oracle's T3-2 server has a distinct motherboard design (Figure 12).

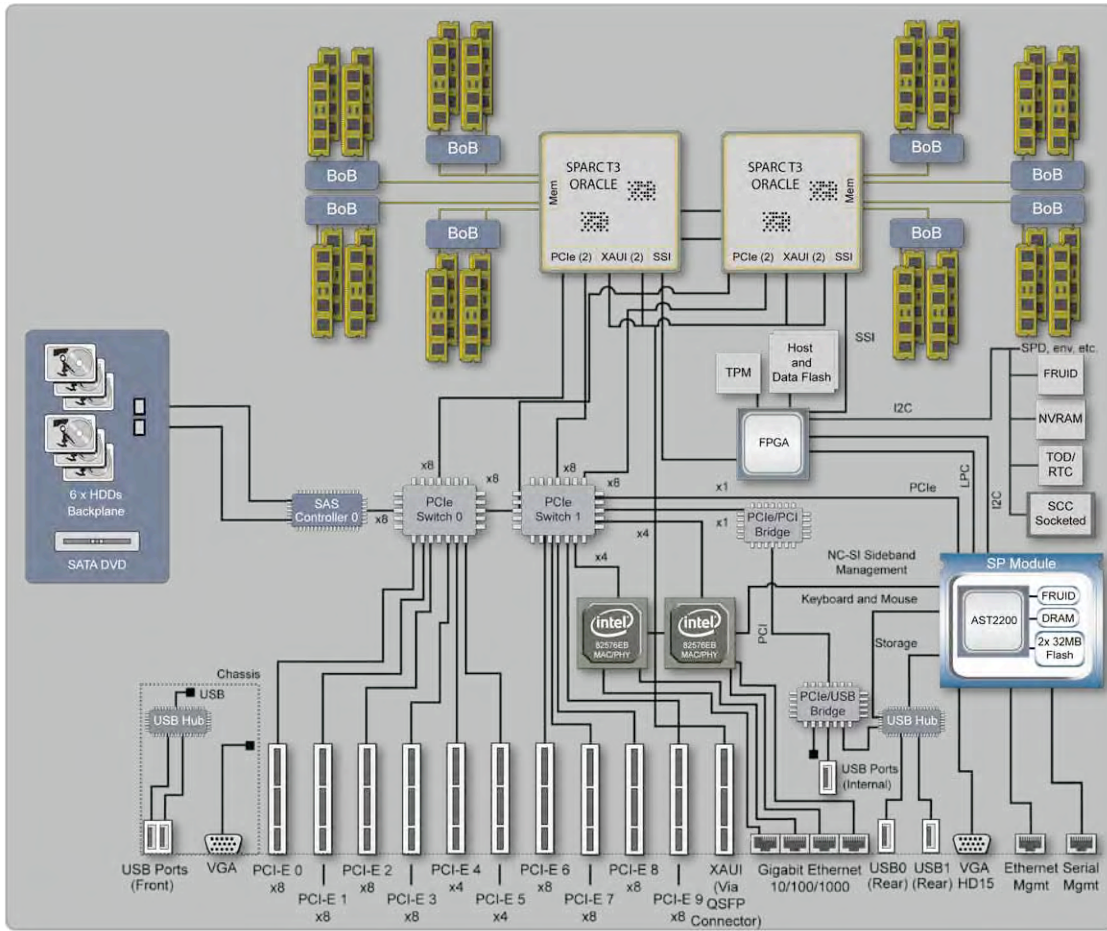


Figure 12. Oracle's SPARC T3-2 server motherboard design.

Enclosure

Oracle SPARC T3-2 servers feature a compact, expandable 3RU rack mount chassis, giving companies the flexibility to scale processing and I/O needs without wasting precious space (Table 4).

TABLE 4. DIMENSIONS AND WEIGHT OF ORACLE'S SPARC T3-2 SERVER

SERVER/DIMENSION	U.S.	INTERNATIONAL
Height	5.11 inches (3 RU)	129.85 millimeters
Width	17.18 inches	436.5 millimeters
Depth	28.81 inches	732 millimeters
Weight (Without PCIe cards or rack mounts)	80 pounds	36.28 kilograms

Oracle's SPARC T3-2 server includes the following major components.

- Dual SPARC T3 processors with sixteen cores per processor operating at 1.65 GHz
- Up to 256 GB of memory in 32 DDR3 DIMM slots (4 GB and 8 GB DDR3 DIMMs)
- Four onboard 10/100/1000 Mb/sec Ethernet ports
- Ten dedicated low-profile PCIe slots
- One dedicated slot for 4x 10GbE XAUI port (This port cannot be shared with any PCIe card).
- Five USB 2.0 ports (two forward, two rear facing, 1 internal restricted to thumb drive)
- 6 available disk drives slots that support SAS2 commodity disk drives
- ILOM 3.0 system controller
- Two (N+1) hot-plug/hot-swap high-efficiency 2000 watt AC power supplies
- Six fan assemblies under environmental monitoring and control, N+1 redundancy

Front and Rear Perspectives

Figure 13 illustrates the front and back panels of Oracle's SPARC T3-2 server.




Figure 13. The front and back panels of Oracle's SPARC T3-2 server.

External features of Oracle's SPARC T3-2 server include the following.

- Front and rear system and component status indicator lights that provide locator (white), service required (amber), and activity status (green) for the system
- 6 hot-plug SAS2 disk drives insertable through the front panel of the system
- One slimline DVD +/-RW drive accessible through the front panel
- Four USB 2.0 ports, two on the front panel, and two on the rear

- Two hot-plug/hot-swap N+1 power supplies with integral plugs and fans insertable from the rear (rear power supply indicator lights convey the status of each power supply)
- Four 10/100/1000Base-T autosensing Ethernet ports
- A VGA video port
- A total of 10 PCIe card slots
- Two management ports for use with the ILOM 3.0 system controller; RJ-45 serial management port provides default connection to the ILOM 3.0 controller (network management port supports an optional RJ-45 10/100Base-T connection to the ILOM 3.0 system controller)

Oracle's SPARC T3-4 Server Overview

With support for up to four SPARCT3 processors and up to 512 threads, the compact Oracle SPARC T3-4 server provides breakthrough computational power in a space-efficient, 5RU rack mount package. With breakthrough levels of price/performance, this server is ideally suited to the delivery of horizontally scaled transaction and Web services as well as medium to large database applications and presents many opportunities as a consolidation and virtualization serve due to its' large capacity. The server is designed to address the challenges of modern data centers that require a compact footprint combined with a greatly increased performance capability as compared to the previous generation of Oracle's Sun SPARC Enterprise T5440 four processor system. Depending on the model selected, the SPARC T3-4 server features dual- or quad-SPARC T3 processors.

Oracle's T3-4 server has a unique motherboard design (Figure 14).

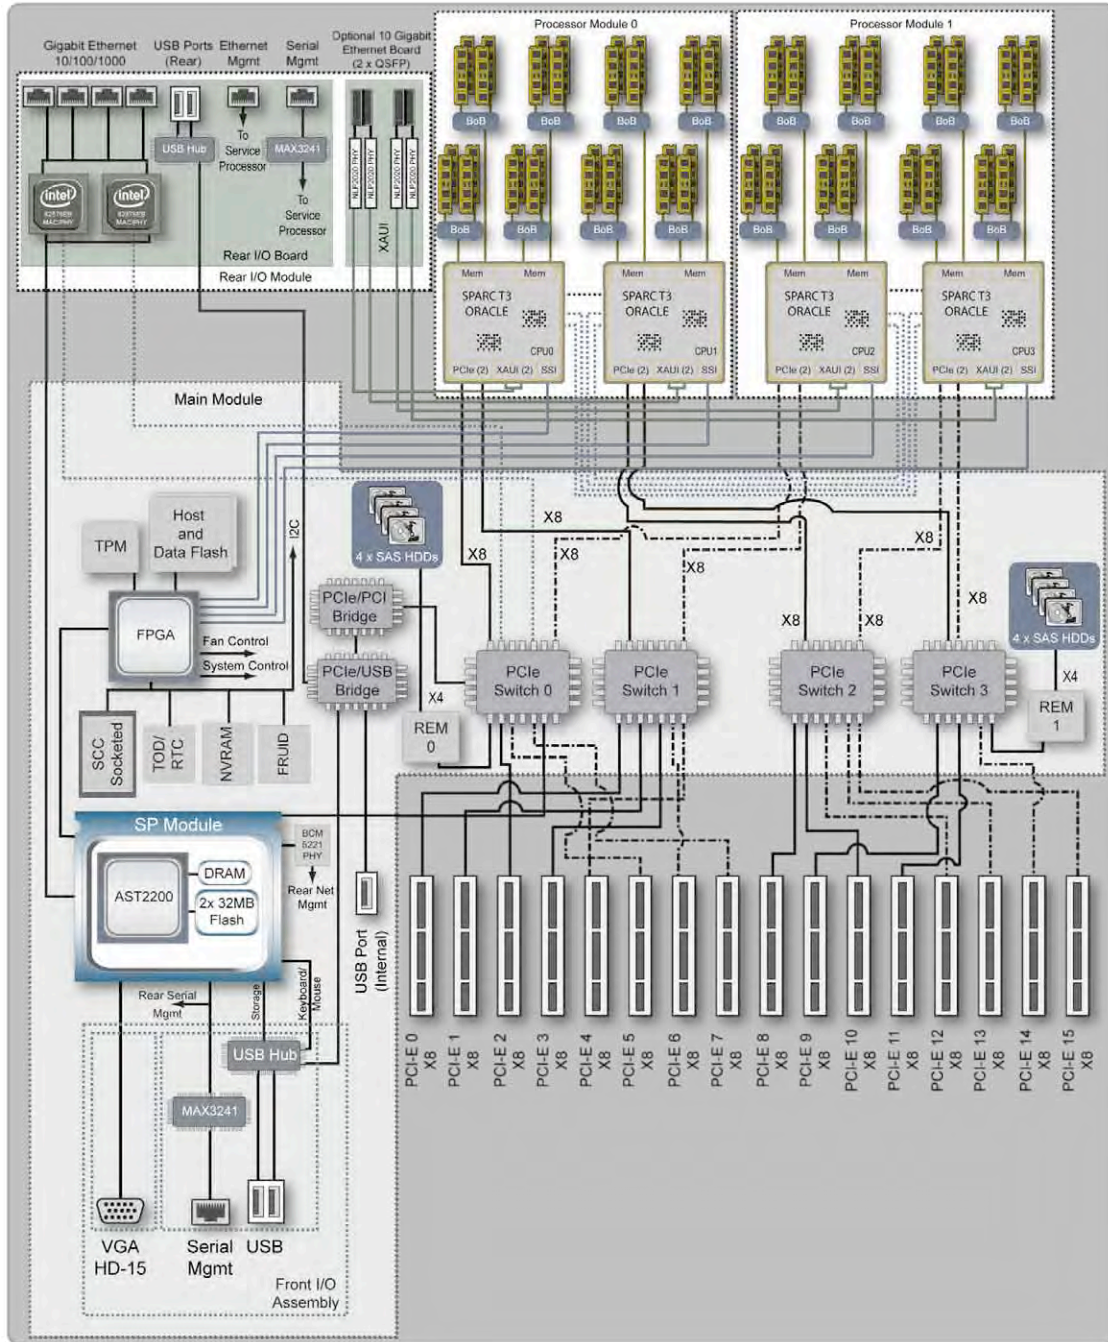


Figure 14. Oracle's SPARC T3-4 server motherboard design.

Enclosure

The 5RU Oracle SPARC T3-4 server enclosure is designed for use in a standard 19-inch rack (Table 5).

TABLE 5. DIMENSIONS AND WEIGHT OF ORACLE'S SPARC T3-4 SERVER

DIMENSION	U.S.	INTERNATIONAL
Height	8.62 inches (5RU)	219 millimeters
Width	17.5 inches	445 millimeters
Depth	27.6 inches	700 millimeters
Weight (Without PCIe cards or rack mounts)	175 pounds	79 kilograms

The Oracle SPARC T3-4 server includes the following major components.

- Two or four SPARC T3 processors sixteen cores operating at 1.65 GHz
- Up to 512 GB of memory in 64 DDR3 DIMM slots (4 GB and 8 GB DDR3 DIMMs are supported)
- Four onboard 10/100/1000 Mb/sec Ethernet ports
- Sixteen PCIe Generation 2 slot x8 (via EMs)
- Eight XAUI via 2 QFSP quad connectors
- Four USB 2.0 ports (two forward, two rear facing)
- Eight available disk drives slots that support SAS2 commodity disk drives
- ILOM 3.0 system controller
- Four (N+N) hot-swappable, high-efficiency 2060 watt AC power supplies
- Five fan assemblies, under environmental monitoring and control, 2 + 2 redundancy

Front and Rear Perspectives

Figure 15 illustrates the front and rear panels of Oracle's SPARC T3-4 server.



Figure 15. The front and rear panels of Oracle's SPARC T3-4 server.

External features of Oracle's SPARC T3-4 server include the following.

- Front
 - Serviced front/rear only (no sliding rack rails)
 - 2 Processor Modules, each with 2x RF CPUs, with 16 or 32 DIMMs, DDR3 DIMMs: 4G, 8G
 - 8 HDDs
 - 4x Power Supplies (2+2)
 - A VGA video port
 - 2 USB ports
 - Console serial port
- Rear
 - Serviced front/rear only (no sliding rack rails)
 - 16 Hot-swap Express Modules
 - Rear IO Module with 4x1G network, 8x10G XAUI network ports
 - 5 Fans (N+1)
 - 4 (2+2) AC cords (200-240V) (any 2 required)

- 2 USB ports
- Console serial port (duplicate of front)
- Console 10/100 network port
- 2 QSFP ports (10G XAUI network)
- 4 1GbE network ports
- LEDs and indicators

Oracle's SPARC T3-1B Server Overview

The Oracle SPARC T3-1B server is optimized to deliver streaming media, virtualization and consolidation, Java application servers, OLTP databases, ERP, CRM, and other back-office applications as well as SOA and business integration. With support for a SPARC T3 processors in a blade server, it is ideal for expansion capabilities, and integrated virtualization technologies, Oracle's SPARC T3-1B server is also an ideal platform for consolidating Tier 1 and Tier 2 workloads.

Oracle's T3-1B server has a unique motherboard design (Figure 16).

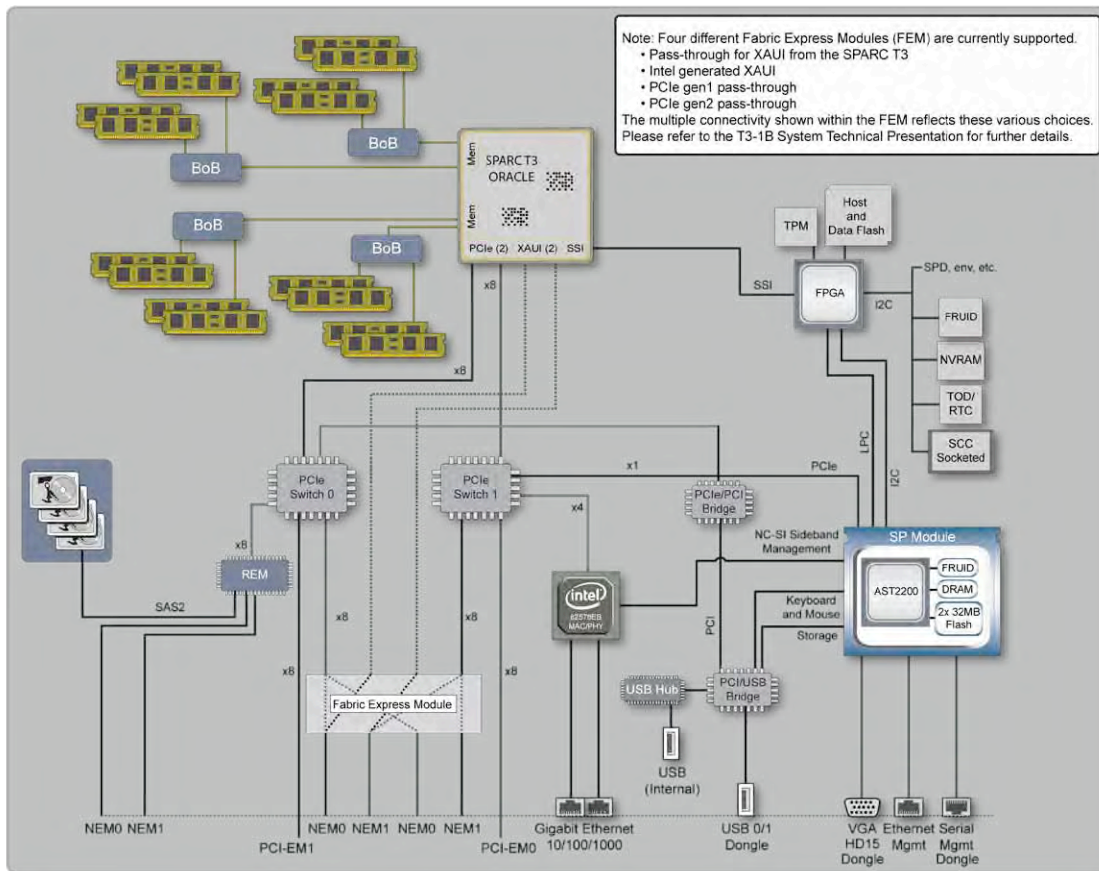


Figure 16. Oracle's SPARC T3-1B server motherboard design.

Enclosure

Oracle's SPARC T3-1B server features a compact blade server, giving organizations the flexibility to scale their processing and I/O by simply adding the SPARC T3-1B to an existing Oracle Sun Blade Chassis 6000 (Table 6).

TABLE 6. DIMENSIONS AND WEIGHT OF ORACLE'S SPARC T3-1B SERVER

SERVER/DIMENSION	U.S.	INTERNATIONAL
Height	1.75 inches (blade)	44.45 millimeters
Width	12.88 inches	327.15 millimeters
Depth	19.56 inches	496.82 millimeters
Weight (With 4 disks, full memory)	20 pounds	9.1 kilograms

Oracle's SPARC T3-1B blade server includes the following major components.

- One SPARC T3 processor with eight or 16 cores operating at 1.65 GHz
- Up to 128 GB of memory in 16 DDR3 DIMM slots (2 GB, 4 GB, and 8 GB DDR3 DIMMs supported)
- Two onboard 10/100/1000 Mb/sec Ethernet ports
- Four dedicated low-profile x8 PCIe slots (all x8 electrically with one x16 physical connector)
- Optional Fabric Expansion Modules (FEM) for PCIe x8 slots
- Three USB 2.0 ports (two external via dongle, 1 internal restricted thumb drive)
- Up to four available disk drive slots supporting commodity SAS2 disk drives
- ILOM 3.0 system controller

Front and Rear Perspectives

Figure 17 illustrates the front and back panels of Oracle's SPARC T3-1B server.



Figure 17. The front and back panels of the SPARC T3-1B server.

External features of the SPARC T3-1B server include the following.

- Front and rear system and component status indicator lights that provide locator (white), service required (amber), and activity status (green) for the system
- Hot-plug SAS2 disk drives insertable through the front panel of the system
- Two USB 2.0 ports accessible via dongle
- A VGA video port
- Two management ports for use with the ILOM 3.0 system controller; RJ-45 serial management port provides default connection to the ILOM 3.0 controller (network management port supports an optional RJ-45 10/100Base-T connection to the ILOM 3.0 system controller)

Enterprise-Class Management and Software

Although new technology often requires time for tools and applications to arrive, delivering agile and highly available services that take advantage of available resources requires stable development tools, operating systems, middleware, and management software. Fortunately, in spite of the breakthrough SPARC T3 processor technology, SPARC T3-1, T3-2, T3-4, and T3-1B servers provide full binary compatibility with earlier SPARC systems and are delivered ready to run with preloaded tools and the solid foundation of Oracle Solaris. Moreover, these systems are provided with a wealth of sophisticated tools that let organizations develop and tune applications as they consolidate and manage workloads while effectively using the resources of the SPARC T3 processor.

System Management Technology

As the number of systems grow in any organization, managing increasingly complex infrastructure throughout its lifecycle becomes difficult. Effective system management requires both integrated hardware that can sense and modify the behavior of key system elements, as well as advanced tools that can automate key administrative tasks.

Integrated Lights Out Manager

Provided across many of Oracle's x64 servers, the Integrated Lights Out Manager service processor acts as a system controller, facilitating remote management and administration of SPARC T3-1, T3-2, T3-4, and T3-1B servers. The service processor is fully featured and is similar in implementation to that used in Oracle's other modular and rack mount x64 servers. As a result, these servers integrate easily with existing management infrastructure. Critical to effective system management, the Integrated Lights Out Manager service processor:

- Implements an IPMI 2.0 compliant services processor, providing IPMI management functions to the server's firmware, OS and applications, and to IPMI-based management tools accessing the service processor via the ILOM 3.0 Ethernet management interface, providing visibility to the environmental sensors (both on the server module, and elsewhere in the chassis)

- Manages inventory and environmental controls for the server, including CPUs, DIMMs, and power supplies, and provides HTTPS, CLI, and SNMP access to this data
- Supplies remote textual console interfaces
- Provides a means to download upgrades to all system firmware

The Integrated Lights Out Manager service processor also allows the administrator to remotely manage the server, independent of the operating system running on the platform and without interfering with any system activity. Integrated Lights Out Manager can also send e-mail alerts of hardware failures and warnings, as well as other events related to each server. The Integrated Lights Out Manager circuitry runs independently from the server, using the server's standby power. As a result, ILOM 3.0 firmware and software continue to function when the server operating system goes offline, or when the server is powered off. Integrated Lights Out Manager monitors the following SPARC T3-1, T3-2, T3-4, and T3-1B conditions.

- CPU temperature conditions
- Hard drive presence
- Enclosure thermal conditions
- Fan speed and status
- Power supply status
- Voltage conditions
- Oracle Solaris watchdog, boot time-outs, and automatic server restart events

Oracle Enterprise Manager Ops Center

Beyond local and remote management capabilities, data center infrastructure needs to be agile and flexible, allowing not only fast deployment but streamlined redeployment of resources as required. Oracle Enterprise Manager Ops Center technology provides an IT infrastructure management platform for integrating and automating management of thousands of heterogeneous systems. To improve lifecycle and change management, Oracle Enterprise Manager Ops Center supports the management of applications and the servers on which they run, including the SPARC Enterprise T3-1, T3-2, T3-4, and T3-1B servers. Oracle Enterprise Manager Ops Center takes a step-by-step approach to unraveling the challenges of getting systems operational quickly.

- **Discover.** As systems are added to the management network, administrators can use Oracle Enterprise Manager Ops Center to discover bare metal systems based on a given subnet address or IP range.
- **Group.** Given the number of systems to manage and the constant repurposing of systems, it is critical for IT organizations to find ways to group resources together. Oracle Enterprise Manager Ops Center lets users logically group systems together and perform actions across a group of systems as easily as performing actions on a single system. Systems can be grouped by function (for

example, Web servers versus clusters servers), administrative responsibility, or other categorization based on organizational needs.

- **Provision.** Oracle Enterprise Manager Ops Center remotely installs operating systems onto selected systems. Administrators can use this functionality to provision operating systems onto bare-metal systems or reprovision existing systems. As the infrastructure lifecycle continues, Oracle Enterprise Manager Ops Center can update firmware and provision software packages and patches to selected systems.
- **Monitor.** When systems are up and running, administrators can use Oracle Enterprise Manager Ops Center to monitor system health, helping to ensure that everything is running at the optimal levels. The software provides detailed hardware monitoring for attributes such as fans, temperature, disk, and voltage usage, including bare-metal systems. Oracle Enterprise Manager Ops Center also monitors OS attributes such as swap space, CPU, memory, and file systems. Administrators can define specific threshold levels and set preferred notification methods, including e-mail, pager, or Simple Network Management Protocol (SNMP) traps, for each monitored component as business needs demand.
- **Manage.** Businesses require that infrastructure lifecycle management extend beyond just deploying and monitoring systems. Oracle Enterprise Manager Ops Center includes Lights Out Management capabilities, such as powering systems on and off, and remote serial console access to help IT organizations manage their IT infrastructure from remote locations. Leveraging a role-based access control (RBAC) feature, organizations can grant permissions to specific users to perform specific management tasks.
- **Hybrid user interface.** Oracle Enterprise Manager Ops Center offers users a hybrid user interface (UI), accessible from the Web, that integrates both the GUI and CLI into one console. With this hybrid UI, operations performed in the GUI are simultaneously reflected in the CLI, and vice versa.

Scalability and Support for Multicore/Multithreading Technology

Oracle Solaris 10 is specifically designed to deliver the considerable resources of SPARC T3 processor-based systems. In fact, Oracle Solaris 10 provides key functionality for virtualization, optimal use high availability, unparalleled security, and extreme performance for both vertically and horizontally scaled environments. Oracle Solaris 10 runs on a broad range of SPARC and x86/x64-based systems, and compatibility with existing applications is guaranteed. One of the most attractive features of systems based on SPARC T3 processors is that they appear as a familiar SMP system to the Solaris OS and the applications it supports. In addition, Oracle Solaris 10 has incorporated many features to improve application performance on Oracle's multicore/multithreaded architectures.

- **Multicore/multithreaded awareness.** Oracle Solaris 10 is aware of the SPARC T3 processor hierarchy so that the scheduler can effectively balance the load across all the available pipelines. Even though it exposes each of these processors as 64 logical processors, Oracle Solaris understands the correlation between cores and the threads they support, and provides a fast and efficient thread implementation.

- **Fine-granularity manageability.** For the SPARC T3 processor, Oracle Solaris 10 has the ability to enable or disable individual cores and threads (logical processors). In addition, standard Oracle Solaris features such as processor sets provide the ability to define a group of logical processors and schedule processes or threads on them.
- **Binding interfaces.** Oracle Solaris allows considerable flexibility in that processes and individual threads can be bound to either a processor or a processor set, if required or desired.
- **Support for virtualized networking and I/O, and accelerated cryptography.** Oracle Solaris contains technology to support and virtualize components and subsystems on the SPARC T3 processor, including support for the on-chip 10 GbE ports and PCIe interface. As a part of a high-performance network architecture, Oracle multicore/multithreaded-aware device drivers are provided so that applications running within virtualization frameworks can effectively share I/O and network devices. Accelerated cryptography is supported through the cryptographic framework in Oracle Solaris.
- **Nonuniform memory access optimization in Oracle Solaris.** With memory managed by each SPARC T3 processor on the SPARC T3-2/T3-4 servers, the implementation represents a nonuniform memory access (NUMA) architecture. In NUMA architectures, the speed needed for a processor to access its own memory is slightly different than that required to access memory managed by another processor. Oracle Solaris provides technology that can specifically help applications improve performance on NUMA architectures.
- **Memory Placement Optimization.** Oracle Solaris 10 uses Memory Placement Optimization (MPO) to improve the placement of memory across the physical memory of a server, resulting in increased performance. Through MPO, Oracle Solaris 10 works to help ensure that memory is as close as possible to the processors that access it, while still maintaining enough balance within the system. As a result, many database and HPC applications are able to run considerably faster with MPO.
- **Hierarchical Lgroup Support.** Hierarchical Lgroup Support (HLS) improves the MPO feature in Oracle Solaris. HLS helps Oracle Solaris optimize performance for systems with more-complex memory latency hierarchies. HLS lets the Solaris OS distinguish between the degrees of memory remoteness, allocating resources with the lowest-possible latency for applications. If local resources are not available by default for a given application, HLS helps Oracle Solaris allocate the nearest remote resources.
- **Oracle Solaris ZFS.** Oracle Solaris ZFS offers a dramatic advance in data management, automating and consolidating complicated storage administration concepts and providing unlimited scalability with the world's first 128-bit file system. Oracle Solaris ZFS is based on a transactional object model that removes most of the traditional constraints on I/O issue order, resulting in dramatic performance gains. Oracle Solaris ZFS also provides data integrity, protecting all data with 64-bit checksums that detect and correct silent data corruption.
- **A secure and robust enterprise-class environment.** Best of all, Oracle Solaris does not require arbitrary sacrifices. Existing SPARC applications continue to run unchanged on SPARC T3

platforms, protecting investments. Certified multilevel security protects Oracle Solaris environments from intrusion. The fault management architecture in Oracle Solaris means that elements such as Oracle Solaris predictive self-healing can communicate directly with the hardware to help reduce both planned and unplanned downtime. Effective tools such as Oracle Solaris DTrace help organizations tune their applications to get the most of the system's resources.

End-to-End Virtualization Technology

Virtualization technology is increasingly popular as organizations strive to consolidate disparate workloads onto fewer, more-powerful systems, while increasing use. SPARC T3-1, T3-2, T3-4, and T3-1B servers are designed specifically for virtualization, providing very fine-grained division of multiple resources—from processing to virtualized networking and I/O. Most important, Oracle's virtualization technology is provided as a part of the system, not an expensive add-on.

A Multithreaded Hypervisor

Like the prior UltraSPARC T1, UltraSPARC T2, and UltraSPARC T2 Plus processors, the SPARC T3 processor offers a multithreaded hypervisor—a small firmware layer that provides a stable virtual machine architecture that is tightly integrated with the processor. Multithreading is crucial, because the hypervisor interacts directly with the underlying multicore/multithreading processor. This architecture is able to context switch between multiple threads in a single core, a task that would require additional software and considerable overhead in competing architecture.

Corresponding layers of virtualization technology are built on top of the hypervisor, as shown in Figure 18. The strength of Oracle's approach is that all the layers of the architecture are fully multithreaded, from the processor up through applications that use the fully threaded Java application model. Far from new technology, Oracle Solaris has provided multithreading support since 1992. This experience has helped to inform technology decisions at other levels, ultimately resulting in a system that parallelizes and virtualizes at every level. In addition to the processor and hypervisor, Oracle provides fully multithreaded networking and the fully multithreaded Oracle Solaris ZFS file system. Oracle VM Server for SPARC (previously called Sun Logical Domains), Oracle Solaris Containers, and multithreaded applications are able to receive exactly the resources they need.

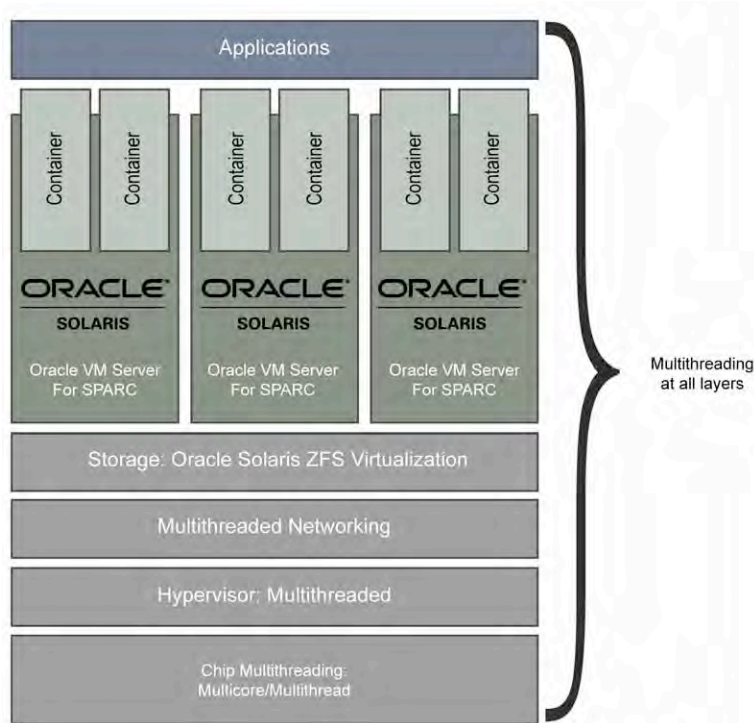


Figure 18. Oracle provides parallelization and virtualization at every level of the technology stack.

Oracle VM Server for SPARC

Supported in all servers from Oracle using Oracle's multicore/multithreaded technology, Oracle VM Server for SPARC provides full virtual machines that run an independent operating system instance, and contain virtualized CPU, memory, storage, console, and cryptographic devices. Within the Oracle VM Server for SPARC architecture, operating systems such as Oracle Solaris 10 are written to the hypervisor, which provides a stable, idealized, and virtualizable representation of the underlying server hardware to the operating system in each domain. Each domain is completely isolated, and the maximum number of virtual machines created on a single platform relies upon the capabilities of the hypervisor, rather than the number of physical hardware devices installed in the system. For example, Oracle's SPARC T3-1 server with a single SPARC T3 processor supports up to 128 domains¹, and each individual domain can run a unique OS instance.

By taking advantage of domains, organizations gain the flexibility to deploy multiple operating systems simultaneously on a single platform. In addition, administrators can leverage virtual device capabilities to transport an entire software stack hosted on a domain from one physical machine to another.

¹ Though possible, this practice generally is not recommended.

Domains can also host Oracle Solaris Containers to capture the isolation, flexibility, and manageability features of both technologies. Deeply integrating Oracle VM Server for SPARC with the SPARC T3 processor, Oracle Solaris 10 increases flexibility, isolates workload processing, and improves the potential for maximum server utilization.

The Oracle VM Server for SPARC architecture includes underlying server hardware; hypervisor firmware; virtualized devices; and guest, control, and service domains. The hypervisor firmware provides an interface between each hosted operating system and the server hardware. An operating system instance controlled and supported by the hypervisor is called a guest domain. Communication to the hypervisor, hardware platform, and other domains for creation and control of guest domains is handled by the control domain. Guest domains are granted virtual device access via a service domain, which controls both the system and hypervisor, and also assigns I/O.

To support virtualized networking, Oracle VM Server for SPARC implements a virtual Layer 2 switch, to which guest domains can be connected. Each guest domain can be connected to multiple vswitches and multiple guest domains can also be connected to the same vswitch. Vswitches can either be associated with a real physical network port, or they may exist without an associated port, in which case the vswitch provides only communications between domains within the same server. This approach also gives guest domains a direct communication channel to the network (Figure 19). Each guest domain believes it owns the entire NIC and the bandwidth it provides, yet in practice only a portion of the total bandwidth is allotted to the domain. As a result, every NIC can be configured as demand dictates, with each domain receiving bandwidth on an as-needed basis. Dedicated bandwidth can be made available by tying a vswitch device to a dedicated physical Ethernet port.

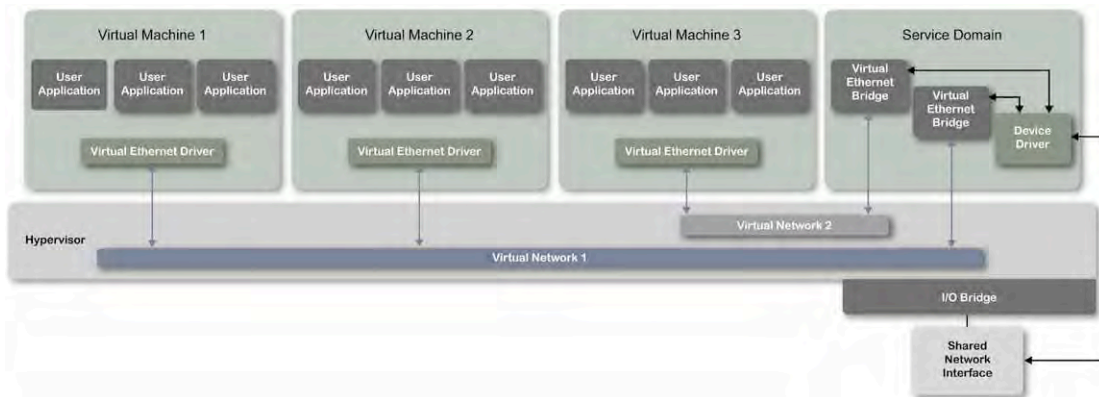


Figure 19. Data moves directly between a Virtual Machine and a virtualized device.

Oracle Solaris Containers

Oracle Solaris 10 provides a unique partitioning technology called Oracle Solaris Containers that can be used to create an isolated and secure environment for running applications. A Container is a virtualized operating system environment created within a single instance of Oracle Solaris. Containers can be used to isolate applications and processes from the rest of the system. This isolation helps enhance

security and reliability since processes in one zone are prevented from interfering with processes running in another zone.

Resource management tools provided with the Oracle Solaris help allocate resources such as CPUs to specific applications. CPUs in a multiprocessor system (or threads in the SPARC T3 processor) can be logically partitioned into processor sets and bound to a resource pool, which in turn can be assigned to a Container. Resource pools provide the capability to separate workloads so that consumption of CPU resources do not overlap, and also provide a persistent configuration mechanism for processor sets and scheduling class assignment. In addition, the dynamic features of resource pools enable administrators to adjust system resources in response to changing workload demands.

Oracle Solaris 10 cryptographic frameworks at user and kernel level provides applications a direct interface to several security ciphers. These ciphers can be implemented at software or hardware levels. With the Oracle's multicore/multithreaded processors and with these ciphers being implemented in the chip, the applications running on Solaris can transparently get to the crypto service at the hardware level. This drastically reduces not only the development of secured applications, but also the load on the system (CPU) when deployed. This helps developers create secured applications for a large scale deployment.

Fault Management and Predictive Self-Healing

Oracle Solaris 10 introduced a new architecture for building and deploying systems and services capable of fault management and predictive self-healing. The predictive self-healing feature in Oracle Solaris 10 is an innovative capability that automatically diagnoses, isolates, and recovers from many hardware and application faults. As a result, business-critical applications and essential system services can continue uninterrupted in the event of software failures, major hardware component failures, and even software misconfiguration problems.

- **Oracle Solaris fault manager.** The fault manager feature in Oracle Solaris collects data relating to hardware and software errors. This facility automatically and silently detects and diagnoses the underlying problem, with an extensible set of agents that automatically respond by taking the faulty component offline. Easy-to-understand diagnostic messages link to articles in Oracle's knowledge base to clearly guide administrators through corrective tasks that require human intervention. The open design of the fault manager feature also permits administrators and field personnel to observe the activities of the diagnostic system. With fault manager, the overall time from a fault condition to automated diagnosis to any necessary human intervention is greatly reduced, increasing application uptime.
- **Oracle Solaris service manager.** The service manager feature in Oracle Solaris creates a standardized control mechanism for application services by turning them into first-class objects that administrators can observe and manage in a uniform way. These services can then be automatically restarted if they are accidentally terminated by an administrator, if they are aborted as the result of a software programming error, or if they are interrupted by an underlying hardware problem. In addition, service manager reduces system boot time by as much as 75 percent by starting services in parallel according to their dependencies. An "undo" feature helps safeguard against human errors by

permitting easy change rollback. The service manager feature is also simple to deploy; developers can convert most existing applications to take full advantage of service manager capabilities by simply adding a simple XML file to each application.

Predictive self-healing and fault management provide the following specific capabilities on the SPARC T3-1, T3-2, T3-4, and T3-1B servers.

- CPU offlining takes a core or threads offline that has been deemed faulty. Offlined CPUs are stored in the resource cache and stay offline on reboot unless the processor has been replaced, in which case the CPU is cleared from the resource cache.
- Memory page retirement retires pages of memory that have been marked as faulty. Pages are stored in the resource cache and stay retired on reboot unless the offending DIMM has been replaced, in which case affected pages are cleared from the resource cache.
- I/O retirement logs errors and faults.
- `fmlog` logs faults detected by the system.

Multicore/Multithreading Tools: Performance and Rapid Time to Market

No matter how compelling new hardware or OS platforms may be, organizations must be ensured that the costs and risks of adoption are in line with the rewards. In particular, organizations want to be able to continue to leverage the considerable advantages of popular commercial and open source software. Developers don't want to have to switch compilers and basic development tools. Administrators can scarcely afford a more-complex support matrix or more time spent getting applications to run effectively in a new environment. Previous generations of Oracle's multicore/multithreaded servers have relied upon individual tools that provide recommendations as to whether an existing application would execute with maximum efficiency on an Oracle multicore/multithreaded server. Developers could employ a variety of individual tools to optimize applications for execution within a Oracle multicore/multithreaded environment. However, Oracle, in its' ongoing efforts to produce a tightly integrated software platform for optimum execution of existing applications or for development of new applications, has integrated much of what was separate tools into Oracle Solaris or Oracle Solaris Studio.

Application Selection

Application selection previously helped identify those applications that stand to benefit from Oracle's multicore/multithreading technology. The Chip Multithreading Selection Tool (`coolst`) still exists, however, is not planned to be updated for the SPARC T3 processor although it will function similarly in the SPARC T3 environment, just as it did for UltraSPARC T2 and T2 Plus processors. The engineering rationale for this is that the SPARC T3 processor both provides a much higher degree of performance and throughput for general purpose codes and effectively quadruples the maximum floating point capability as compared to the UltraSPARC T2 and T2 Plus processors. However, for compatibility purposes and use with previous generations of processors, the tool will be continue to be located on the below-referenced website. To be sure, Oracle's intentions for the SPARC T-Series

processors are for both superior single threaded performance while combining the historically continuing T-Series processors throughput performance.

Development

Developers need to be able to build, test, and evaluate applications, producing the most effective code while advancing their productivity with their chosen tools. Certainly, previous standalone functionality was necessary, yet as Oracle continues to migrate more and more standalone functionality into existing Oracle Solaris and Oracle Solaris Studio, the individual tools have themselves disappeared. However, the functionality continues to exist but as integrated directly into either Oracle Solaris or Oracle Solaris Studio. This is an advantage to developers of applications since developers no longer are required to insure that a given application is processed by multiple tools, as well as ensuring that a given standalone tool is at the latest revision level.

The following tools have been integrated into Oracle Solaris Studio 12.2 release:

- **Oracle Solaris Studio 12.2.** Oracle Solaris Studio 12.2 provides SPARC T3 developers with the latest record-setting, high-performance, optimizing C, C++, and FORTRAN compiler compilers for Oracle Solaris on SPARC and x86/x64 platforms. Oracle Solaris Studio 12.2 also takes advantage of the latest microarchitectural changes implemented in the SPARC T3 processor. Command-line tools and a NetBeans-based integrated development environment (IDE) are provided for application performance analysis and debugging of mixed source language applications. In addition to providing multiplatform support, Oracle Solaris Studio 12.2 compilers are compatible with GCC, Visual C++, C99, OpenMP, and FORTRAN 2003.
- **GCC for SPARC systems (GCC4SS).** Specifically tuned and optimized for SPARC systems, GCC4SS complements the popular GCC compiler suite, delivering up to three times the performance of compiled applications with even greater levels of reliability. At the same time, GCC4SS is 100 percent compatible with GCC, supporting all ABIs, language extensions, and flags
- **Binary Improvement Tool and Simple Performance Optimization Tool.** Used for code coverage analysis, the Binary Improvement Tool (BIT) provides instruction and call count data at runtime, helping to significantly improve developer productivity and application performance. The BIT does not require source code, and works with both executables and libraries. The Simple Performance Optimization Tool (SPOIT) also helps deliver improved developer productivity by automating the gathering and reporting of code data.
- **Oracle Memory Error Discovery Tool (Discover).** Memory access errors can be one of the hardest types of errors to detect, because symptoms of the error typically appear arbitrarily far from the point where the error occurred. The Oracle Memory Error Discovery Tool (Discover) is designed to detect and report common memory access errors. Reported errors include accessing uninitialized memory, writing past the end of an array, or accessing memory after it has been freed.

Tuning and Debugging

Administrators and developers alike need to monitor, analyze, and tune applications under real-world conditions. The following tools aid with tuning and debugging.

Standalone Corestat (version 1.2.4):

- **Corestat.** Corestat provides an online monitoring tool for core use of the SPARC T3 processor, providing a more-accurate measure of processor and system use than tools that only measure the use of individual threads. Implemented as a Perl script and updated for SPARC T3 processor awareness, Corestat aggregates instructions executed by all the threads on a single core, revealing the cycles per instruction of key workloads and indicating where more tuning is needed.
- **Automatic Tuning and Trouble-Shooting System (ATS).** In the interest of automating application tuning, ATS automatically reoptimizes and recompiles binaries with no need for source code. ATS identifies the inadequate optimization and then automatically rebuilds the application with the correct options for optimization. ATS was a plug-in for GCC4SS and previous releases prior to Oracle Solaris Studio 12 . Oracle Solaris Studio 12 now performs this function. It is also currently located at <http://cooltools.sunsource.net/index.html>. This tools on this site will be relocated to <http://opensparc.net/> in December, 2010.

Deployment

Multicore/Multithreading Tool deployment elements provide applications that are already optimized for Multicore/Multithreading technology, and save critical time in configuring systems for performance and consolidation. Deployment elements include:

- **Cool Tuner.** Cool Tuner provides an onsite “virtual” tuning expert, delivering system performance improvements by automatically applying current best practices including both patching and tuning. Depending on administrator experience, Cool Tuner can save hours to weeks of effort tuning servers. This functionality has been integrated into Oracle Solaris Studio 12.2 release.
- **Cool Stack.** Cool Stack represents a collection of the most commonly used free and open source applications, preoptimized for servers based on UltraSPARC T2 and T2 Plus processor technology running Oracle Solaris. Including such popular applications as Apache, Perl, PHP, Squid, Tomcat, and MySQL database software, these applications have been recompiled with previously released Oracle Solaris Studio 12 compilers (pre-SPARC T3) to deliver a 30 to 200 percent performance improvement over standard binaries compiled with GCC. This has been replaced by/renamed Oracle Webstack and relocated to <http://opensolaris.org/projects/webstack>. While compilations do exist for UltraSPARC T2 and T2 Plus (as mentioned above), the source code can easily be re-compiled for the SPARC T3 processor.
- **Consolidation Tool for Oracle's SPARC T-Series Servers.** Powerful Oracle Solaris Containers offer myriad consolidation possibilities and the Consolidation Tool for T-Series Servers speeds their deployment. With a wizard-based GUI, this tool simplifies and automates the installation of consolidated applications, enabling even novice administrators to create a fully virtualized and

consolidated environment using Containers. The result is fast and high-quality consolidated deployments.

Conclusion

Delivering on the demands of IT services applications and virtualized, eco-efficient data centers requires a comprehensive approach that includes innovative processors, system platforms, and operating systems, along with leading application, middleware, and management technology. With its strong technology positions and R&D investments in all of these areas, Oracle is in a unique position to deliver on this vision. Far from futuristic, Oracle has effective solutions today that can help organizations cope with the need for performance and capacity while effectively managing space, power and heat.

Building upon the successful base of UltraSPARC T1 processor technology with additional improvements offered by the UltraSPARC T2 and UltraSPARC T2 Plus processors, the SPARC T3 delivers approximately twice the throughput and efficiency, and serves as the industry's next generation massively threaded System-On-a-Chip. With 128 threads per processor, on-chip memory management, two 10 GbE interfaces, PCIe, and on-chip cryptographic acceleration, the SPARC T3 processor fundamentally redefines the capabilities of a modern processor. By incorporating cache coherency for multiprocessor support, SPARC T3 processors allow these capabilities to be multiplied incrementally. Oracle's SPARC T3-1, T3-2, T3-4, and T3-1B servers leverage these strengths to provide powerful and highly scalable server platforms while delivering even higher levels of performance in a compact rack mount chassis. The result is data center infrastructure that can truly scale to meet new challenges with a very small footprint.

Oracle's SPARC T3-1, T3-2, T3-4, and T3-1B servers provide the computational, networking, and I/O resources needed by the most demanding databases, IT services, enterprise applications and Web services while facilitating highly effective consolidation efforts. With end-to-end support for multithreading and virtualization, these systems can consolidate workloads and effectively use system resources even as they preserve investments in SPARC and Oracle Solaris technology and provide tools for open source software environments. With innovations such as Oracle VM Server for SPARC, Oracle Solaris Containers, and Java technology, organizations can adopt these radical new systems for their most important projects—acting responsibly toward the environment and the bottom line.



Oracle's SPARC T3-1, SPARC T3-2, SPARC T3-4, and SPARC T3-1B Server Architecture
February 2011
Version 1.2

Oracle Corporation
World Headquarters
500 Oracle Parkway
Redwood Shores, CA 94065
U.S.A.

Worldwide Inquiries:
Phone: +1.650.506.7000
Fax: +1.650.506.7200

oracle.com



Oracle is committed to developing practices and products that help protect the environment

Copyright © 2011, Oracle and/or its affiliates. All rights reserved. This document is provided for information purposes only and the contents hereof are subject to change without notice. This document is not warranted to be error-free, nor subject to any other warranties or conditions, whether expressed orally or implied in law, including implied warranties and conditions of merchantability or fitness for a particular purpose. We specifically disclaim any liability with respect to this document and no contractual obligations are formed either directly or indirectly by this document. This document may not be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose, without our prior written permission.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. UNIX is a registered trademark licensed through X/Open Company, Ltd. 1010

Hardware and Software, Engineered to Work Together