



An Oracle White Paper
June 2012

Oracle's SPARC T4-1, SPARC T4-2, SPARC T4-4, and SPARC T4-1B Server Architecture

Introduction	1
SPARC T4 Processor	4
The First Eight-Core Threaded SPARC System-on-a-Chip to Achieve High Single-Threaded Performance.....	4
Taking Oracle's Multicore/Multithreaded Design to the Next Level	5
SPARC T4 Processor Architecture	8
SPARC T4 Processor Cache Architecture	10
SPARC T4 Core Architecture	11
Enterprise-Class Management	20
System Management Technology	20
Oracle Solaris for Multicore Scalability.....	21
Fault Management and Predictive Self Healing	23
Oracle Solaris Cryptographic Frameworks	24
End-to-End Virtualization Technology.....	24
A Multithreaded Hypervisor	24
Oracle VM Server for SPARC.....	25
Oracle Solaris Containers.....	26
Conclusion	26
Appendix: Server Architectures	28
System-Level Architecture.....	28
Chassis Design Innovations	29
Oracle's SPARC T4-1 Server Overview.....	30
Oracle's SPARC T4-2 Server Overview.....	33
Oracle's SPARC T4-4 Server Overview.....	36
Oracle's SPARC T4-1B Server Overview	40

Introduction

Employing Oracle's new SPARC T4 processor, Oracle SPARC T4-1, SPARC T4-2, SPARC T4-4 and SPARC T4-1B servers offer breakthrough performance and energy efficiency to help simplify data center infrastructures and address other demanding challenges. New levels of performance and scalability across a variety of workloads mean that these versatile systems can deliver a virtualized infrastructure for the entire enterprise while also enabling IT managers to deploy and manage fewer types of platforms and fewer numbers of servers.

The SPARC T4 processor takes the industry's first massively threaded System-On-a-Chip (SoC) to the next level by designing an entirely new core from the ground up. Fifth-generation multicore, multithreading technology supports up to 64 threads in as little as two rack units (2RU), providing increased computational density while staying within constrained envelopes for power and cooling. Very high levels of integration help reduce latency, lower costs, and improve security and reliability. The optimized system design provides support for a wide range of IT services and application types. Uniformity of management interfaces and adoption of standards also help reduce administrative costs, while an innovative chassis design shared across Oracle's volume servers provides density, efficiency, and economy for modern data centers.



Figure 1. Oracle's SPARC T4-1, SPARC T4-2, SPARC T4-4, and SPARC T4-1B servers are designed to leverage the considerable resources of the SPARC T4 processor.

Table 1 compares the SPARC T4-1, T4-2, T4-4, and T4-1B servers.

TABLE 1. SPARC T4-1/T4-2/T4-4/T4-1B SERVER FEATURES

FEATURE	SPARC T4-1 SERVER	SPARC T4-2 SERVER	SPARC T4-4 SERVER	SPARC T4-1B BLADE SERVER
CPU	<ul style="list-style-type: none"> 8-core 2.85 GHz SPARC T4 processor 	<ul style="list-style-type: none"> 8-core 2.85 GHz SPARC T4 processor (Dual) 	<ul style="list-style-type: none"> 8-core 3.00 GHz SPARC T4 processor (Dual or Quad) 	<ul style="list-style-type: none"> 8-core 2.85 GHz SPARC T4 processor
Threads	<ul style="list-style-type: none"> Up to 64 	<ul style="list-style-type: none"> Up to 128 	<ul style="list-style-type: none"> Up to 256 	<ul style="list-style-type: none"> Up to 64
Memory Capacity	<ul style="list-style-type: none"> Up to 256 GB (16 GB DDR3 DIMMs) 	<ul style="list-style-type: none"> Up to 512 GB (16 GB DDR3 DIMMs) 	<ul style="list-style-type: none"> Up to 1.024 TB (16 GB DDR3 DIMMs) 	<ul style="list-style-type: none"> Up to 256 GB (16 GB DDR3 DIMMs)
Maximum Internal Disk Drives	<ul style="list-style-type: none"> Up to 8 HDD (2.5-inch SAS-2 300/600 GB disk drives) RAID 0/1, (5/6 + BBWC) 	<ul style="list-style-type: none"> Up to 6 HDD (2.5-inch SAS2-S 300/600 GB disk drives) RAID 0/1 	<ul style="list-style-type: none"> Up to 8 HDD (2.5-inch SAS-2 300/600 GB disk drives) RAID 0/1 	<ul style="list-style-type: none"> Up to 2 HDD (2.5-inch SAS-2 300/600 GB disk drives) RAID 0/1 (via optional RAID Expansion Module)¹
Video	<ul style="list-style-type: none"> One HD-15 VGA port 	<ul style="list-style-type: none"> One HD-15 VGA port 	<ul style="list-style-type: none"> One HD-15 VGA port 	<ul style="list-style-type: none"> One HD-15 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 (Accessed via rKVMS) • Four USB 2.0 ports 	<ul style="list-style-type: none"> • No DVD (Accessed 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 Module² • 2 EM x8 PCIe Gen2 slots
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) • Four 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 10Gb Ethernet ports via optional XAUI pass-through FEM³
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 2060 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 Sun Blade 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 Sun Blade 6000 Modular System Chassis
Operating System	<ul style="list-style-type: none"> • Oracle Solaris 10 8/11, Oracle Solaris 10 9/10 + Oracle Solaris 10 8/11 Patch Bundle, Oracle Solaris 10 10/09 + Oracle Solaris 10 8/11 Patch Bundle⁴ 	<ul style="list-style-type: none"> • Oracle Solaris 10 8/11, Oracle Solaris 10 9/10 + Oracle Solaris 10 8/11 Patch Bundle, Oracle Solaris 10 10/09 + Oracle Solaris 10 8/11 Patch Bundle⁴ 	<ul style="list-style-type: none"> • Oracle Solaris 10 8/11, Oracle Solaris 10 9/10 + Oracle Solaris 10 8/11 Patch , Oracle Solaris 10 10/09 + Oracle Solaris 10 8/11 Patch Bundle⁴ 	<ul style="list-style-type: none"> • Oracle Solaris 10 8/11, Oracle Solaris 10 9/10 + Oracle Solaris 10 8/11 Patch Bundle, Oracle Solaris 10 10/09 + Oracle Solaris 10 8/11 Patch Bundle⁴

¹ A REM is required if selecting 1-4 HDDs

² To connect to optional Network Expansion Modules

³ And appropriate Network Expansion Module

⁴ Solaris 10 8/11 preloaded at factory.

SPARC T4 Processor

The SPARC T4 processor is the industry's most highly integrated system-on-a-chip, supplying the most high-performance threads of any multicore processor available, and integrating all key system functions.

The First Eight-Core Threaded SPARC System-on-a-Chip to Achieve High Single-Threaded Performance

The SPARC T4 processor eliminates the need for expensive custom hardware and software development by integrating computing, security, and I/O onto a single chip. Achieving binary compatibility with earlier SPARC processors, no other processor delivers so much performance in so little space and with such small power requirements. It enables organizations to rapidly scale the delivery of new network services with maximum efficiency and predictability. The SPARC T4 processor is shown in Figure 2.




Figure 2. The SPARC T4 processor allows organizations to rapidly scale the delivery of new network services and compute-intensive workloads with maximum efficiency and predictability.

Table 2 provides a comparison between the SPARC T4, SPARC T3 and UltraSPARC T2 Plus processors.

TABLE 2. SPARC T4, SPARC T3 AND ULTRASPARC T2 PLUS PROCESSOR FEATURE COMPARISON

FEATURE	SPARC T4 PROCESSOR	SPARC T3 PROCESSOR	ULTRASPARC T2 PLUS PROCESSOR
CPU Frequency	• 2.85/3.0 GHz	• 1.65 GHz	• 1.4 GHz
Out-of-Order Exec.	• Yes	• No	• No
Dual Instr. Issue	• Yes	• No	• No
Data/Instr. Prefetch	• Yes	• No	• No
Level 3 Cache	• Yes	• No	• No

Threads/Core	• 8	• 8	• 8
Cores/Processor	• 8	• Up to 16	• Up to 8
Threads/Processor	• 64	• 128	• 64
Hypervisor	• Yes	• Yes	• Yes
Sockets Supported	• 1, 2, or 4	• 1, 2 or 4	• 2 or 4 ⁴
Memory	• Two memory controllers • Up to 16 DDR3 DIMMs	• Two memory controllers • Up to 16 DDR3 DIMMs	• Two memory controllers • Up to 16 or 32 FB-DIMMs
Caches	• 16 KB instruction cache • 16 KB data cache • 128 KB L2 cache • 4 MB L3 cache (8 banks, 16-way, set associative)	• 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)
Technology	• 40 nm technology	• 40 nm technology	• 65 nm technology
Floating Point	• 1 FPU with Mul/Add per core • 8 FPUs per chip	• 1 FPU with Mul/Add per core • 16 FPUs per chip ⁷	• 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, integrated within pipeline • 14 most popular ciphers	• Stream processing unit/core • 12 most popular ciphers	• Stream processing unit/core • 10 most popular ciphers
Additional On-chip Resources	• Dual PCIe Gen2 interface (x8) • Dual 10 GbE XAUI Interfaces (x8) • Coherency logic and links (6 x 9.6 Gb/second)	• Dual 10 GbE Gen2 interfaces (x8) • Dual 10 GbE PCIe XAUI interface (x8) • Coherency logic and links (6 x 9.6 Gb/second)	• PCIe Gen1 interface (x8) • Coherency logic and links (4.8 Gb/second)

* Two-socket implementation represents SPARC T4-2 server, whereas SPARC T4-4 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 the following key goals in mind:

- Radically increase the single-threaded computational capabilities over that of the SPARC T3 processor for workloads that require this level of performance.

- Maintain the computational capabilities to meet the growing demand from Web applications by providing equivalent throughput of the SPARC T3 processor.
- Support larger and more diverse workloads with greater floating-point performance.
- Provide networking performance equivalent to the SPARC T3 CPU to serve network-intensive workloads.
- Provide end-to-end data center encryption with significantly higher performance as well as adding new ciphers implemented within hardware.
- Increase service levels and reduce planned and unplanned downtime.
- Improve 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 SPARC T3 architecture and resulted in a complete redesign of the core itself. For example, simply increasing the number of cores would have gained additional throughput, but would have not have addressed floating-point or single-threaded performance.

The SPARC T4 processor design recognizes that memory latency is truly the bottleneck to improving performance. By redesigning the cores within each processor, designing a new floating-point pipeline, and by further increasing network bandwidth, this processor is able to provide approximately 5X the single-threaded throughput of the SPARC T3 processor.

Each SPARC T4 processor provides eight cores, with each core able to switch between up to eight threads (64 threads per processor) using a modified LRU (Least Recently Used) 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. Unlike SPARC T3, SPARC T4 fetches 1 of 8 threads for instruction propagation through stages of the pipeline to present to the 'Select' stage by the 'Fetch3' stage. Thread instructions are grouped into 2-instruction 'decode groups' and proceed through 'Decode', 'Rename' and 'Pick' stages before proceeding to 'Issue', after which they are sent to one of four subsequent execution pipelines, depending upon the type of instruction to be performed.

Up to this point, each instruction from any thread has proceeded through the pipeline independent of the type of instruction. Two instructions are issued for execution per cycle by the Issue stage per cycle, unlike SPARC T3 which issued only one instruction per cycle

Figure 3 provides a simplified high-level illustration of the thread model supported by an 8-core SPARC T4 processor.

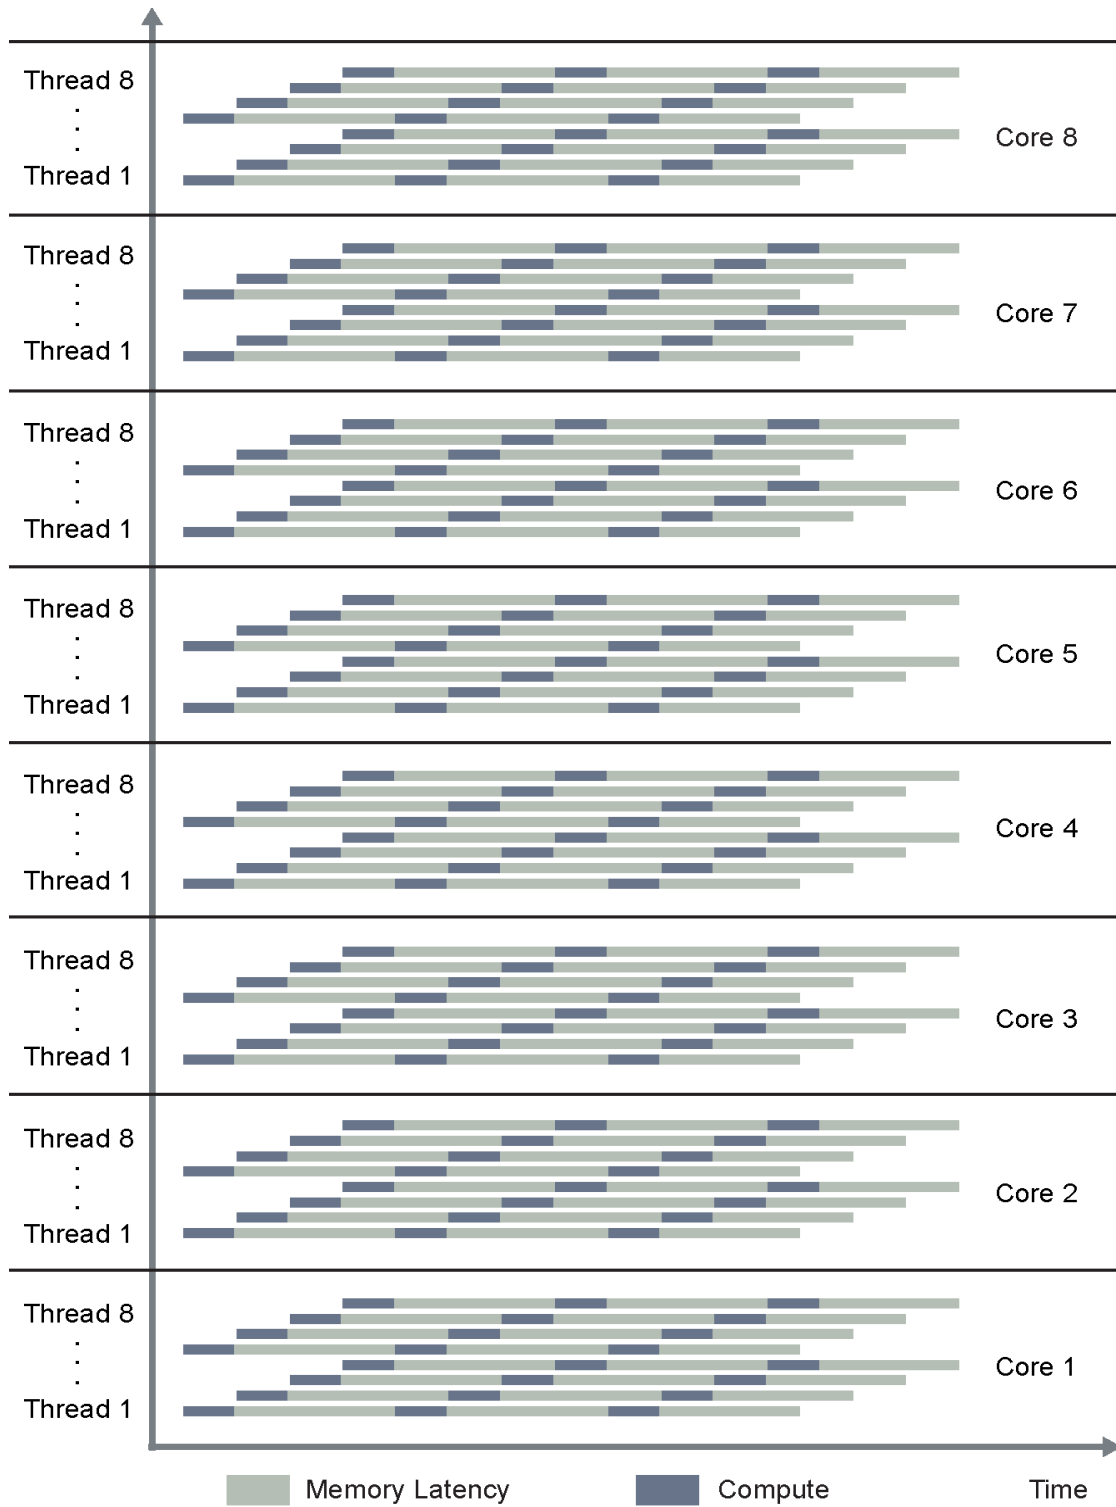


Figure 3. A single 8-core SPARC T4 processor supports up to 64 threads, with up to two threads running in each core simultaneously.

SPARC T4 Processor Architecture

The SPARC T4 processor extends Oracle's multicore/multithreaded initiative with an elegant and robust architecture that delivers real performance to applications. Figure 4 provides a block-level diagram of the SPARC T4 processor.

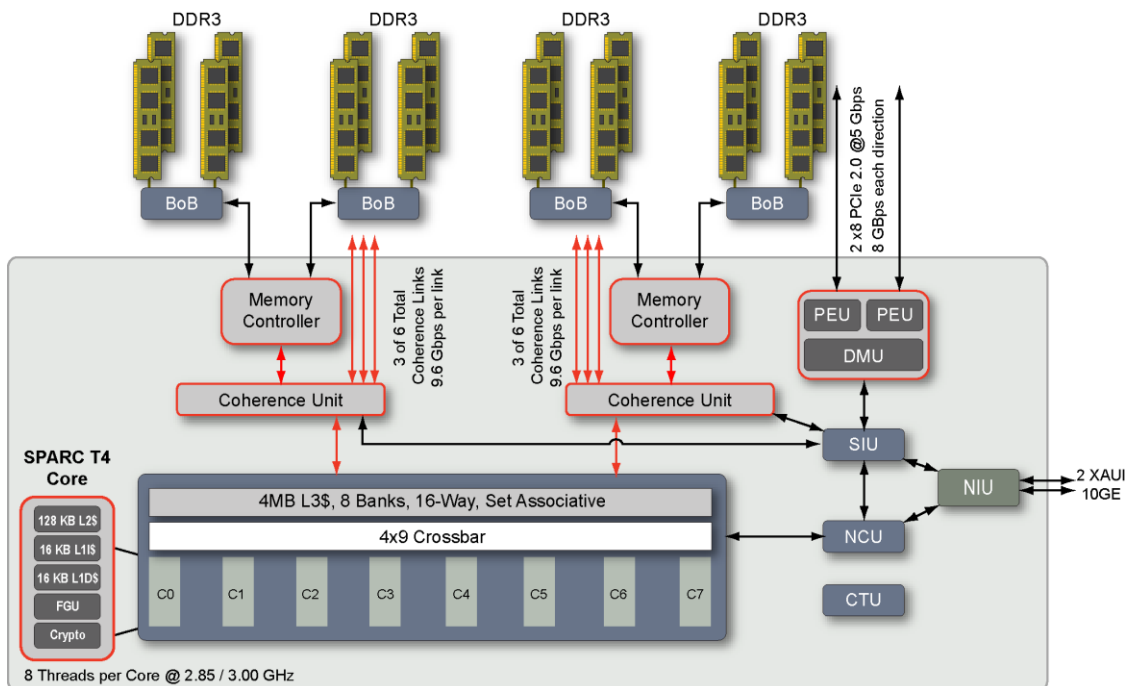


Figure 4. The SPARC T4 processor provides six coherence links to connect to up to four other processors.

The SPARC T4 has coherence link interfaces to allow communication between up to four SPARC T4 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 T4 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 T4 is maintained at 6.4 Gb/sec, identical to that of the SPARC T3.

The SPARC T4 processor can support one-, two- and four-socket implementations. A typical two-socket implementation is shown in Figure 5. Dual-socket, as well quad-socket SPARC T4 implementations interconnect the processors' six coherence links. No additional circuitry is required.

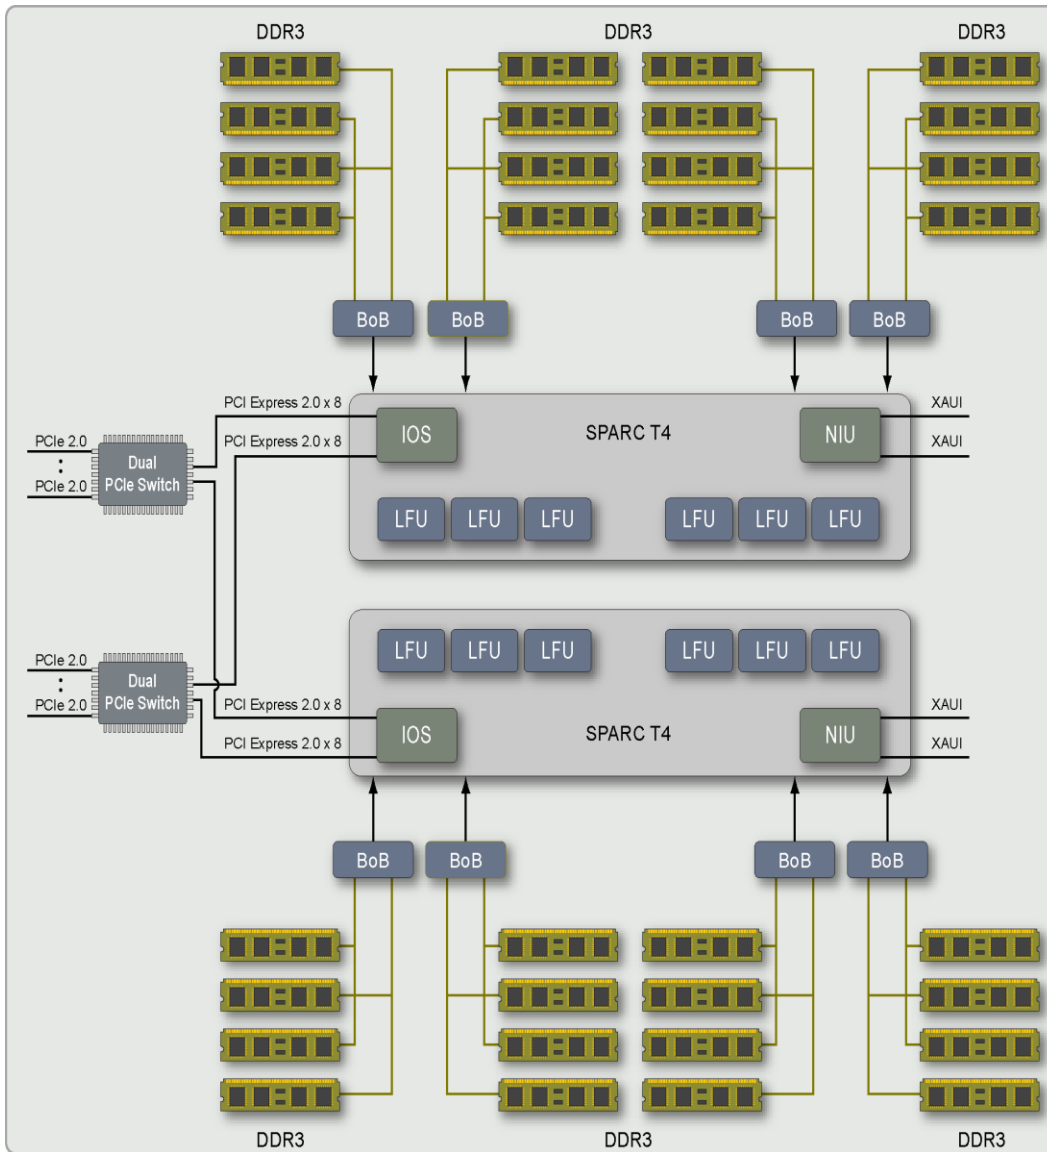


Figure 5. A typical dual-socket SPARC T4 configuration.

SPARC T4 Processor Cache Architecture

The SPARC T4 processor has a three-level cache architecture. Levels 1 and 2 are specific to each core, that is, these two levels of cache are not shared with other cores. Level 3 is shared across all cores of a given processor. Cache sharing does not occur across another processor even though that processor may be in the same physical system. SPARC T4 has Level 1 caches that consist of separate data and instruction caches. Both are 16KB in size and are per core. A single Level 2 cache, again per core, is 128KB in size. The Level 3 cache is *shared* across all eight cores of the SPARC T4 processor and is 4 MB in size, has eight banks, and is 16-way set associative. Figure 6 illustrates the relationship between L2 and L3 caches and shows them connected by a 4x9 crossbar:

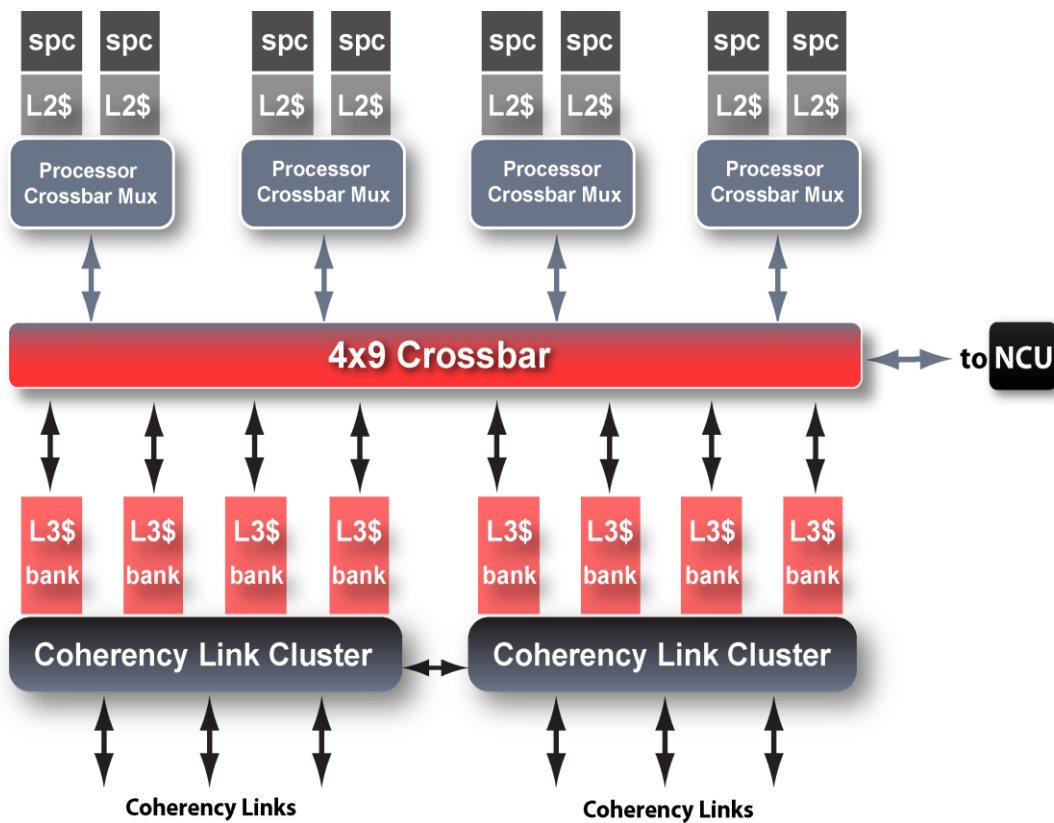


Figure 6. The relationship between Level 2 and Level 3 caches.

SPARC T4 Core Architecture

The SPARC T4 represents a fundamental redesign of the core within a SPARC multicore architecture. Now included within the core are the following aspects that are more conventionally associated with superscalar designs:

- OOO (Out-of-Order) instruction execution
- Sophisticated branch prediction
- Prefetching of both instructions and data
- Much deeper pipelines (relative to previous versions of multicore processors from Sun/Oracle)
- Three levels of cache
- Support for a much larger Memory Management Unit page size (2GB)
- Multiple instruction issue.

All of these characteristics in the SPARC T4 have yielded improvements in single-thread performance by 5X while retaining networking and throughput performance equal to that of previous multicore processors from Sun/Oracle.

There are many functional units, pipelines and associated details that are present within the SPARC T4 core but beyond the scope of this paper. This paper, however (due to the significantly new characteristics and features of the SPARC T4 core), does attempt to touch upon the major exposed (i.e., having visibility to either programmers or users of SPARC T4-based systems) features or characteristics.

One aspect by which the designers of the SPARC T4 architecture were able to achieve a physical space savings of chip real estate was to re-use many physical pieces of a given core for widely varying functionality. For example, for each of the four major pipelines present within each core, the first 14 stages of each pipeline are actually shared. This represents a major space utilization efficiency by making each of the first 14 stages identical. Thus, they can be used by one of two integer instructions, a floating-point graphics instruction, or a load-store instruction. In Figure 7, the first six blocks represent the 14 identical stages, specifically defined in Figure 8.

Dynamic Threading

SPARC T4 is dynamically threaded. While software can activate up to eight strands on each core at a time, hardware dynamically and seamlessly allocates core resources such as instruction, data, and L2 caches and TLBs, as well as out-of-order execution resources such as the 128-entry re-order buffer in the core. These resources are allocated among the active strands. Software activates strands by sending an interrupt to a HALTed strand. Software deactivates strands by executing a HALT instruction on each strand that is to be deactivated. No strand has special hardware characteristics. All strands have identical hardware capabilities.

Since the core dynamically allocates resources among the active strands, there is no explicit "single-thread mode" or "multi-thread mode" for software to activate or deactivate. If software effectively halts all strands except one on a core via Critical Thread Optimization (described later in this document), the core devotes all of its resources to the sole running strand. Thus, that strand will run as quickly as possible. Similarly, if software declares six out of eight strands as non-critical, the two active strands share the core execution resources.

The extent to which strands compete for core resources depends upon their execution characteristics. These characteristics include cache and TLB footprints, inter-instruction dependencies in their execution streams, branch prediction effectiveness, and others. Consider one process that has a small cache footprint and a high correct branch prediction rate such that when running alone on a core, it achieves two instructions per cycle (SPARC T4's peak rate of instruction execution). This is termed a high IPC process. If another process with similar characteristics is activated on a different strand on the same core, each of the strands will likely operate at approximately 1 instruction per cycle. In other words, the single-thread performance of each process has been cut in half. As a rule of thumb, activating N high-IPC strands will result in each strand executing at $1/N$ of its peak rate, assuming each strand is capable of executing close to two instructions per cycle.

Now consider a process that is largely memory-bound. Its native IPC will be small, possibly 0.2. If this process runs on one strand on a core with another clone process running on a different strand, there is a good chance that both strands will suffer no noticeable performance loss, and the core throughput will improve to 0.4 IPC. If a low-IPC process runs on one strand with a high-IPC process running on another strand, it's likely that the IPC of either strand will not be greatly perturbed. The high-IPC strand may suffer a slight performance degradation (as long as the low-IPC strand does not cause a substantial increase in cache or TLB miss rates for the high-IPC strand).

The guidelines above are only general rules-of-thumb. The extent to which one strand affects another strand's performance depends upon many factors. Processes that run fine on their own, but suffer from destructive cache or TLB interference when run with other strands may suffer unacceptable performance losses. Similarly, it is also possible for strands to cooperatively improve performance when run together. This may occur when the strands running on one core share code or data. In this case, one strand may prefetch instructions or data that other strands will use in the near future.

The same discussion can apply between cores running in the chip. Since the L3 cache and memory controllers are shared between the cores, activity on one core can influence the performance of strands on another core.

Figure 7 illustrates a block-level diagram representing a single SPARC core on the SPARC T4 processor. Eight cores are supported per processor.

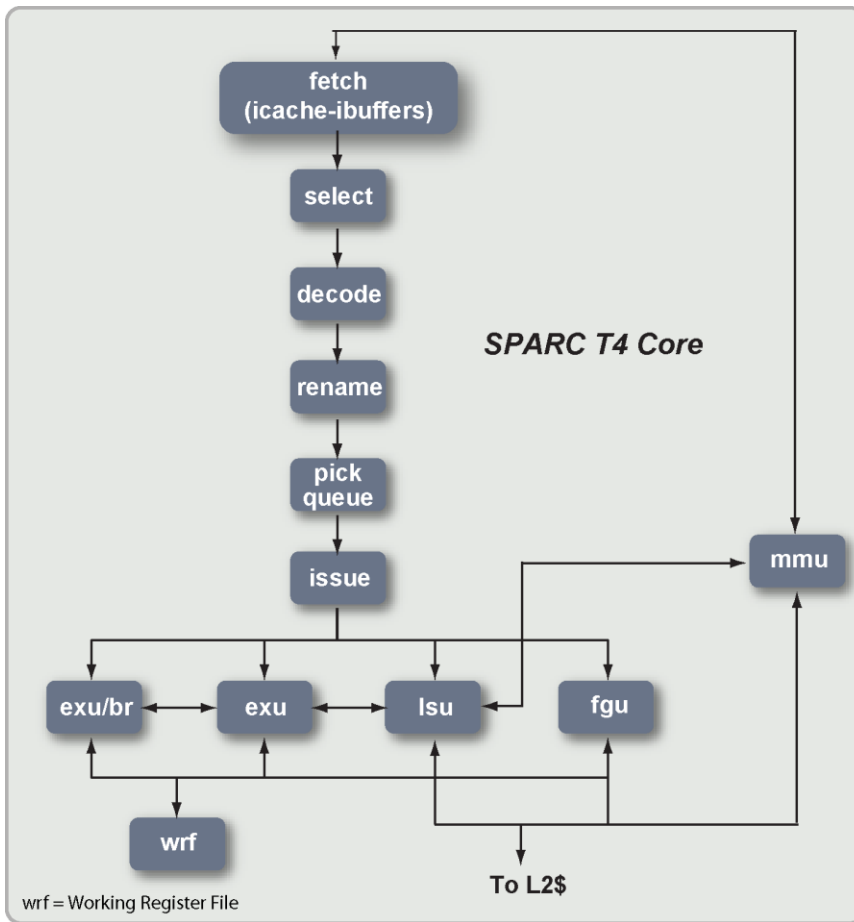


Figure 7. Block-level diagram of a single core of the SPARC T4 processor.

Components implemented in each core include the following:

- **Trap logic unit (not shown).** The trap logic unit (TLU) updates the machine state as well as handling exceptions and interrupts.
- **Instruction fetch unit.** The instruction fetch function is responsible for selecting the thread, fetching instructions from icache for the selected thread, and providing up to four instructions to the Select stage every cycle. On the SPARC T4, it performs the following major functions:
 - Select the thread to be fetched.
 - Fetch instructions from icache for the selected thread, and place them in the Instruction Buffers for the Select Unit.
 - Predict direction and target of delayed control transfer instructions (DCTI) on the thread being fetched.
 - On Icache miss, fetch data from level 2 cache (L2\$), pre-decode it, and store it in icache.

- **Select Unit.** The primary responsibility of the select unit is to schedule a thread for execution on SPARC T4 processor's pipeline for each cycle. For each cycle up to one thread out of eight threads total can be selected for execution. A thread is in one of two states: Ready or Wait. Threads can be in a Wait state due to postsync conditions, mispredicted branches, lack of valid instructions, or other instruction related wait conditions. For each cycle, the select unit selects one thread for execution from among the ready threads using an LRU algorithm for fairness. For the selected thread, up to two instructions are sent to the decode unit per cycle.
- **Decode Unit.** The decode unit on SPARC T4 is responsible for the following:
 - Identifying illegal instructions.
 - Decoding integer and FP sources and sinks for up to two instructions per cycle as well as detecting source/sink dependencies.
 - Generating flat mapping of integer and FP registers.
 - Decoding condition-code sources and destinations.
 - Generating micro-ops for complex instructions.
 - Generating instruction slot assignments.
 - Detecting DCTI (delayed control transfer instruction) couples.
 - Creating NOOPs when exceptions or annulling are detected.
 - Maintaining speculative copies of window registers and executing certain window register instructions.
 - Decoding up to two instructions every cycle
 - Preparing the data and the addressing for the Logical Map Tables (LMTs), which are part of the rename unit (RU).
- **Rename Unit.** The rename unit is responsible for renaming the destinations of instructions and resolving destination-source dependencies between instructions within a thread as well as creating age vector dependency based on issue-slot. Renaming takes three cycles, R1, R2, and R3. For each cycle, the rename unit gets up to two instructions from the decode unit at the end of the D2 cycle. Each group of instructions is called a decode group. The Rename Unit does not break the decode group of instructions received from decode.
- **Pick Unit.** The pick unit schedules up to three instructions per cycle out of a 40 entry pick queue (PQ). Up to three instructions (two instructions plus one store data acquisition op) are written into the PQ during the second phase of the R3. The PQ is read during the first phase of the pick cycle.
- **Issue unit.** The primary responsibility of the issue unit is to provide instruction sources and data to the execution units. The SPARC T4 has six execution units corresponding to the three issue slots as shown in Table 3.

TABLE 3. RELATIONSHIP BETWEEN ISSUE SLOTS AND EXECUTION UNITS

<i>Issue Slot</i>	<i>Unit</i>
0	Load/Store Unit Integer Execution Unit 0
1	Integer Execution Unit 1 Branch Unit FGU SPU
2	Store data operation

- 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. In addition, the integer Fused Mul/Add instruction from the SPARC64 VII instruction set has been added. This also performs part of the cryptographic calculations based upon the algorithm being executed.

A newly designed core for the SPARC T4 implements a sixteen-stage Integer pipeline and a twenty-stage Load-Store pipeline, and a twenty seven-stage Floating-point graphics pipeline. All are present in each of the eight cores of a SPARC T4 processor (Figure 8).

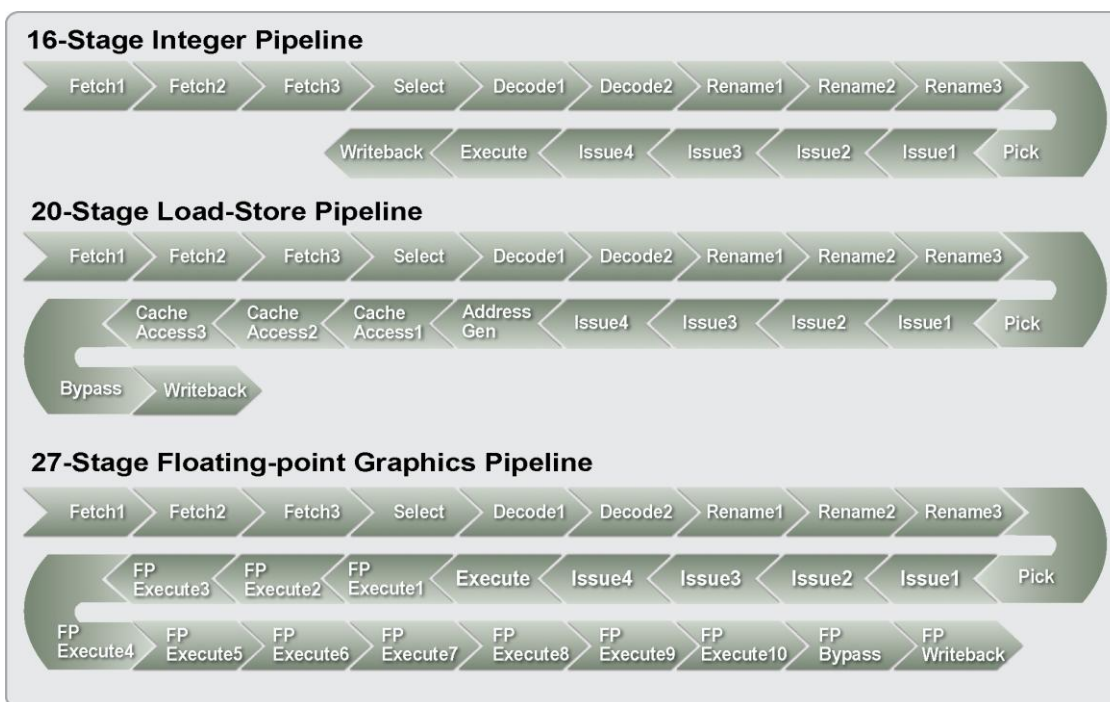


Figure 8. A 16-stage Integer pipeline, a 20-stage Load-Store pipeline, and a 27-stage Floating-point graphics pipeline are provided by each SPARC T4 processor core.

- **Stream Processing Unit.** Each core contains a stream processing unit (SPU) that provides cryptographic processing. This functionality has been implemented within the core pipelines in the SPARC T4 and is accessible by 29 new user-level instructions.
- **Load/Store Unit.** The load-store unit (LSU) is responsible for processing all memory reference instructions and properly ordering all memory references. The LSU will receive load and store instructions out of order as they are picked by the Pick unit. Loads may be issued out of order with respect to other loads and stores may be issued out of order with the respect to other loads and stores. However, loads will not be issued ahead of previous stores. In addition to the memory references required by the instruction set, the LSU also contains a hardware prefetcher, which prefetches data into the L1 cache based upon detected access patterns.
- **Memory management unit.** The memory management unit (MMU) provides a hardware table walk (HWTW) and supports 8 KB, 64 KB, 4 MB, 256 MB and 2 GB pages.
- **Integer execution unit.** The integer execution unit (EXU) is capable of executing up to two instructions per cycle. Single-cycle integer instructions are executed in either the EXU0 (slot0) or EXU1 (slot1) pipeline. Load and store address operations go to EXU0 (slot0). Branch instructions are executed in EXU1 (slot1). Floating point, multi-cycle integer, and SPU instructions go thru the EXU1 (slot1) pipeline. Store data operations go to EXU0 (slot2), but are not considered separate instructions by the EXU since the store address operation must also occur for the same instruction.

To illustrate how the dual integer pipelines function, Figure 9 depicts the dual EXUs with the Working Register Files (WRF), Floating-point Register Files (FRF) and Integer Register Files (IRF) shown along with the various data paths.

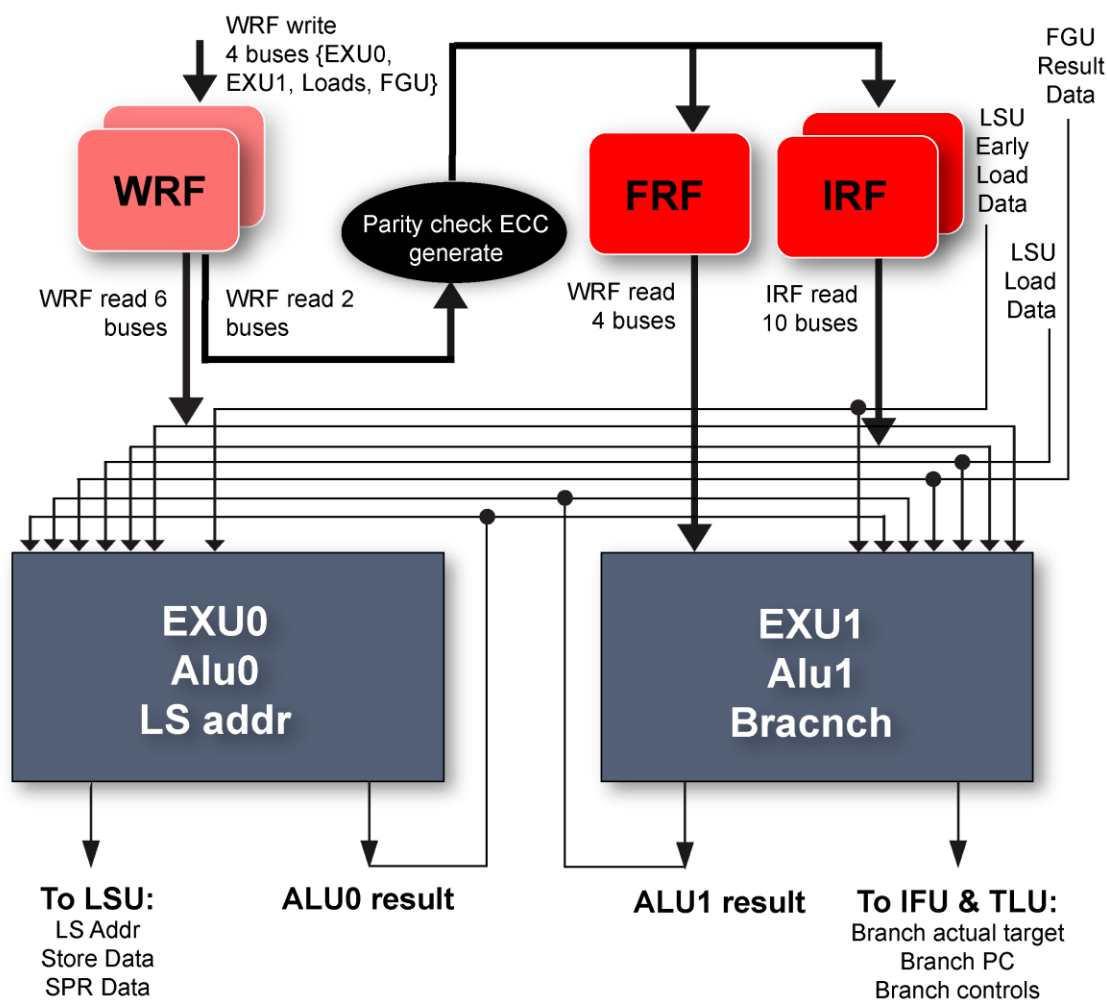


Figure 9. Threads are interleaved between the two integer pipelines and are restricted to EXU0 or EXU1 according to which type of integer operation is to be executed.

Integrated Networking

By providing integrated on-chip networking, both the SPARC T3 and SPARC T4 processors provide integrated on-chip networking, enabling better networking performance than earlier processors such as the UltraSPARC T2 Plus. 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 T4 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 is implemented within the core as part of the pipelines themselves, and operates at the same clock speed as the core. The SPARC T4 supports the following cryptographic algorithms;

- DES/3DES
- AES-128/192/256
- Kasumi, Camellia
- MD5
- SHA-1
- SHA-256
- SHA-512
- RSA via MPMUL/MONTMUL/MONTSQR instructions.

The SPU is designed to achieve wire-speed encryption and decryption on the processor's 10 GbE ports. A cryptographic algorithm (that is supported in hardware from the group previously listed) actually uses parts of the FGU and the Integer pipelines. Figure 10 illustrates the basic logical pipeline of the SPU.

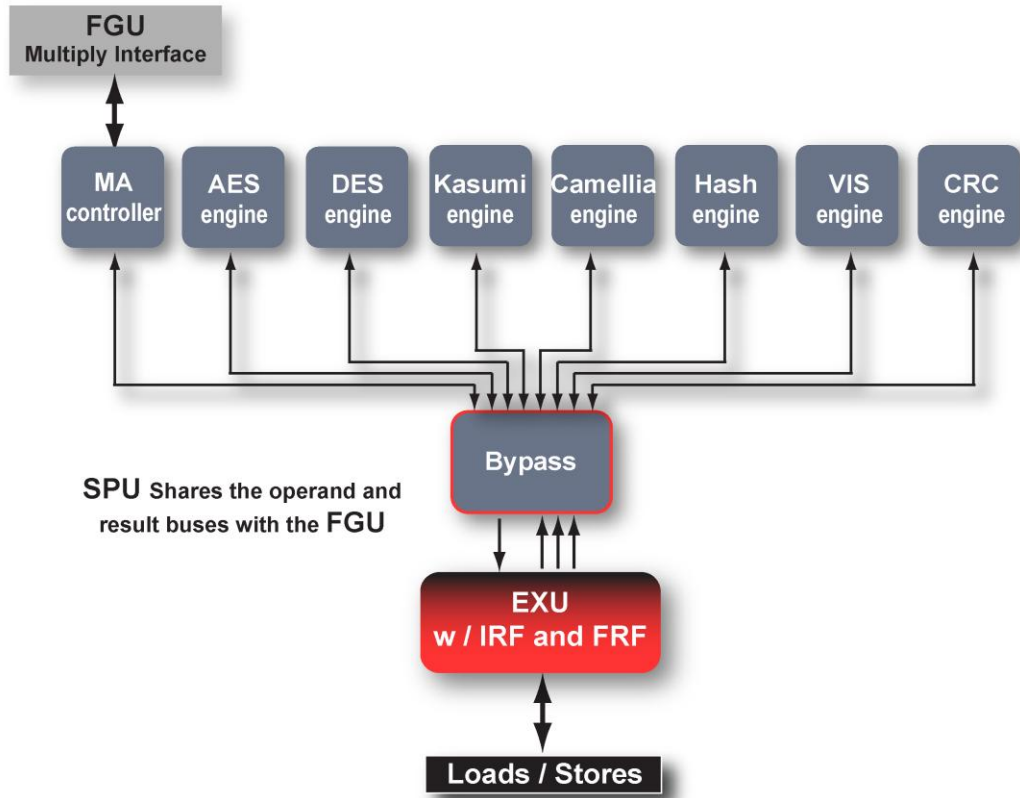


Figure 10. A logical depiction of the Stream Processing Unit pipeline that is in each core.

Integral PCI Express Generation 2 Support

SPARC T4 processors provide dual on-chip PCIe Generation 2 interfaces. Each operates at 5 Gbps 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 an x8 lane) is 4 GB/sec, with a maximum payload size of 256 bytes per PCIe Gen2 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 T4 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

Enterprise-Class Management

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 T4 processor technology, SPARC T4-1, T4-2, T4-4, and T4-1B 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 T4 processor.

System Management Technology

As the number of systems grows 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 all of Oracle's x86 servers, the Integrated Lights Out Manager service processor acts as a system controller, facilitating remote management and administration of SPARC T4-1, T4-2, T4-4, and T4-1B servers. The service processor is fully featured and is similar in implementation to that used in Oracle's other modular and rack mount x86 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 service 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. The service processor also provides 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 send e-mail alerts of hardware failures and warnings, as well as other events related to each server. Its 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 T4-1, T4-2, T4-4, and T4-1B server conditions.

- CPU temperature conditions
- Hard drive presence
- Enclosure thermal conditions
- Fan speed and status
- Power supply status
- Voltage conditions
- Oracle Solaris predictive self healing, 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 also streamlined redeployment of resources as required. Oracle Enterprise Manager Ops Center 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 T4-1, T4-2, T4-4, and T4-1B servers. Oracle Enterprise Manager Ops Center takes a step-by-step approach to unraveling the challenges of getting systems operational quickly. For further information and a detailed description of Oracle Enterprise Manager Ops Center, please refer to <http://www.oracle.com/ru/products/enterprise-manager/044497.html>

Oracle Solaris for Multicore Scalability

Oracle Solaris 10 is specifically designed to deliver the considerable resources of SPARC T4 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-based systems, and compatibility with existing applications is guaranteed. One of the most attractive features of systems based on SPARC T4 processors is that they appear as a familiar SMP system to the Oracle Solaris OS and the applications it supports. In addition, Oracle Solaris 10 and 11 (initial release) have incorporated many features to improve application performance on Oracle's multicore/multithreaded architectures.

- **Accelerated Cryptography.** Accelerated cryptography is supported through the cryptographic framework in Oracle Solaris (initial release) as well as the SPARC T4 processor. The SPARC T4 permits access to cryptographic cypher hardware implementations. For the first time, through user-level instructions, the cyphers are implemented within the appropriate pipeline itself rather than as a co-processor. This means both a more efficient implementation of the hardware-based cyphers as well as no privilege level changes, resulting in large increase in efficiency in cryptographic algorithm calculations. In addition, database operations can make much more efficient use of the various cryptographic cyphers that are implemented within the instruction pipeline itself.
- **Critical Thread Optimization.** Oracle Solaris 10 and 11 have the ability to permit either a user or programmer to allow the Oracle Solaris Scheduler to recognize a 'critical thread' by means of raising its priority to 60 or above through the use of either the Command Line Interface or system calls to a function. If this is done, that thread will run by itself on a single core, garnering all resources of that core for itself. The one condition that would prevent this single thread from executing on a single core is when there are more runnable threads than available CPUs. This limit was put into place to prevent resource starvation to other threads. There will be further enhancements to Critical Thread Optimization done for Oracle Solaris 11 initial release).
- **Multicore/multithreaded awareness.** Oracle Solaris 10 is aware of the SPARC T4 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 T4 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 as required or desired.
- **Support for virtualized networking and I/O.** Oracle Solaris contains technology to support and virtualize components and subsystems on the SPARC T4 processor, including support for the on-chip 10 GbE ports and PCIe interface. As 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.
- **Non-uniform memory access optimization in Oracle Solaris.** With memory managed by each SPARC T4 processor on the SPARC T4-2 and T4-4 servers, these implementations represent a non-uniform memory access (NUMA) architecture. In NUMA architectures, the time needed for a processor to access its own memory is slightly shorter than that required to access memory managed by another processor. Oracle Solaris provides technology that can specifically help to decrease the impact of NUMA on applications and 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 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 Oracle Solaris 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 T4 platforms, protecting software 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.

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.

Predictive self healing and fault management provide the following specific capabilities on the SPARC T4-1, T4-2, T4-4, and T4-1B servers:

- CPU offlining takes offline any cores or threads that have 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.

Oracle Solaris Cryptographic Frameworks

Oracle Solaris cryptographic frameworks at user and kernel level provide applications with a direct interface to several security ciphers. These ciphers can be implemented at software or hardware levels. With Oracle's multicore/multithreaded processors and with these ciphers 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 time of secured applications, but also the load on the system when deployed. This helps developers create secured applications for a large-scale deployment.

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 T4-1, T4-2, T4-4, and T4-1B 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, UltraSPARC T2 Plus and SPARC T3 processors, the SPARC T4 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 architectures.

Corresponding layers of virtualization technology are built on top of the hypervisor, as shown in Figure 11. 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, or LDOMs), Oracle Solaris Zones, and multithreaded applications are able to receive exactly the resources they need.

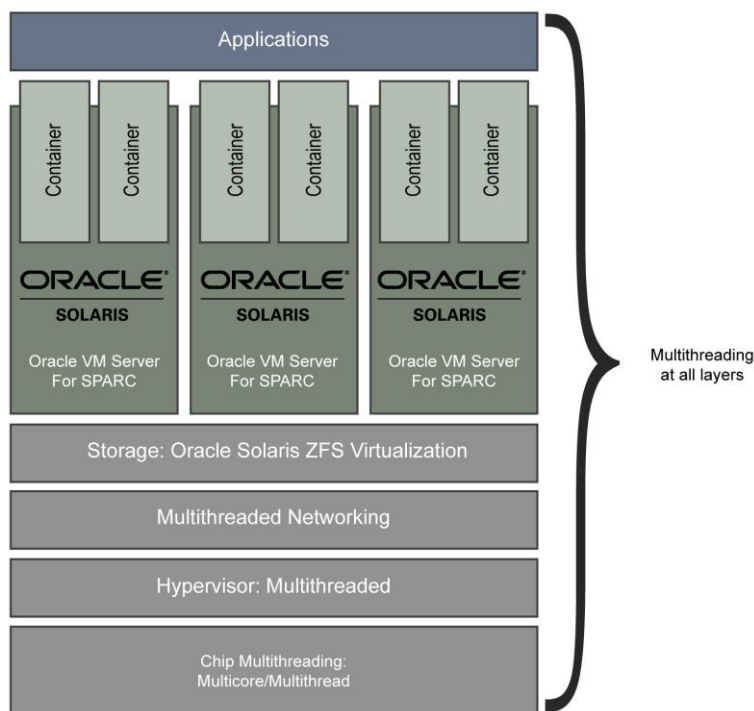


Figure 11. Oracle provides parallelization and virtualization at every level of the technology stack.

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. Each operating system instance contains 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 T4-1 server with a single SPARC T4 processor supports up to 128 domains⁸, and each individual domain can run a unique OS instance.

Oracle VM Server for SPARC 2.1 introduced the ability to perform a live migration from one domain to another. As the name implies, the source domain and application no longer need to be halted or stopped. Migration of a running application from one domain to another is now possible with Oracle

⁸ Though possible, this practice generally is not recommended.

VM Server for SPARC 2.1. Please refer to any of the T4 System Technical presentations for more information or for a more detailed description, refer to [Oracle VM Server for SPARC](#).

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. Domains can also host Oracle Containers to capture the isolation, flexibility, and manageability features of both technologies. Deeply integrating Oracle VM Server for SPARC with the SPARC T4 processor, Oracle Solaris 10 increases flexibility, isolates workload processing, and improves the potential for maximum server utilization.

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. An Oracle Solaris Container is a virtualized operating system environment created within a single instance of Oracle Solaris. 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 Oracle Solaris Container are prevented from interfering with processes running in another Oracle Solaris Container.

CPUs in a multiprocessor system (or threads in the SPARC T4 processor) can be logically partitioned into processor sets and bound to a resource pool, which in turn can be assigned to an Oracle Solaris Container. Resource pools provide the capability to separate workloads so that consumption of CPU resources do not overlap. They 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.

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 technologies. 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, UltraSPARC T2 Plus and SPARC T3 processors, the SPARC T4 serves as the industry's next generation massively threaded System-On-a-Chip. It delivers equivalent network throughput and efficiency compared to SPARC T3 while offering 5X in single-threaded performance gains. With 64 threads per processor, on-chip memory management, two 10 GbE interfaces, two PCIe Generation 2 root complexes, and on-chip cryptographic acceleration, the SPARC T4 processor fundamentally redefines the capabilities of a modern processor. By

incorporating cache coherency for multiprocessor support, SPARC T4 processors allow these capabilities to be multiplied incrementally. Oracle's SPARC T4-1, T4-2, T4-4, and T4-1B 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 T4-1, T4-2, T4-4, and T4-1B 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. 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.

Appendix: Server Architectures

Oracle's SPARC T4-1, T4-2, T4-4 and T4-1B 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 T4 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 T4-1, T4-2, T4-4, and T4-1B Servers

A separate motherboard design is used in each of Oracle's T4-1, T4-2, T4-4 and T4-1B servers. Although each system deploys the same technologies (SPARC T4 processor, DDR3 Memory, PCIe Generation 2 switches, SAS-2), each motherboard has its own unique design and form factor. Common features of the SPARC T4 motherboard (Processor Module for the SPARC T4-4 server) are:

- A minimum of one socket for a SPARC T4 processor
- Memory slots to supply memory for the SPARC T4 server
- A remote to local memory latency ratio of 1.47 (in the case of a 2 or 4 processor system running the `1mbench` read access test)

Memory Subsystem

In Oracle's SPARC T4-1, T4-2, T4-4, and T4-1B servers, the SPARC T4 processor provides on-chip memory controllers that communicate indirectly to DDR3 DIMMs via newly designed Buffers-on-Board (BoB) memory interfaces through four high-speed serial links. Two dual-channel memory controller units (MCUs) are provided on the SPARC T4 processor. Each MCU can transfer data at an aggregate rate of 6.4 Gb/sec. There are 16 motherboard memory socket locations on Oracle's SPARC T4-1 server motherboard. The SPARC T4-2 motherboard has 32 DIMM slots broken up over four memory risers, each of which plug into the motherboard. There are up to 64 motherboard memory socket locations on the SPARC T4-4 motherboard and 16 motherboard memory socket locations on the SPARC T4-1B motherboard, providing sufficient board space for four 1066 MHz DDR3 DIMMs per channel.

I/O Subsystem

Each SPARC T4 processor interfaces through two (x8) PCIe Gen2 ports capable of operating at 5 GB/s bi-directionally. In each server, these ports natively interface through the I/O devices through PCIe Gen2 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 T4-1 and T4-2 servers. (DVDs are not supported in Oracle's SPARC T4-4 or T4-1B servers.)
- **Modular disk backplanes.** Depending on the system, a 6- or 8-disk backplane is attached to the LSI disk controller by one or more x4 SAS-2 links. Oracle's SPARC T4-1 server supports a SAS-2 capable 8-disk backplane. The SPARC T4-2 server supports a SAS-2 capable 6-disk backplane, and the SPARC T4-4 server supports a SAS-2 capable 8-disk backplane.
- **Gigabit Ethernet.** Oracle's SPARC T4-1, T4-2, and T4-4 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 T4-1B blade server.
- **Dual 10 Gigabit Ethernet.** Oracle's SPARC T4 processor provides dual 10 Gigabit Ethernet Attachment Unit Interfaces (XAUI) ports. [Note that the T4-1B server does *not* make use of the SPARC T4 processor-based 10 GbE XAUI ports.]
- **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 T4-1, T4-2, T4-4, and T4-1B servers share basic chassis design elements. This approach not only provides a consistent look and feel across the product line, but it also 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 T4-1, T4-2, T4-4 and T4-1B 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-aisle/cold-aisle multi-racked 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 to fan modules enables fans to 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 T4-1, T4-2, and T4-4 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 T4-1, T4-2, and T4-4 servers all support USB 2.0 interfaces, two each on the front and back of the chassis. Oracle SPARC T4-1, T4-2, T4-4, and T4-1B servers also have an internal USB port. Oracle's SPARC T4-1B servers support two external USB 2.0 interfaces on the front of the blade server via a dongle cable.

Oracle's SPARC T4-1 Server Overview

The compact Oracle SPARC T4-1 server provides significant computational power in a space-efficient, low-power 2RU rack mount package. With excellent price to performance ratios, 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. 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 T4-1 server has a unique motherboard design (Figure 12).

- Up to 256 GB of memory in 16 DDR3 DIMM slots (4GB, 8GB and 16GB DDR3 DIMMs supported)
- Four onboard 10/100/1000 Mb/sec Ethernet ports
- Dedicated low-profile PCIe slots (x8) w/two combination XAUI or low-profile PCIe x4 slots
- Five USB 2.0 ports (two forward, two rear facing, 1 internal)
- Eight disk drive slots support SAS-2 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 13 illustrates the front and rear panels of Oracle's SPARC T4-1 server, showing 8-disk backplane.




Figure 13. The front and rear panels of Oracle's SPARC T4-1 server.

External features of Oracle's SPARC T4-1 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 hot-plug SAS-2 disk drives that are insertable through the front panel
- One Slimline, slot-accessible DVD+R/-W 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 HD-15 VGA video port
- A total of six PCIe card slots, two of which can alternately support XAUI cards connected to the SPARC T4 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 T4-2 Server Overview

The expandable SPARC T4-2 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 T4-2 server is also an ideal platform for consolidated Tier 1 and Tier 2 workloads.

Oracle's T4-2 server has a distinct motherboard design (Figure 14).

Oracle's SPARC T4-2 server includes the following major components:

- Dual SPARC T4 processors with eight cores per processor operating at 2.85 GHz
- Up to 512GB of memory in 32 DDR3 DIMM slots (4GB, 8 GB and 16 GB DDR3 DIMMs)
- Four onboard 10/100/1000 Mb/sec Ethernet ports
- Ten dedicated low-profile PCIe x8 slots
- Slots 3 and 6 accept x16 PCIe cards, but only run at x8 speeds
- Slots 4 & 5 run at x4 speeds, but accept x8 cards
- 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)
- Six available disk drives slots that support SAS-2 commodity disk drives
- ILOM 3.0 system controller
- Two (N+1) hot-plug/hot-swap high-efficiency 2060 watt AC power supplies
- Six fan assemblies under environmental monitoring and control, N+1 redundancy

Front and Rear Perspectives

Figure 15 illustrates the front and back panels of Oracle's SPARC T4-2 server.




Figure 15. The front and back panels of Oracle's SPARC T4-2 server.

External features of Oracle's SPARC T4-2 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
- Six hot-plug SAS-2 disk drives insertable through the front panel of the system
- One Slimline DVD +R/□ W 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 HD-15 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 T4-4 Server Overview

With support for up to four SPARC T4 processors and up to 256 threads, the compact Oracle SPARC T4-4 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. It presents many opportunities as a consolidation and virtualization server 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 Oracle SPARC T3-4 system. Depending on the model selected, the SPARC T4-4 server features dual- or quad-SPARC T4 processors. Uniquely setting this processor apart from past multicore processors is its' ability to deliver a 5X performance increase on single-threaded workloads.

Oracle's T4-4 server has a unique motherboard design (Figure 16).

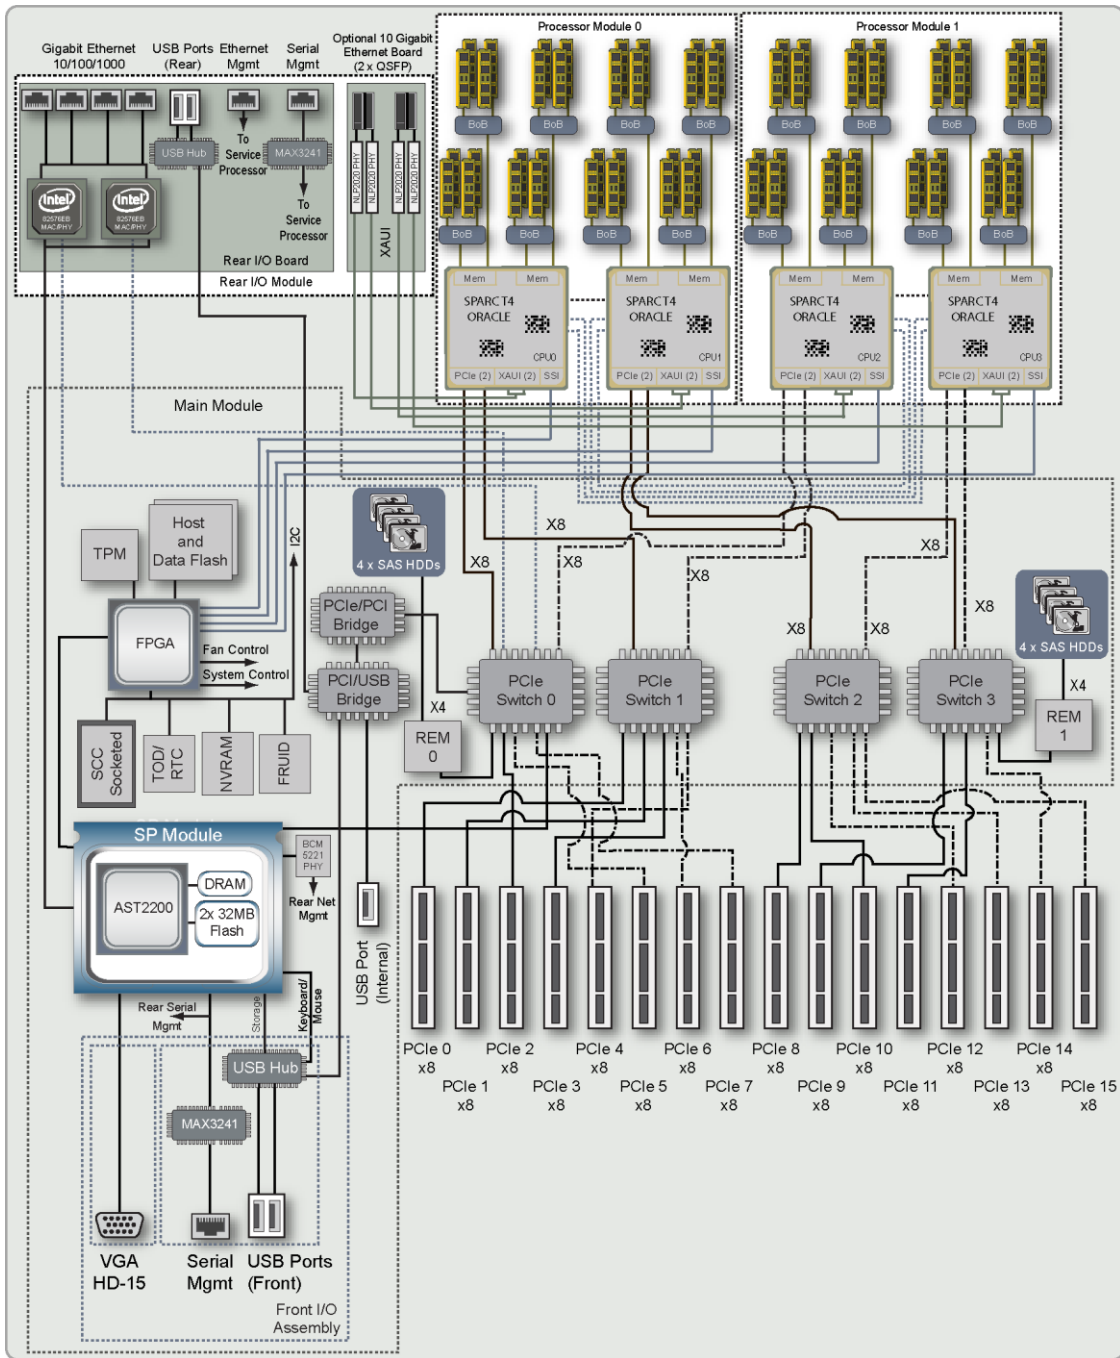


Figure 16. Oracle's SPARC T4-4 server motherboard design

Enclosure

The 5RU Oracle SPARC T4-4 server enclosure is designed for use in a standard 19-inch rack (Table 6).

TABLE 6. DIMENSIONS AND WEIGHT OF ORACLE'S SPARC T4-4 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 T4-4 server includes the following major components:

- Two or four SPARC T4 processors eight cores per processor operating at clock speed of 3.00 GHz
- Up to 1.024 TB of memory in 64 DDR3 DIMM slots (4 GB, 8GB and 16GB DDR3 DIMMs are supported)
- Four onboard 10/100/1000 Mb/sec Ethernet ports
- Sixteen x8 PCIe Generation 2 slots (via EMs)
- Eight XAUI ports via 2 QFSP quad connectors
- Four USB 2.0 ports (two forward, two rear facing)
- Eight available disk drives slots that support SAS-2 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 17 illustrates the front and rear panels of Oracle's SPARC T4-4 server.




Figure 17. The front and rear panels of Oracle's SPARC T4-4 server

External features of Oracle's SPARC T4-4 server include the following.

- Front
 - Serviced front/rear only (no sliding rack rails)
 - Two Processor Modules, each with 2x SPARC T4 CPUs, with 16 or 32 DDR3 DIMMs: 4GB, 8GB or 16GB
 - Eight HDDs
 - Four power supplies (2+2)
 - A HD-15 VGA video port
 - Two 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
 - Five Fans (N+1)
 - Four (2+2) AC cords (200-240V) (any 2 required)
 - Two USB ports
 - Console serial port (duplicate of front)
 - Console 10/100 network port
 - Two QSFP ports (10GbE XAUI network)
 - Four 1GbE network ports
 - LEDs and indicators

Oracle's SPARC T4-1B Server Overview

The Oracle SPARC T4-1B server is optimized to deliver streaming media, virtualization and consolidation, other back office applications such as Java application servers, OLTP databases, ERP, CRM, SOA, and business integration. With support for a SPARC T4 processor in a blade server, it is ideal for expansion capabilities and integrated virtualization technologies. Oracle's SPARC T4-1B server is also an ideal platform for consolidating Tier 1 and Tier 2 workloads.

Oracle's T4-1B server has a unique motherboard design (Figure 18).

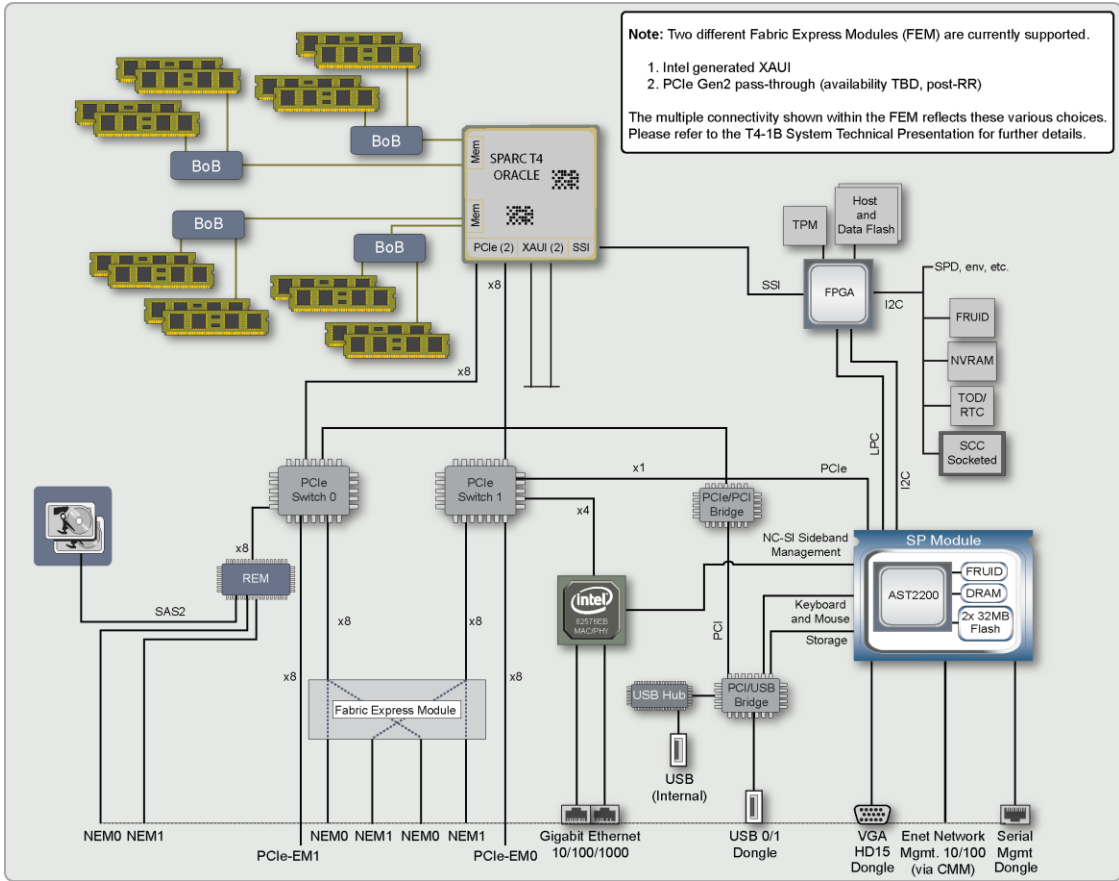


Figure 18. Oracle's SPARC T4-1B server motherboard design

Enclosure

Oracle's SPARC T4-1B server features a compact blade server, giving organizations the flexibility to scale their processing and I/O by simply adding the SPARC T4-1B to an existing Oracle Sun Blade Chassis 6000 (Table 7).

TABLE 7. DIMENSIONS AND WEIGHT OF ORACLE'S SPARC T4-1B 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 2 disks, full memory)	20 pounds	9.1 kilograms

Oracle's SPARC T4-1B blade server includes the following major components:

- One SPARC T4 processor with eight cores operating at 2.85 GHz
- Up to 256 GB of memory in 16 DDR3 DIMM slots (4GB, 8GB and 16GB DDR3 DIMMs supported)
- Two onboard 10/100/1000 Mb/sec Ethernet ports
- Two dedicated x8 PCIe Express Module slots
- Two x8 PCIe slots for use by optional Fabric Expansion Modules (use with appropriate Network Expansion Modules)
- Three USB 2.0 ports (two external via dongle, 1 internal restricted thumb drive)
- Up to two available disk drive slots supporting commodity SAS-2 disk drives
- ILOM 3.0 system controller

Front and Rear Perspectives

Figure 19 illustrates the front and back panels of Oracle's SPARC T4-1B server.




Figure 19. The front and back panels of the SPARC T4-1B server.

External features of the SPARC T4-1B server include the following.

- Front component status indicator lights that provide locator (white), service required (amber), and activity status (green) for the system
- Two Hot-plug SAS-2 disk drives accessible through the front panel of the system
- Two USB 2.0 ports accessible via dongle
- A HD-15 VGA video port
- Two management ports for use with the ILOM 3.0 system controller:
 - RJ-45 serial management port (attached via dongle) provides default connection to the ILOM 3.0 controller.
 - Network management port, connected to the SB6000 Chassis Management Module (CMM) switch and accessed via the network management port on the CMM.



Oracle's SPARC T4-1, SPARC T4-2, SPARC T4-4, and SPARC T4-1B Server Architecture
June 2012
Version 1.1

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