An Oracle White Paper
April 2012

Sun Blade 6000 Server Module Architecture

ORACLE

Introduction

While many blade systems promise dense packaging to support consolidated infrastructures, most are limited in flexibility and cannot be expanded to handle the scalability requirements of today's high-end Web services and enterprise IT applications. Sun Blade 6000 modular systems are designed with an open architecture and the latest technology so that they can provide high performance, scalable capacity, and massive levels of I/O for the most demanding application environments.

Flexible and Expandable Sun Blade 6000 Modular Systems

Sun Blade 6000 modular systems support flexible and expandable configurations, built from a range of standard hot-plug, hot-swap modules, including:

- SPARC T4-1B and T3-1B server modules, Sun Blade X6270 M2, X6270 M3, and X6275 M2 server modules, or Sun Blade Storage Module M2 in any combination up to 10 modules per chassis
- Blade-dedicated PCI Express (PCIe) ExpressModules (EM), which support industry-standard PCIe interfaces and provide each blade with its own unique I/O configuration (just like rack servers)
- NEMs, providing access and an aggregated "chassis-wide" common interface to all of the server modules in the Sun Blade 6000 chassis
- Sun Blade Storage Module M2 that expands high-performance, chassis-integrated storage capacity beyond individual server module capacities
- Integral chassis monitoring module (CMM) for transparent management access to server and storage modules
- Hot-swap (N+N) power supply modules and hot-swap redundant (N+1) fan modules

With common system components and a wide choice of compute and I/O modules, organizations can scale capacity of both processing power and I/O throughput with fine or coarse granularity. Sun Blade 6000 modular systems also give organizations a choice of operating systems and virtualization environments. They are optimized to run Oracle Linux, Oracle Solaris and Oracle VM, and are also certified to run Red Hat Enterprise Linux, SUSE Linux Enterprise Server, Windows Server and VMware.

Focus of This Paper

This paper provides an in-depth review of the architecture of Sun Blade 6000 server modules, describing how Oracle's Sun Blade 6000 server modules are designed to protect investments through industry-leading scalability and an open architecture that leverages industry standard components.

This paper is part of a series of white papers about the architecture of Sun Blade 6000 modular systems. The other two white papers in the series include:

- *Oracle's Sun Blade 6000 Modular Systems* — An overview of the Sun Blade 6000 modular systems architecture.

- *Sun Blade 6000 I/O and Management Architecture* — An in-depth review of the I/O, storage, and management components of the Sun Blade 6000 modular systems architecture.

Sun Blade 6000 Server Modules Overview

The ability to host demanding compute, memory, and I/O-intensive applications is ultimately dependent on the characteristics of the actual server modules. The innovative Sun Blade 6000 chassis provides customers with considerable flexibility in terms of delivering powerful server modules for a broad range of applications.

Except for labeling, all server modules feature a physically similar front panel design. This design is intentional since any server module can be used in any slot of the chassis, no matter what the internal architecture of the server module. As mentioned, all server modules use the same midplane connectors and have similar I/O capabilities.

Up to 10 Sun Blade 6000 family server modules can be inserted vertically beneath the power supply modules in the front of the chassis. Depending on the server module, up to four slots for storage (HDDs or SSDs) media are provided for easy hot-swap access from the front of the server module. Power efficient solid-state drives (SSDs) based on innovative enterprise Flash technology can be used instead of SAS or SATA hard disk drives (HDDs). Indicator LEDs and a high density I/O port are also provided on the front of the server modules for easy access. A number of connectors are provided through the high-density front panel communications port of each server module. These ports are broken out using an available server module adaptor multiport "dongle" cable. Depending on the server module, available ports include a VGA HD-15 monitor port, two USB 2.0 ports, and a DB-9 or RJ-45 serial port that connects to the server module and integral service processors.

A Choice of Processors

By providing a choice of Oracle SPARC or Intel Xeon processors, the Sun Blade 6000 modular systems can serve a wide range of applications and workloads. Organizations are free to choose the platform that best satisfies their requirements. Server modules of different architectures (SPARC and x86) can co-reside within the same Sun Blade 6000 chassis.

Supported Server Modules and Their Capacities

Table 1 lists the server modules supported in Sun Blade 6000 modular systems and identifies their capacity in terms of number of processors, cores, threads, and maximum memory.

TABLE 1. CAPACITIES FOR SUN BLADE 6000 SERVER MODULES

SERVER MODULE	PROCESSOR(S)	MAXIMUM CORES/THREADS	MEMORY SLOTS
SPARC T4-1B Server Module	1 SPARC T4 processor	• 8 cores 64 threads	16 DDR3 DIMM slots
SPARC T3-1B Server Module	1 SPARC T3 processor	• 8 cores 64 threads, or 16 cores 128 threads	16 DDR3 DIMM slots

Sun Blade X6270 M2 Server Module	Up to 2 Intel Xeon Processor 5600 Series CPUs	• Quad Core: 8 cores, 16 threads • Six Core: 12 cores, 24 threads	18 DDR3 DIMM slots – 1333 MT/s support
Sun Blade X6270 M3 Server Module	2 Intel Xeon Processor E5-2600 product family CPUs	• Quad Core: 8 cores, 16 threads • Six Core: 12 cores, 24 threads • Eight Core: 16 cores, 32 threads	24 DDR3 DIMM slots – 1600 MT/s support
Sun Blade X6275 M2 Server Module	2 separate nodes, each with 2 Intel Xeon Processor 5600 Series CPUs	• Per node, 5600 Series: 12 cores, 24 threads • Per server module, 5600 Series: 24 cores, 48 threads	12 DDR3 DIMM slots per compute node 24 DDR3 DIMM slots per server module

Enterprise-Class Features

Unlike most traditional blade servers, Sun Blade server modules provide a host of enterprise features that help ensure greater reliability and availability:

- Each server module supports hot-plug capabilities
- Disk Capacity varies between server modules, but when provided (two- or four-disk models), disks are hot-pluggable and provide the option of RAID controllers[1]
- Redundant hot-swap chassis-integrated fans mean greater reliability for the server modules due to a reduction in parts
- Redundant hot-swap chassis-integrated power supply modules mean that no power supplies are located in individual server modules, thereby increasing the reliability of the server modules
- Leveraging chassis-based infrastructure components (power supplies, fans etc.) means there are fewer moving parts on the individual server modules which greatly helps increase reliability and reduce power/cooling requirements

A Choice of Operating Systems and Virtualization Support

In order to provide maximum flexibility and investment protection, the server modules support a choice of operating systems, including:

- Oracle Solaris
- Multiple versions of the Linux operating system, including Oracle Linux, 64-bit Red Hat, SUSE
- Microsoft Windows Server

[1] RAID expansion modules (REMs) are supported by the T3-1B and X6270 M2 server modules.

Virtualization Software:

- Oracle VM
- VMware

Oracle Solaris Support on all Server Modules

Among the available operating systems, Oracle Solaris is ideal for large-scale enterprise deployments. Supported on all the Sun Blade server modules, Oracle Solaris has certain hardware-specific features that can enhance performance, efficiency, and reliability — with different features mapped to the CPU architecture as noted.

Additional information about the Oracle Solaris operating system are provided in the architecture overview white paper, "Oracle's Sun Blade 6000 Modular Systems," and on the Web at: *http://www.oracle.com/technetwork/server-storage/solaris/overview/*.

Oracle VM Server for SPARC Processor Support in SPARC T-Series Server Modules

Supported in all Sun SPARC servers that utilize Sun processors with chip multithreading (CMT) technology, Oracle VM Server for SPARC (formerly Sun Logical Domains) provides a full virtualization layer. Multiple independent virtual machines, each with their own operating system instance, enable virtualized CPU, memory, storage, console, and cryptographic devices. Within the Oracle VM Server for SPARC architecture, a small firmware layer known as the Hypervisor provides a stable, virtualized machine architecture to which an operating system can be written. As such, each logical domain is completely isolated. The maximum number of virtual machines created on a single platform relies upon the capabilities of the underlying hardware architecture and the Hypervisor as opposed to the number of physical hardware devices installed in the system. For example, the SPARC T4-1B supports up to 64 logical domains, while the SPARC T3-1B Server supports up to 128 logical domains, and each individual logical domain can run a unique instance of the operating system[2].

By taking advantage of Oracle VM Server for SPARC, organizations gain the flexibility to deploy multiple operating systems simultaneously on a single server module. In addition, administrators can exploit virtual device capabilities to transport an entire software stack hosted on a logical domain from one physical machine to another. Logical domains can also host Oracle Solaris Zones (formerly Oracle Solaris Containers) to capture the isolation, flexibility, extensive granularity and manageability features of both technologies. By deeply integrating logical domains with both the industry-leading CMT capabilities of the SPARC T4 and SPARC T3 processors and the Oracle Solaris OS, Oracle VM Server for SPARC technology increases flexibility, isolates workload processing, and improves the potential for maximum server utilization.

[2] Though technically possible, this practice is not generally recommended.

Server Module Architectures

The T4-1B, T3-1B, X6270 M2, and X6275 M2 server modules are described in the subsections that follow. Both the Sun Blade X6270 M2 and X6275 M2 server modules employ Intel Xeon Processor 5600 Series CPUs whereas the SPARC T4-1B and SPARC T3-1B server modules utilize SPARC T3 and SPARC T4 processors.

The server module descriptions that follow contain some information about I/O processing and storage capabilities that are specific to the server modules. However, they do not cover I/O and storage, and networking technologies that can be deployed in a top-of-rack (ToR) configuration to support all server modules within a Sun Blade 6000 rack. For additional information about I/O and management components related to these blade servers, refer to the white paper, "Sun Blade 6000 I/O and Management Architecture," which describes the PCIe ExpressModules (EMs), NEMs, Sun Blade Storage Module M2, and the chassis monitoring module (CMM).

SPARC T4-1B Server Module

The SPARC T4-1B server module is a single socket 2.85GHz SPARC T4 processor-based blade module with eight cores and 64 threads. Unique and new to the SPARC T4 processor, The SPARC T4 processor offers superior performance for a broad range of workloads from single-treaded workloads such as database, OLTP, and batch, to multithreaded workloads such as middleware applications. The SPARC T4 processor delivers 5 x the single-thread performance of the previous generation SPARC T3 processor, while maintaining industry leading multi-thread performance.

The physical layout out of the SPARC T4-1B server module is shown in Figure 1.

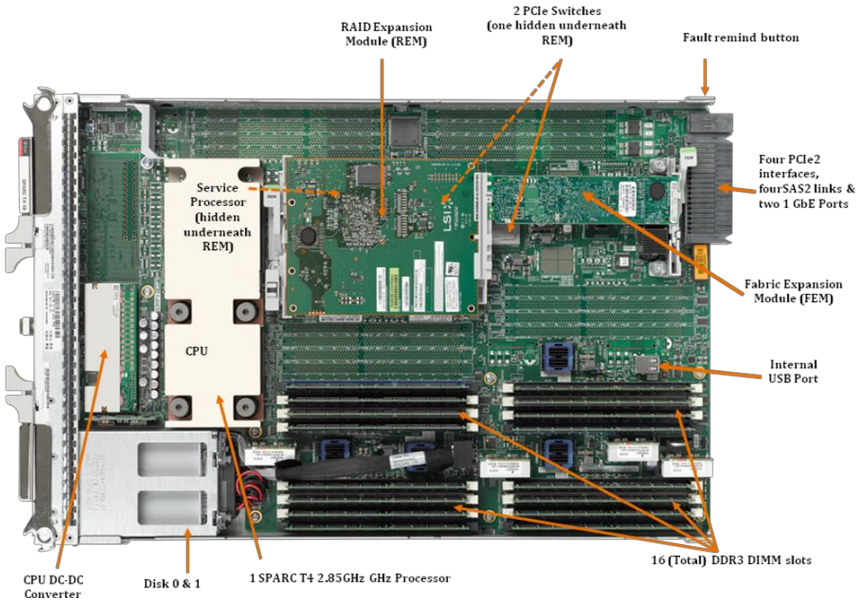

Figure 1. The SPARC T4-1B server module supports a single SPARC T4 processor.

The SPARC T4-1B blade module is based on the latest generation SPARC T4 2.85GHz processor with eight cores, and 64 simultaneous threads. With integrated 10GbE networking, built-in PCIe Generation 2 expansion, and integrated on-chip cryptographic acceleration, the SPARC T4-1B offers enterprise performance and security delivered in a compact, economical blade form factor.

The SPARC T4-1B includes 16 DIMM slots for up to 256 GB memory and four drive slots for hot-pluggable 2.5 inch SAS or SATA drives. It supports one optional Fabric Expansion Module (FEM), one RAID Expansion Module (REM), and the ORACLE ILOMORACLE ILCOM service processor as shown in Figure 1.

The Sun Blade 6000 chassis supports up to 10 SPARC T4-1B server modules for a total of 80 cores and 640 threads, and I/O throughput up to 6.4 Terabit-per-second.

The SPARC T4 Processor

- One T4 processor in SPARC T4-1B, eight cores, 64 threads, 2.85GHz
- Private 128K Level 2 Cache per core and shared 4MB Level 3 Cache
- Dual, multithreaded, on-chip 10GbE ports
- Cryptographic stream processing unit (SPU) integrated in the pipeline
- Dual PCIe Generation 2 x8 interfaces integrated in silicon
- Security features include support for AES, Camellia, CRC32c, DES, 3DES, Kasumi, MD5, RSA, ECC, DSA, SHA-1, SHA-224, SHA-256, SHA-384, and SHA-512

Through its system–on-a-chip design, the SPARC T4 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 T4 to the world of compute-intensive applications as well as the traditionally CMT-friendly data center throughput applications. No-cost security and cryptographic acceleration is provided by the on-chip, per-core streaming accelerators. In addition, the SPARC T4 processor is significantly aided by two integrated PCIe Generation 2 interfaces and dual 10GbE interfaces that help move data in and out quickly. The SPARC T4 processor also implements cache coherency logic and links on the processor silicon to facilitate a multisocket, glueless system design.

SPARC T4-1B Server Module Architecture

The motherboard design for the T4-1B server module is shown in Figure 2.

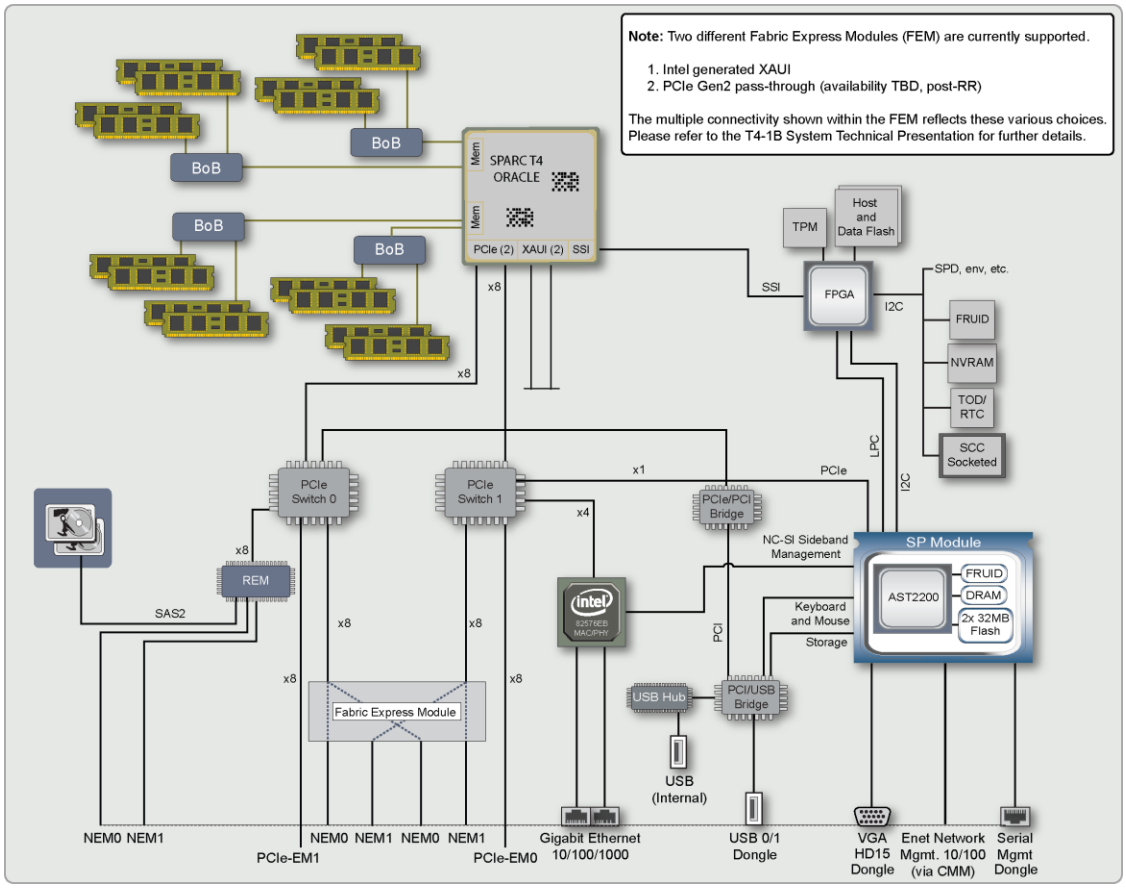

Figure 2. Oracle's SPARC T4-1B server module motherboard design

Enclosure

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

Oracle's SPARC T4-1B server module 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 (FEMs) (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
- Oracle ILOM 3.0 system controller

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 3 provides a block-level diagram of the SPARC T4 processor.

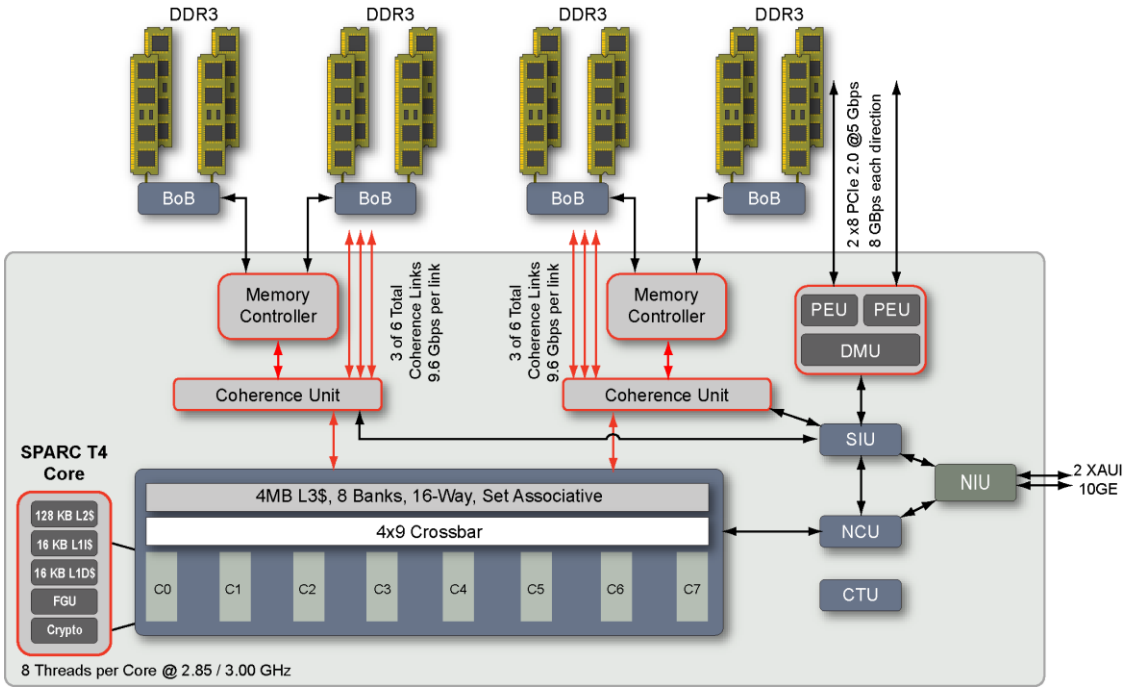

Figure 3. The SPARC T4 processor provides six coherence links to connect to up to four other processors.

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, yielding a maximum frame rate of 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 the SPARC T3 processor.

Integral PCIe 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

SPARC T3-1B Server Module

The SPARC T3 processor is based SPARC T-Series chip multithreading technology (CMT) with industry-leading performance. Ideal workloads for the SPARC T3-1B server are multithreaded workloads with large instruction and data sets, middleware, and application tier workloads, and especially Java environments.

The SPARC T3-1B server module is available with two different T3 processor-based system options, either an 8 core (64 thread), or 16 core (128 thread), single T3 processor. The physical layout out of the SPARC T3-1B is the same for both base systems, as shown in Figure 4.

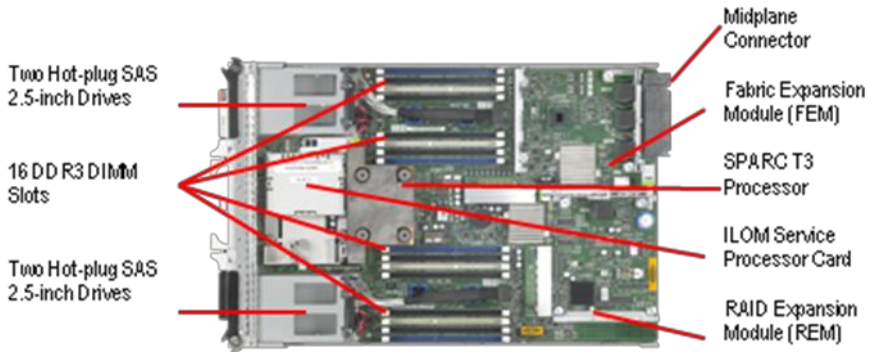

Figure 4. The SPARC T3-1B server module supports a single SPARC T3 processor.

The SPARC T3-1B server module is based on the SPARC T3 processor with up to 16 cores, and 128 simultaneous threads. For high-speed networking, the SPARC T3 uniquely offers dual multithreaded 10 Gigabit Ethernet (10GbE) PCIe integrated onto the processor chip. With integrated 10GbE networking, built-in PCIe Generation 2 expansion, and integrated on-chip cryptographic acceleration, the SPARC T3-1B takes enterprise performance and security to a whole new level delivered in a compact, economical blade form factor.

The SPARC T3-1B includes 16 DIMM slots for up to 128 GB memory, four drive slots for hot-pluggable 2.5 inch SAS drives. The T3-1B supports one optional Fabric Expansion Module (FEM), up to two RAID Expansion Modules (REM), and the ORACLE ILOMORACLE ILCOM Service processor as shown in Figure 4.

The Sun Blade 6000 chassis supports up to 10 SPARC T3-1B server modules for a total of 160 cores and 1,280 threads, and I/O throughput up to 6.4 Terabit-per-second.

SPARC T3-1B Server Module Architecture

Oracle's T3-1B server module motherboard design is similar to that of the T4-1B server module and is illustrated in Figure 5.

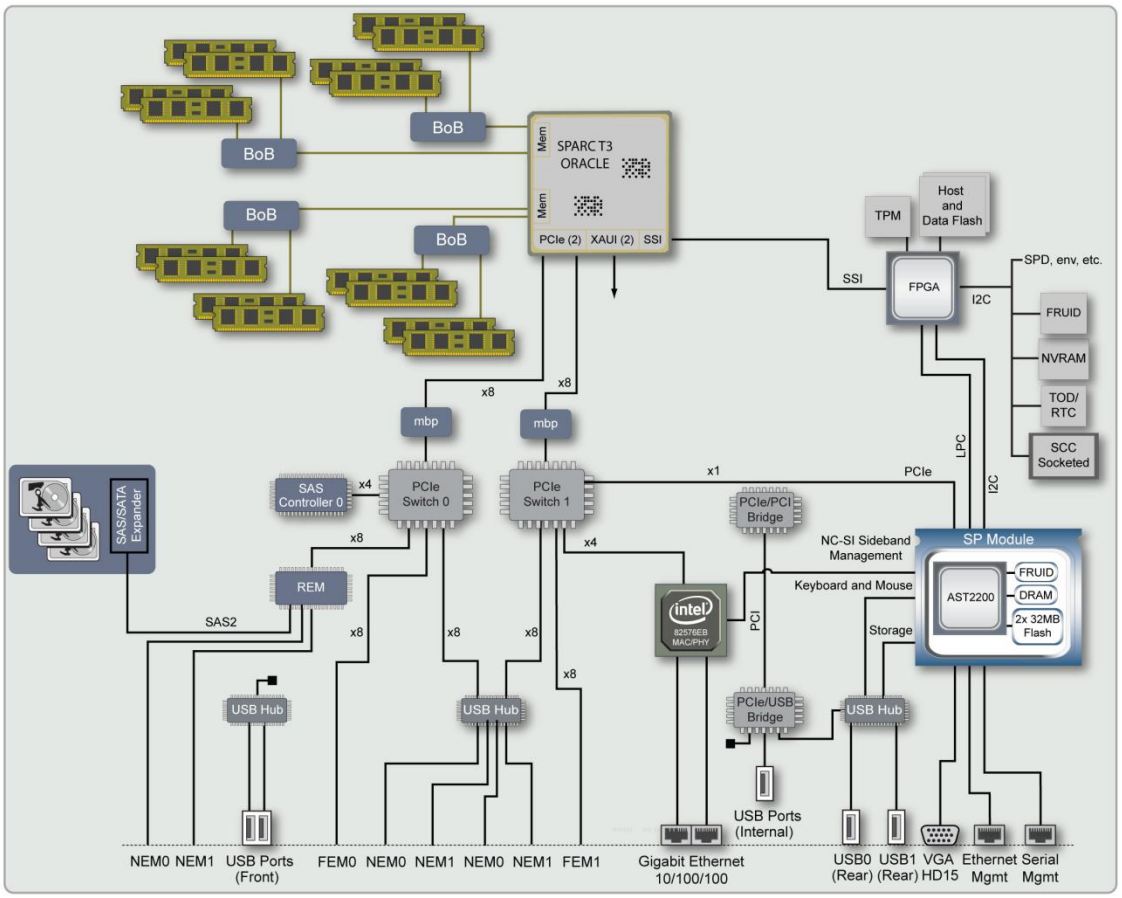

Figure 5. Oracle's SPARC T3-1B server motherboard design.

For I/O, each SPARC T3 processor interfaces through two (x8) PCIe Gen 2 ports capable of operating at 5 Gbps per lane 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.
- **Gigabit Ethernet** — Two 10/100/1000 Mb/sec Ethernet interfaces are provided on the SPARC T3-1B server.
- **Dual 10 Gigabit Ethernet** — Oracle's SPARC T3 processor provides dual 10GbE 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.

The SPARC T3 Processor

Unlike complex single-threaded processors, CMT processors use the available transistor budget to implement multiple hardware multithreaded processor cores on a chip die. SPARC T3 processors take the CMT 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 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 10GbE interfaces
- Two on-chip PCIe Generation 2 (PCIe Gen2) interfaces
- Six on-chip cache coherency links and logic

SPARC T3 Processor Architecture

The SPARC T3 processor extends Oracle's CMT initiative with an elegant and robust architecture that delivers real performance to applications. Figure 6 provides a block-level diagram of the SPARC T3 processor.

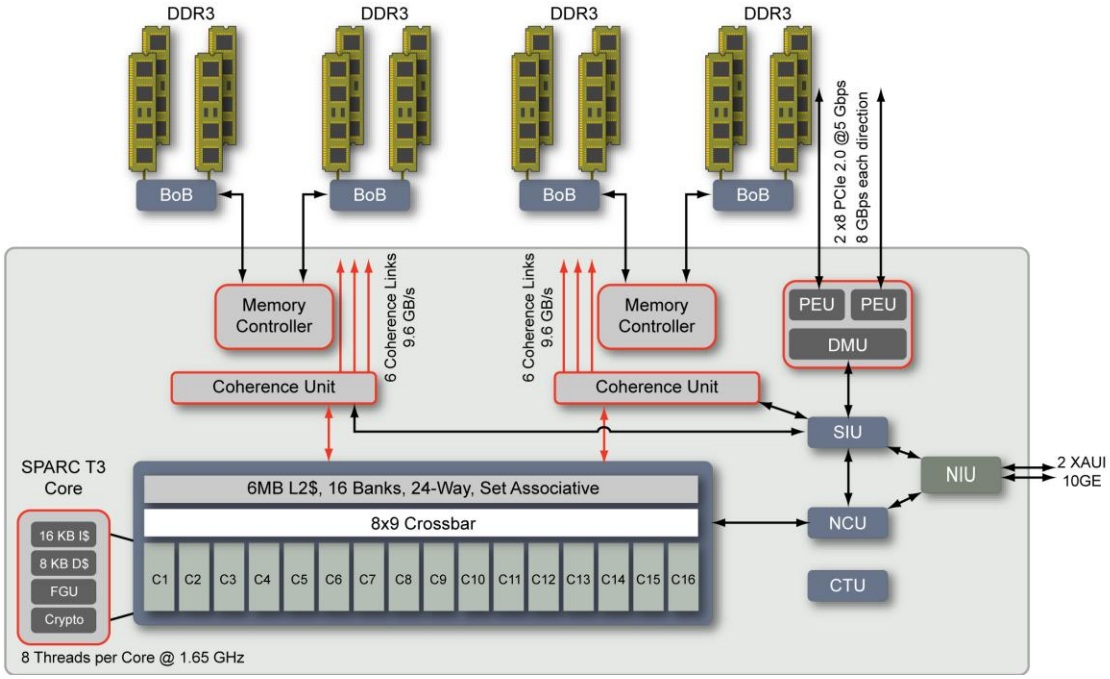

Figure 6. The SPARC T3 processor provides six coherence links to connect to up to four other processors.

The SPARC T3 coherence link interfaces were the foundation for design of future SPARC T4 processors. Like the SPARC T4 processor, the SPARC T3 processor also provides dual on-chip PCIe Generation 2 interfaces. Beyond the inherent efficiencies of CMT design, the SPARC T-Series processor incorporates unique power management features at both the core and memory levels of the processor.

More information regarding SPARC T-Series processor coherency link interfaces, PCIe Gen 2 support, integrated networking architectures, and power management features can be found in the previous section about the SPARC T4 processor.

Sun Blade X6270 M3 Server Module

The Sun Blade X6270 M3 server module improves on previous-generation Sun blade X6270 M2 server modules by offering support for two Intel Xeon Processor E5-2600 product family CPUs. In addition to the increased number of cores offered by these latest CPU's, this server module supports 24 DIMM sockets using low voltage DDR3 DIMMs to accommodate memory-intensive applications. Up to four HDDsor SSDs can be accommodated through front accessible SAS-2 based bays in the front panel. FEMs and RAID Expansion Modules are also supported, allowing for flexible storage and I/O connectivity options.

The Sun Blade X6270 M3 server module builds on the success of previous-generation Sun Blade X6270 M2 server modules by offering a number of enhancements, including:

- Support for Intel Xeon Processor E5-2600 product family CPUs, now with up to eight cores per CPU
- SAS Gen-2 signaling capabilities throughout

A top-view of the Sun Blade X6270 M3 server module is shown in Figure 7.

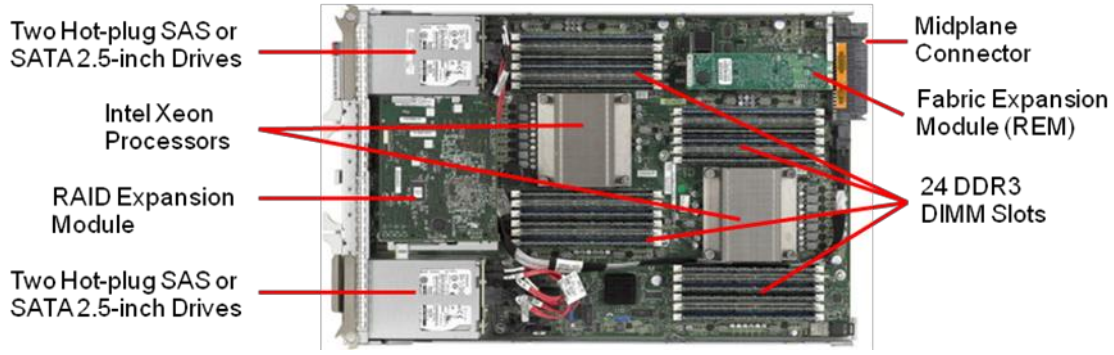

Figure 7. The Sun Blade X6270 M3 server module provides two sockets for quad-core, six-core, or eight-core Intel Xeon Processor E5-2600 product family CPUs

Intel Xeon Processor E5-2600 Product Family

With this latest introduction, Intel continued with a 32 nm manufacturing process, but increased the core computing power from six processor cores to eight processor cores and also increased cache in the die. The new Intel Core micro architecture is extremely modular, enabling a range of implementations to meet a variety of application needs and price points. Available in four-, six-, and eight-core versions, the Intel Xeon Processor E5-2600 product family provide up to 20 MB of shared L3 cache and integrated PCIe 3.0 support. They also utilize Intel Turbo Boost Technology and Intel Hyper-Threading Technology and are equipped with a four-channel DDR3 memory controller, each channel controlling up to three DIMMs. Two bi-directional QuickPath Interconnects are also provided for high-speed CPU-CPU and CPU-I/O subsystem communications.

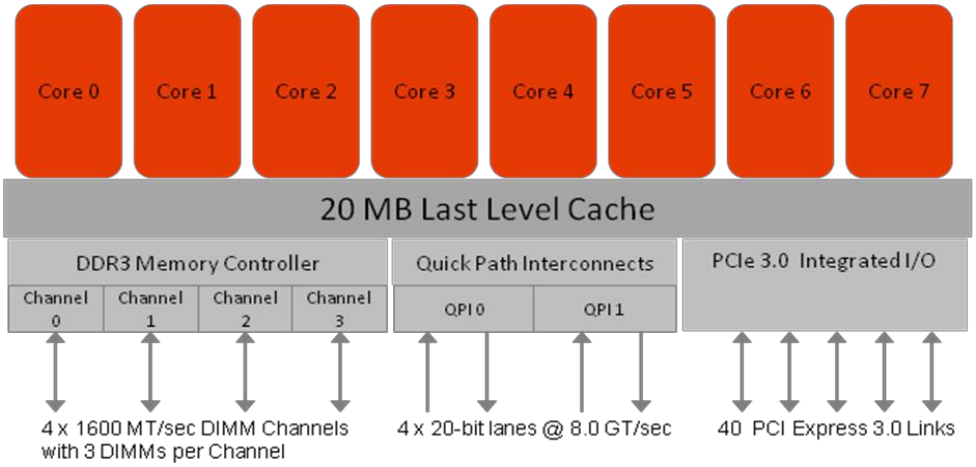

Figure 8. The Intel Xeon Processor E5-2600 product family offers up to eight cores, an integrated memory controller, two bidirectional QuickPath Interconnects, and an integrated PCIe controller.

These new Intel processors are targeted at delivering optimal performance for bandwidth-intensive, threaded applications, with a microarchitecture that features the following significant innovations over previous designs:

- *Intel QuickPath technology* — This technology provides a high-speed, point-to-point interconnect between processors and I/O. The Intel QuickPath Interconnect (QPI) links processors in a design that provides both high-bandwidth and low-latency. Because it is a point-to-point interconnect, processors do not contend for a single bus when accessing memory and I/O, and do not compete for bus bandwidth, enhancing scalability. Each QPI port includes two unidirectional links that support from 6.4 GT/ sec up to 8.0 GT/sec per link, offering up to 16 GB/sec bandwidth in each direction for a total bandwidth of 32 GB/sec — significantly higher than previous bus designs.
- *Multiple processor cores* — The Intel Xeon Processor E5-2600 product family microarchitecture offers four cores, six cores, or eight cores per die.
- *Integrated DDR3 memory controller* — The integrated memory controller provides four 1600 MT/sec channels and each channel supports up to three DIMMs. The processor design creates a NUMA-style memory architecture since each processor in multisocket systems can access local memory (connected to the local memory controller) as well as remote memory that is connected to another processor.
- *Integrated I/O controller* — In previous generation Intel platforms, PCIe I/O was handled by a separate I/O Hub Controller. In the new Intel Xeon Processor E5-2600 product family, the processors have PCIe integrated on the processor die, reducing I/O latency. The new processor also supports the new PCIe 3.0 specification that greatly improves the PCIe bandwidth.
- *Advanced cache model* — The Intel Xeon Processor E5-2600 product family offers a 20 MB Last-Level cache that is shared by the processor cores.

- *Virtualization enhancements* — Embedded virtualization technologies provide hardware-based assistance for I/O device virtualization, improved virtualization efficiency, enhanced connectivity, and improved security within a virtualized server.
- *Intel HyperThreading Technology* — This technology provides two virtual threads per core, increasing performance for highly threaded applications.
- *Intel Turbo Boost Technology* — For both multithreaded and single-threaded workloads, this technology increases performance by taking advantage of processor and system power as well as thermal headroom. The Turbo Boost feature can opportunistically increase performance up to nine speed bins per core (900 MHz) above typical performance levels where thermal headroom exists. Intel Turbo Boost Technology and Hyper-Threading Technology capabilities vary according to specific processor models.
- *Intel Intelligent Power Technology* — When a processor workload decreases, unneeded components — cores, cache, and memory — are put into sleep mode to reduce power consumption.

Sun Blade X6270 M3 Server Module Architecture

As shown in Figure 9, the Sun Blade X6270 M3 server module supports two Intel Xeon Processor E5-2600 CPUs interconnected to each other using Intel's QuickPath Interconnect (QPI) technology. Each Intel Xeon Processor E5-2600 product family CPU has an integrated memory controller with four DDR3 channels and up to three DIMMs per channel. Each processor also provides 40 PCIe 2.0 lanes that connect to the FEM socket that delivers up to two PCIe 2.0 x8 interfaces or XAUI connections to the chassis midplane for NEM connections. Two additional PCIe 2.0 x8 interfaces are delivered to the passive midplane for EM connections.

An x4 PCIe 2.0 connection interfaces to the optional RAID Expansion Module, which provides SAS-2 links to the midplane, and hosts the internal HDDs and/or SSDs. The Intel C602J I/O Controller Hub provides SATA and USB connectivity as well as an interface to the ASPEED service processor. The Intel Xeon processor E5-2600 product family CPUs also interface directly with an Intel I350 Gigabit Ethernet controller (formerly known as Powerville) to provide two Gigabit Ethernet interfaces to the passive midplane. An internal USB connector is also provided by the Intel C602J.

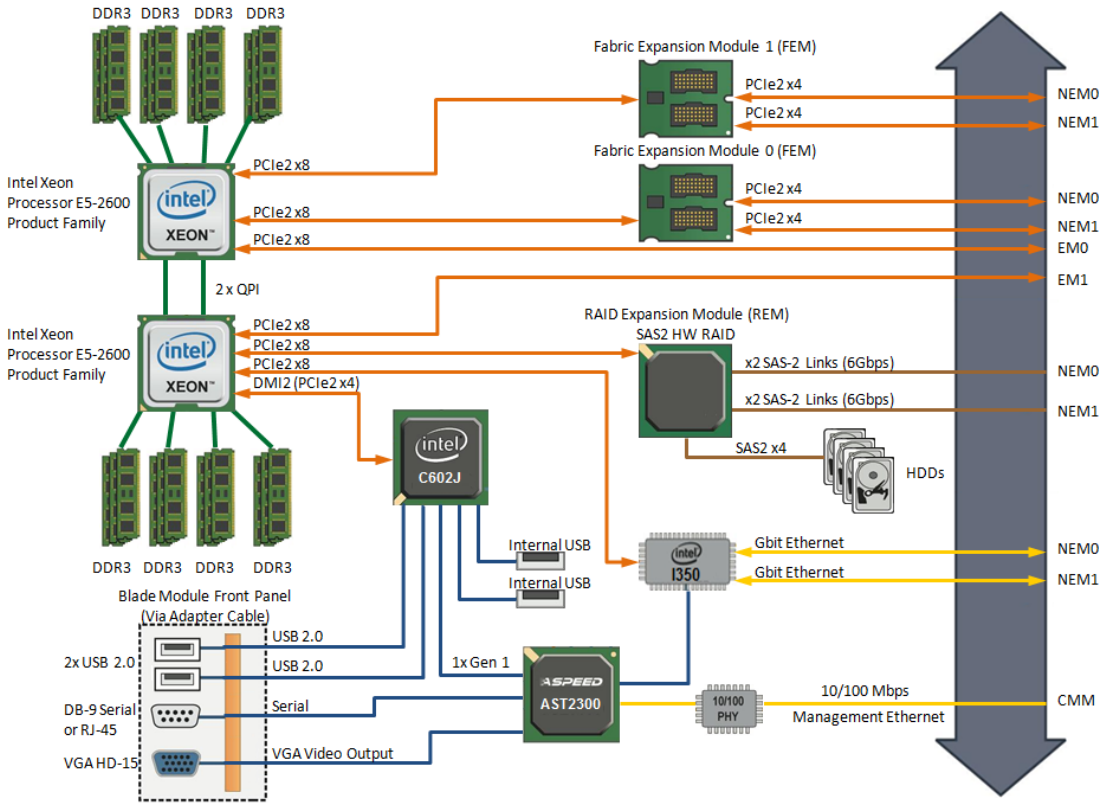

Figure 9. The Sun Blade X6270 M3 server module introduces support for Intel Xeon processor E5-2600 product family CPUs and SAS-2 storage devices

Sun Blade X6270 M2 Server Module

The Sun Blade X6270 M2 server module improves on previous-generation server modules by offering support for up to two Intel Xeon processor 5600 Series CPUs (Westmere-EP). In addition to the increased performance offered by these latest CPU's, this server module supports 18 DIMM sockets using low-voltage DDR3 DIMMs to accommodate memory-intensive applications. Up to four HDDs or SSDs can be accommodated through front-accessible drive bays in the front panel. FEMs and RAID Expansion Modules are also supported, allowing for flexible storage and I/O connectivity options.

The Sun Blade X6270 M2 server module builds on the success of previous-generation server modules by offering a number of enhancements, including:

- Support for Intel Xeon Processor 5600 Series CPUs, now with up to 6 cores per CPU.
- Enhanced Gigabit Ethernet support
- SAS Gen-2 signaling capabilities throughout

A top-view of the Sun Blade X6270 M2 server module is shown in Figure 10.

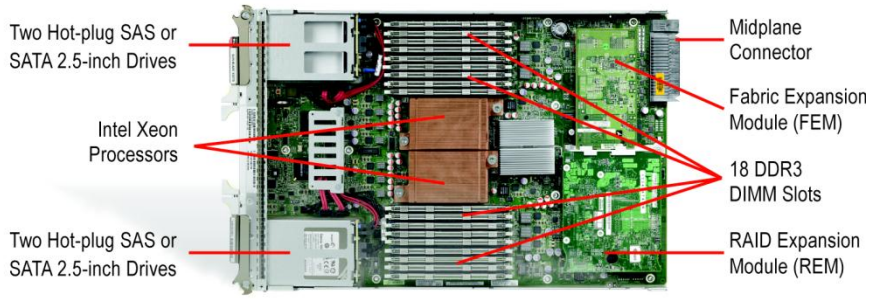

Figure 10. The Sun Blade X6270 M2 server module provides two sockets for quad-core or six-core Intel Xeon Processor 5600 Series CPUs

Intel Xeon Processor 5600 Series

With each release of a new processor series, Intel alternates between enhancing the manufacturing process (shrinking the processor die) and redesigning the core microarchitecture. With the Intel Xeon Processor 5600 Series (Westmere-EP, Figure 11), Intel transitioned to a 32 nm process and added processor cores to the base design of the Intel Xeon Processor 5500 series (described later in this section). Available in either four-core or six-core versions, the Intel Xeon Processor 5600 Series CPU's are equipped with a three channel DDR3 Memory Controller, each channel controlling up to three DIMMs. Two bi-directional QuickPath Interconnects are also provided for high-speed CPU-CPU and CPU-I/O subsystem communications.

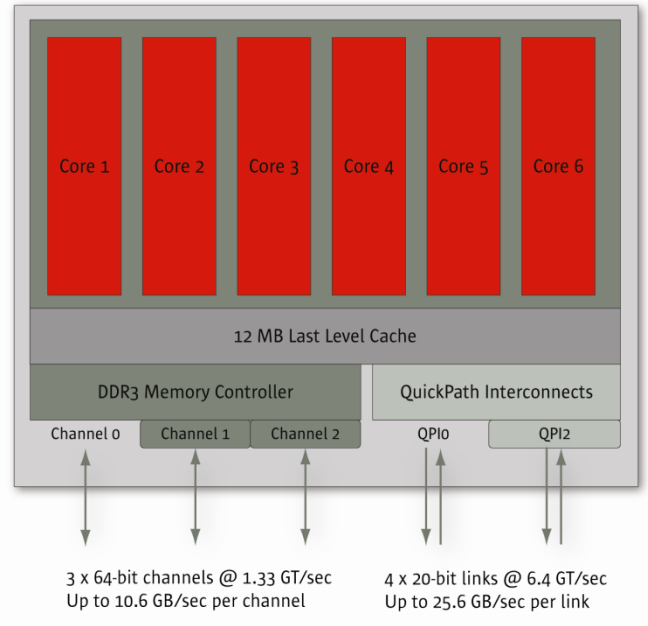

Figure 11. The Intel Xeon Processor 5600 Series offers up to six cores, an integrated memory controller, and two bidirectional QuickPath Interconnects.

These new Intel processors are targeted at delivering optimal performance for bandwidth-intensive, threaded applications, with a microarchitecture that features the following significant innovations over previous designs:

- *Intel QuickPath technology* — This technology provides a high-speed, point-to-point interconnect between processors and I/O. The Intel QuickPath interconnect (QPI) links processors in a design that provides both high-bandwidth and low-latency. Because it is a point-to-point interconnect, processors do not contend for a single bus when accessing memory and I/O, and do not compete for bus bandwidth, enhancing scalability. Each QPI port includes two unidirectional links that support from 4.8 GT/ sec up to 6.4 GT/sec per link, offering up to 12.8 GB/sec bandwidth in each direction for a total bandwidth of 25.6 GB/sec — significantly higher than previous bus designs.
- *Multiple processor cores* — The Intel Xeon Processor 5600 Series microarchitecture offers four cores or six cores per die.
- *Integrated DDR3 memory controller* — Implemented as a separate, remote component in earlier architectures, the memory controller is now integrated on the processor die. The processor design creates a NUMA-style memory architecture since each processor in multisocketed systems can access local memory (connected to the local memory controller) as well as remote memory that is connected to another processor.
- *Advanced cache model* — The Intel Xeon Processor 5600 Series offers a 12 MB Last-Level cache that is shared by the processor cores.
- *Extended SSE4 (Streaming SIMD Extensions)* — These processor extensions improve performance for XML, string, and text processing.

- *Virtualization enhancements* — Embedded virtualization technologies provide hardware-based assistance for I/O device virtualization, improved virtualization efficiency, and enhanced connectivity within a virtualized server.
- *Intel Hyper-Threading (HT) technology* — This technology provides two virtual threads per core, increasing performance for highly threaded applications.
- *Intel Turbo Boost Technology* — For both multithreaded and single-threaded workloads, this technology increases performance by taking advantage of processor and system power as well as thermal headroom. The Turbo Boost feature can opportunistically increase performance up to two or three speed bins per core (266 or 400 MHz) above typical performance levels where thermal headroom exists. Turbo Boost and Hyper-Threading capabilities vary according to specific processor models.
- *Intel Intelligent Power Technology* — When a processor workload decreases, unneeded components — cores, cache, and memory — are put into sleep mode to reduce power consumption.

Sun Blade X6270 M2 Server Module Architecture

As shown in Figure 12, the Sun Blade X6270 M2 server module supports up to two Intel Xeon processors interconnected to each other and the Intel 5600 Chipset using Intel's QuickPath Interconnect (QPI) technology. The Intel Xeon Processor 5600 Series CPUs each have one integrated memory controller with three DDR3 channels and up to three DIMMs per channel. The Intel 5600 Chipset provides 16 PCIe 2.0 lanes to the FEM socket that deliver up to two PCIe 2.0 x8 interfaces or XAUI connections to the chassis midplane for NEM connections. Two additional PCIe 2.0 x8 interfaces are delivered to the passive midplane for EM connections. An x4 PCIe 2.0 connection interfaces to the optional RAID Expansion Module, which provides SAS-2 links to the midplane, and hosts the internal HDDs and/or SSDs.

The Intel ICH10R I/O Controller Hub provides PCI, SATA, and USB connectivity. The Controller Hub interfaces with an Intel 82576EB Gigabit Ethernet controller (formerly known as Kawela) to provide two Gigabit Ethernet interfaces to the passive midplane. An internal USB connector is also provided by the Intel ICH10R.

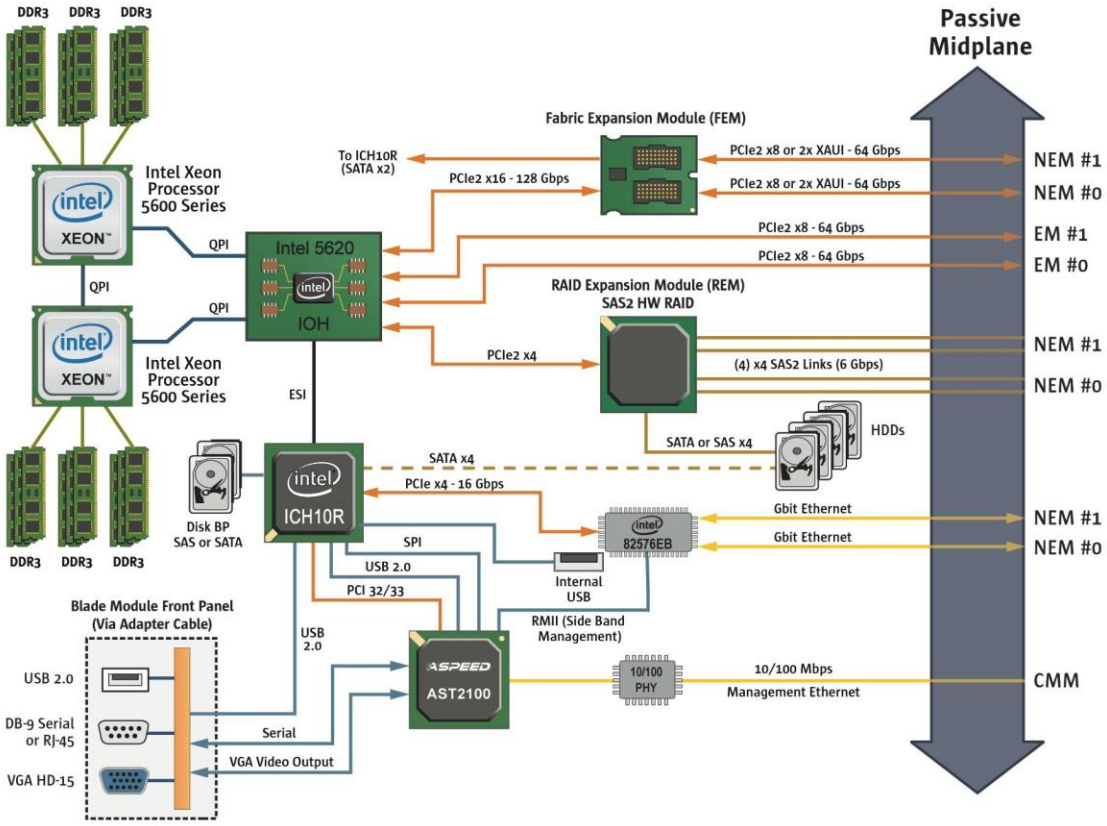

Figure 12. The Sun Blade X6270 M2 server module introduces support for Intel Xeon Processor 5600 Series CPUs and SAS-2 storage devices

Sun Blade X6275 M2 Server Module

The Sun Blade X6275 M2 server module is uniquely designed with two compute nodes configured in a single server module footprint. In addition, the design makes better utilization of the printed circuit board and the dual-node design delivers cost savings as well.

Each of the two compute nodes on the Sun Blade X6275 M2 server module provide two sockets for Intel Xeon Processor 5600 CPUs and 12 DDR3 DIMM slots for memory. An independent Oracle ILOM service processor is also provided for each node. The server module contains no conventional drive slots, but two slots are provided to support Sun Flash Modules, one per compute node.

The Sun Blade X6275 M2 server module has two versions: 10GbE and GbE. Both versions are supported in the Sun Blade 6000 chassis. A photo of the Sun Blade X6275 M2 GbE model is provided in Figure 13, with key components called out.

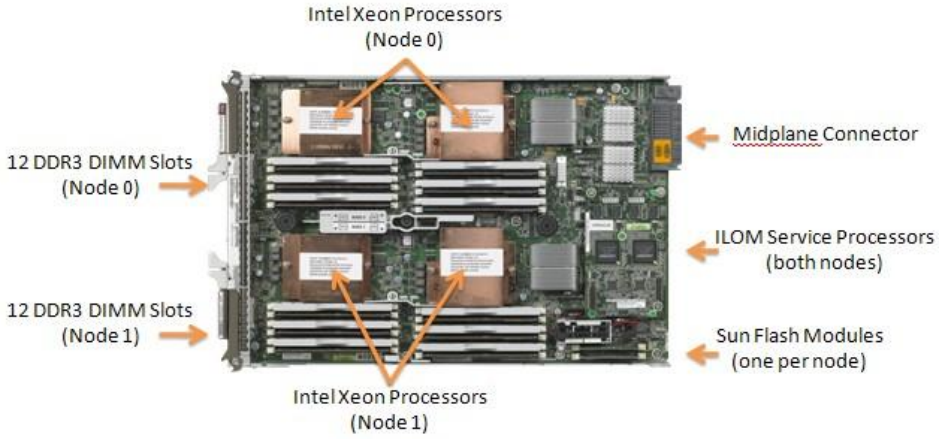

Figure 13. The Sun Blade X6275 M2 server modules provide two nodes on a single server module, each with two sockets for Intel Xeon Processor 5600 Series CPUs.

Sun Blade X6275 M2 Server Module Architecture

Figure 14 shows that the Sun Blade X6275 M2 10GbE server module provides two compute nodes within the standard blade module form-factor. Each node supports two Intel Xeon processors interconnected to each other and the Intel 5600 Chipset using Intel's QuickPath Interconnect (QPI) technology. The Intel Xeon Processor 5600 Series CPUs each have one integrated memory controller with three DDR3 channels and up to three DIMMs per channel. The Intel 5600 Chipset provides eight PCIe 2.0 lanes to the passive midplane for EM connection. The Intel ICH10R I/O Controller Hub provides PCI, SATA, and USB connectivity. The Controller Hub interfaces with an Intel 82576EB Gigabit Ethernet controller (formerly known as Kawela) to provide a Gigabit Ethernet interfaces to the passive midplane. An internal USB connector is also provided by the Intel ICH10R.

The Sun Blade X6275 M2 server module is available in two models: GbE and 10GbE. On the Sun Blade X6275 GbE M2 server module, the 1GbE connection from compute node 0 is routed to NEM #0 and the 1GbE connection from compute node 1 is routed to NEM #1. Additionally, each node is provided with access to a single PCIe ExpresssModule slot. The Sun Blade X6275 M2 10GbE server module also provides two compute nodes within the standard blade form-factor. Each compute node provides a 10GbE network connection to a NEM. The 10GbE connection from compute node 0 is routed to NEM #0 and the 10GbE connection from compute node 1 is routed to NEM #1. The 10GbE model leverages Mellanox ConnectX®-2 controllers. Like the Sun Blade X6275 M2 GbE server module, each node is provided with access to a single PCIe ExpresssModule slot.

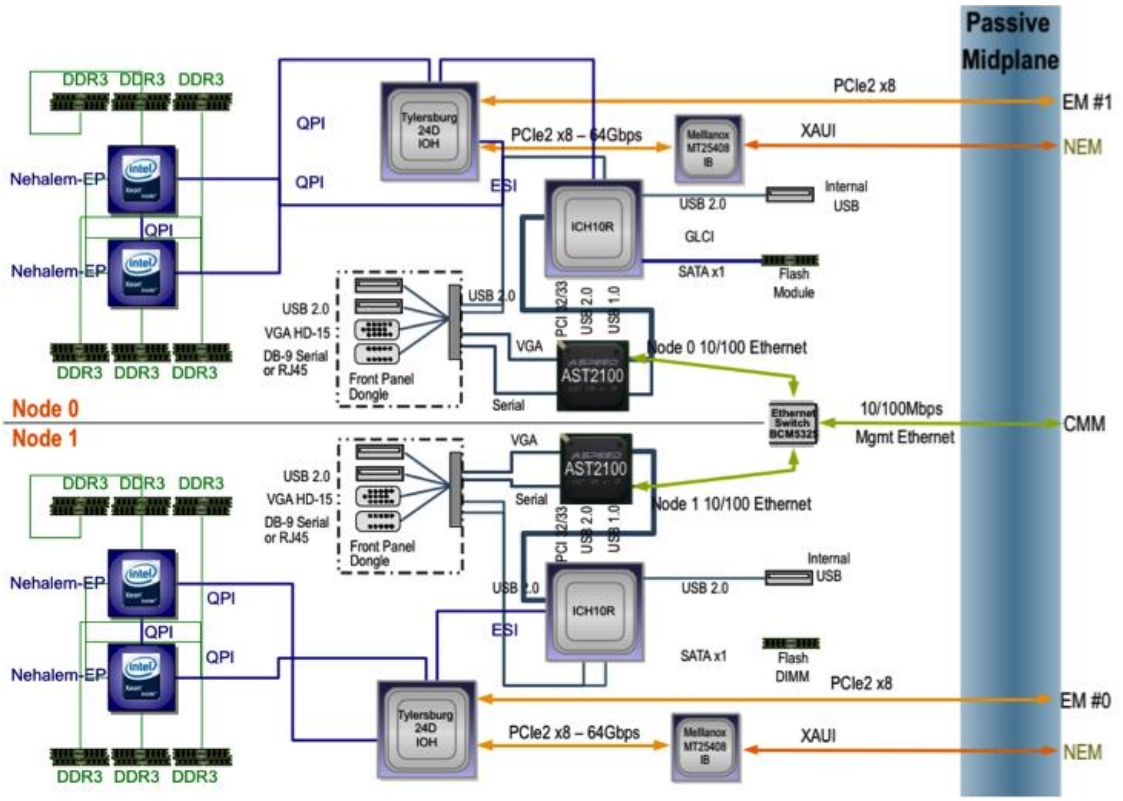

Figure 14. The Sun Blade X6275 M2 server modules provide two nodes on a single server module, each with two sockets for Intel Xeon Processor 5600 Series CPUs. This figure shows the 10GbE version of the Sun Blade X6275 M2 server module

Conclusion

Oracle's innovative technology and open-systems approach make modular systems attractive across a broad set of applications and activities — from consolidating infrastructure through virtualization to deploying dynamic enterprise applications or cloud computing. Oracle's Sun Blade 6000 modular system provides the promised advantages of modular architecture while retaining essential flexibility for how technology is deployed and managed.

With a chassis designed for investment protection into the future, organizations can literally cable once, and change their deployment options as required — mixing and matching server modules as desired. A choice of Oracle x86 and SPARC processor-based server modules along with support for a range of operating systems as well as blade-specific I/O capabilities makes it easy to select the right platform for Oracle and non-Oracle enterprise applications.

For More Information

For more information on Sun Blade 6000 modular systems, please visit *http://www.oracle.com/goto/blades*, which also contains links to additional Sun Blade 6000 modular system white papers. The Web sites listed below also provide more specific references:

- *Sun Blade systems OTN page: http://www.oracle.com/technetwork/server-storage/sun-blade/documentation/index.html*
- Blades power calculator*: http://www.oracle.com/sun-power-calculators/calc/6000chassis*
- Oracle Enterprise Manager Ops Center: *http://www.oracle.com/us/products/enterprise-manager/opscenter/*
- Oracle Solaris operating system*: http://www.oracle.com/technetwork/server-storage/solaris/overview/*.

Sun Blade 6000 Server Module Architecture
April 2012 Version 1.3

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

Hardware and Software, Engineered to Work Together