



SUN SERVER

An Oracle White Paper  
July 2011; Revised July 2012

# Oracle's Sun Server X2-8 System Architecture

---

Introduction .....	1
Oracle's Sun Server X2-8 System .....	2
The Sun Server X2-8 System Architecture .....	5
CPU Module .....	6
Memory .....	9
PCI Express ExpressModule .....	10
Network Express Module .....	12
Internal Drives .....	13
Chassis and External Interfaces .....	13
Cooling .....	14
Power .....	15
System Management.....	16
ILOM Service Processor.....	16
SNMP Support .....	20
Oracle® Enterprise Manager Ops Center Software .....	21
A Choice of Operating Systems.....	21
Oracle Solaris Operating System.....	22
Linux and Microsoft Windows Operating Environments .....	23
Virtualization Platforms .....	23
Conclusion .....	23
For More Information .....	24
Appendix A—Intel® Technology.....	25
Intel Xeon Processor E7-8800 Product Family .....	26
Intel 7500 Scalable Memory Buffer .....	28

## Introduction

The design characteristics of the Sun Server X2-8 system are redefining the x86 enterprise market. The superior architecture and significant price-per-performance advantages offered by the Sun Server X2-8 system make this platform an ideal replacement for many types of systems including HP Itanium and IBM Power servers.

In addition to providing superior performance, the Sun Server X2-8 system offers significant processing, memory, and I/O expansion capabilities. The architecture of the Sun Server X2-8 system accommodates up to eight CPUs, providing a total of up to 80 processing cores. Memory-intensive applications benefit from the ability of this platform to support a 4 TB memory footprint. Furthermore, support for up to eight hot-swappable PCI Express ExpressModules (EMs) help architects create configurations that offer high availability for I/O connections and an ample number of I/O paths, as well as eight additional 10 GbE connections and eight 1-GbE connections available on two NetworkExpressModules (NEMs).

## Oracle's Sun Server X2-8 System

Oracle's Sun Server X2-8 systems (Figure 1) combine Oracle's system design expertise and the latest Intel core microarchitecture to offer exceptional performance and I/O expansion in a compact five rack unit (5 U) chassis. Their high-performance, modular design makes these systems ideal for back-end workloads, such as in-memory databases, data warehouses, consolidation projects, and enterprise high-performance computing (HPC) applications.

To support demanding applications, the Sun Server X2-8 system includes some of the fastest technology available. It allows insertion of up to four CPU modules (CMODs), each supporting two Intel Xeon Processor E7-8800 product family CPUs, 32 DIMM slots, and an Intel 7500 I/O Controller Hub (IOH). Within this architecture, Intel QuickPath Interconnects (QPIs) support high-speed, low-latency communication between processors and from processors to IOHs. Four Scalable Memory Interfaces on the die of Intel Xeon E7-8800 product Family CPUs support fast data transfer to and from memory DIMMs. In addition, the Sun Server X2-8 system utilizes the Intel Xeon 7500 chipset to establish connectivity to I/O components, including two Network Express Modules (NEMs) and up to eight hot-swappable PCI Express ExpressModules (EMs).

Two cooling zones and independent fan speed control within Sun Server X2-8 systems help optimize airflow, minimizing noise and reducing power requirements for cooling. To simplify system management, the Sun Server X2-8 system includes an integrated service processor that offers access to management functions using an Ethernet, serial, or Universal Connector Port (UCP)-enabled video connection. As with all of Oracle's Sun Fire servers, the Sun Server X2-8 system supports simultaneous execution of 32-bit and 64-bit applications and offers a choice of operating systems including, Oracle® Solaris, Oracle Linux, Oracle VM, SUSE Linux, and Microsoft Windows operating environments, providing extensive flexibility and investment protection.




Figure 1. The high-performance, compact Sun Server X2-8 system can support demanding business application workloads.

As shown in Table 1, the characteristics of Sun Server X2-8 system offer benefits for compute and memory intensive workloads, such as database engines, consolidation projects, and enterprise high-performance computing (HPC) applications.

**TABLE 1. THE BENEFITS OF SUN SERVER X2-8 SYSTEMS BY WORKLOAD TYPE**

WORKLOAD TYPE	APPLICATION	BENEFITS
Databases	<ul style="list-style-type: none"> <li>In-memory databases</li> <li>Data warehousing</li> </ul>	<ul style="list-style-type: none"> <li>Large memory footprint</li> <li>Affordability</li> <li>Industry-standard, high-performance processors</li> <li>Significantly increased performance and reduced latency</li> </ul>
Consolidation	<ul style="list-style-type: none"> <li>Enterprise Resource Planning (ERP)</li> <li>Customer Relationship Management (CRM)</li> <li>Database front end</li> </ul>	<ul style="list-style-type: none"> <li>Large memory footprint</li> <li>Increased RAS with hot-swappable I/O components</li> </ul>
Enterprise HPC	<ul style="list-style-type: none"> <li>Memory-intensive simulations</li> <li>Workload balancing</li> </ul>	<ul style="list-style-type: none"> <li>Support for up to ten cores per processor</li> <li>Large memory footprint</li> <li>Significant price-per-performance advantages over proprietary solutions</li> </ul>

To meet the challenges of demanding applications, the Sun Server X2-8 architecture accommodates up to eight high-performance Intel Xeon E7-8800 product family CPUs, up to 4 TB of memory, and up to eight hot-swappable PCI Express ExpressModules (EMs) supported by PCI Express Gen2 data transfer buses. Table 2 details the technical specifications and expansion capabilities of the Sun Server X2-8 system.

**TABLE 2. CHARACTERISTICS OF THE SUN SERVER X2-8 SYSTEM**

COMPONENT	SUN SERVER X2-8 SYSTEM TECHNICAL SPECIFICATIONS
Enclosure	5 rack unit (5 U)
Processor	Four or eight Intel Xeon E7-8800 product family CPUs: <ul style="list-style-type: none"> <li>Intel Xeon processor E7-8870 (10-core, 2.4 GHz, 30 MB cache, 6.4 GT/sec, QPI, 130 Watts)</li> <li>Intel Xeon processor E7-8860 (10-core, 2.26 GHz, 24 MB cache, 6.4 GT/sec. QPI, 130 Watts)</li> <li>Intel Xeon processor E7-8830 (8-core, 2.13 GHz, 24 MB cache, 6.4 GT/sec. QPI, 105 Watts)</li> </ul>
Memory	128 DIMM slots (1066 MHz) Up to 4 TB of low-voltage memory (128 GB minimum) 4 GB, 8 GB, 16 GB or 32 GB DDR3 ECC Registered DIMMs

COMPONENT	SUN SERVER X2-8 SYSTEM TECHNICAL SPECIFICATIONS
Internal Storage	<p>Up to eight SAS internal disk drives</p> <p>300 GB or 600 GB 2.5-inch SAS-2 drive (10,000 RPM)</p> <p>300 GB 2.5-inch SATA-1 SLC SSD drive</p> <p>One SAS-2 RAID Express Module (REM)</p>
System I/O Interfaces	<p>One internal USB 2.0 port</p> <p>Two on-board Gigabit Ethernet (GbE) ports per CPU module (accessible via the NEMs)</p>
Network Express Modules	<p>Two NEMs standard:</p> <ul style="list-style-type: none"> <li>• Four 10 GbE ports (SFP+ pluggable transceivers) per NEM</li> <li>• Two x4 mini SAS-2 ports per NEM</li> <li>• Four GbE ports (RJ-45 connectors) per NEM: <ul style="list-style-type: none"> <li>○ 10/100/1000 Mb/sec. full- and half-duplex operation</li> <li>○ IEEE 802.3ab Auto-Negotiation for speed, duplex, and flow control</li> <li>○ PXE Boot is possible on all four interfaces</li> </ul> </li> </ul>
PCI Express ExpressModules	<p>Up to eight PCI-SIG form-factor compliant EMs</p> <p>Two per CPU module</p> <p>Eight PCI Express Gen2, 5.0 GT/s lanes per EM</p> <p>Maximum of 32 Gb/s bandwidth per EM</p>
Power	Four redundant, hot-swappable power supplies
System Management	<p>Oracle Integrated Lights Out Manager (Oracle ILOM) Service Processor:</p> <ul style="list-style-type: none"> <li>• Remote management with full keyboard, mouse, video, storage (KVMS)</li> <li>• Remote media capability (floppy, CD, and so on)</li> <li>• Full Distributed Management Task Force (DMTF) CLI</li> <li>• Browser UI for control of the system through a graphical interface</li> <li>• IPMI 2.0 compliant for management and control</li> <li>• SNMP v1, V2c, V3 for system monitoring</li> <li>• Monitor and report system and component status on all FRUs</li> </ul> <p>Service Processor I/O interfaces:</p> <ul style="list-style-type: none"> <li>• Two dedicated management 10/100Base-T Ethernet ports</li> <li>• One TIA/EIA-232-F asynchronous RJ-45 port</li> <li>• One UCP interface, supporting a video port, serial port, and two USB ports (requires use of a UCP dongle)</li> </ul>

## The Sun Server X2-8 System Architecture

The Sun Server X2-8 system design is based on a highly-integrated, modular system architecture that emphasizes processor performance, memory capacity, and I/O expansion.

The server implements the New Intel Core Microarchitecture, utilizing Intel Xeon Processor E7-8800 product family (formerly Westmere-EX) CPUs and the Intel Xeon 7500 Chipset. The Intel Xeon Processor E7-8800 product family microarchitecture offers eight or ten cores per die, two integrated memory controllers, and a large Level-3 cache. Advanced capabilities, such as Intel TurboBoost technology can help maximize performance by temporarily allowing processors to operate at higher-than-stated frequency levels when the system is functioning below the target power and Thermal Design Points (TDP). In addition, embedded virtualization technologies provide hardware-based assistance for I/O device virtualization, improved virtualization efficiency, and enhanced connectivity within a virtualized server.

The Sun Server X2-8 system is a multi-socket compute platform built with modular components. The 5 U enclosure supports four CPU modules, four hot-pluggable power supplies, a pair of NEMs, up to eight EMs, up to eight internal solid state or conventional hard disk drives, and a set of four fan modules. In addition, a chassis monitoring module provides Integrated Lights Out Management (ILOM) and monitoring capabilities. Figure 2 provides a high-level illustration of the Sun Server X2-8 architecture.

The sections that follow detail the physical and architectural aspects of this system. Details about the Intel Xeon Processor E7-4800 product family and its advantages can be found in Appendix A.



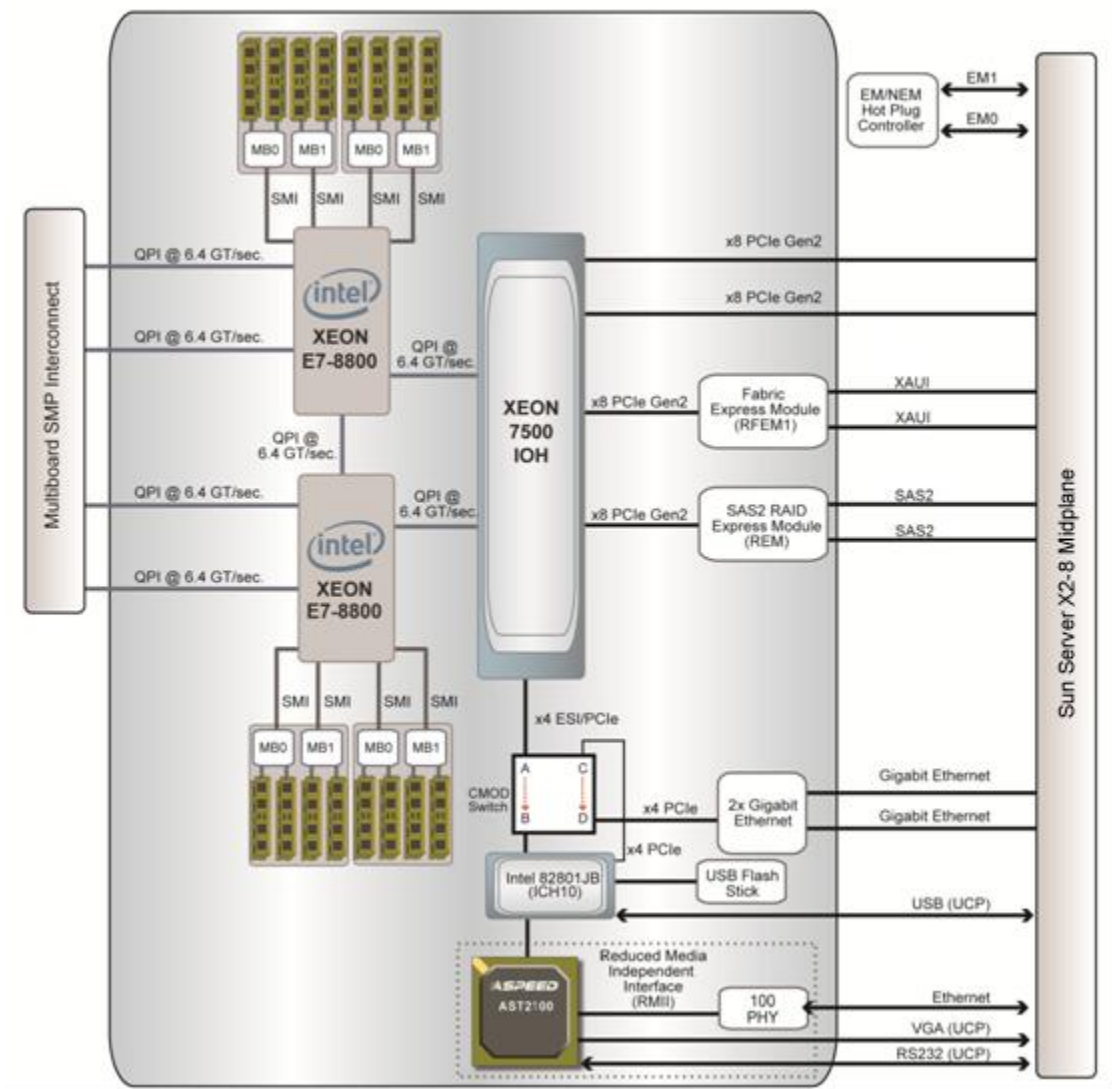
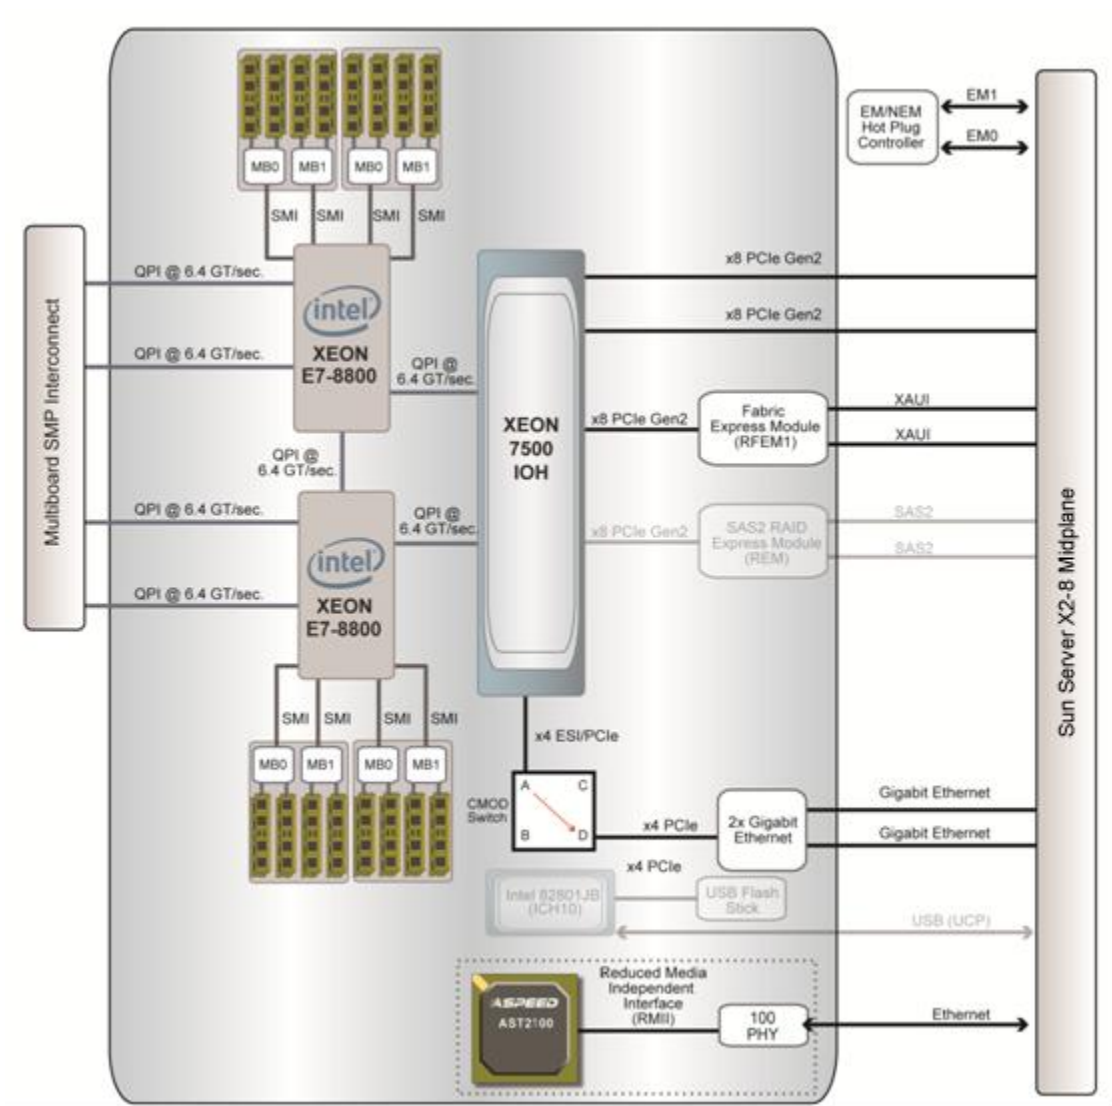


Figure 3. The CPU module in slot 0 acts as a master interface to the REM, USB port, and ILOM service processor.

When a CMOD is placed in slots 1, 2, or 3, the intelligent switch configures itself as shown in Figure 4, essentially disconnecting from the ICH10. In addition, the REM is disabled and the local ILOM service processor becomes a slave to the ILOM service processor on the CMOD in slot 0. Figure 4 depicts a CPU module seated in slot 1, 2, or 3.



**Figure 4.** CPU modules seated in slots 1, 2, and 3 provide CPU and memory capacity, Gigabit Ethernet connections, and QPI links to other CPU modules.

The design of the Sun Server X2-8 CPU module includes the following:

- Two sockets populated with Intel Xeon Processor E7-8800 product family CPUs
- Dual integrated memory controller on each processor
- Intel 7500 I/O Controller Hub (IOH)
- The Intel QuickPath Interconnect architecture, providing 6.4-GT/sec links, delivering up to 25.6 GB/sec of total bandwidth
- SMP midplane interconnect
- Interconnect for two 8x PCIe 2.0 I/O Express Modules

- Two XAUI interconnects to Network Express Modules (NEM), one interconnect per NEM
- Two Gigabit Ethernet ports (available via the NEMs)
- Support for one Fabric Express Module (FEM)

The following architectural features of the CPU module are available for only CMOD0.

- The Intel 82801JB I/O Controller Hub, supporting connectivity to legacy interfaces
- On-board storage:
  - Internal USB 2.0 port for Flash memory stick
  - One SAS-2 RAID Express Module (REM), supporting connectivity to storage fabric
- Connectivity to the ILOM service processor with an ASPEED AST2100 integrated communications processor, providing system management, monitoring, and remote keyboard, video, mouse, storage (KVMS) capabilities

## Memory

Each Sun Server X2-8 CPU module supports up to 32 JEDEC standard Registered DDR3 ECC DIMMs (RDIMMs) for a maximum of 128 RDIMMs for an eight-socket configuration and 64 RDIMMs for a four-socket configuration. The memory architecture of the Sun Server X2-8 system supports a maximum of 4 TB of memory, when using 32 GB DIMMs. Memory interface slots can be populated with 4 GB, 8 GB, 16 GB or 32 GB DDR3 DIMMs running at 1066 MegaTransfer/second (MT/sec). Memory fillers are required to fill empty DIMM slots, helping to optimize system cooling.

With the introduction of the new Intel Xeon E7-8800 product family, low-voltage DIMMs are now supported. With the support of low-voltage DIMMs and integrated reduced-power memory buffers, operating costs can be reduced further by decreasing energy demand.

The following DIMM population rules and best practices can help organizations create configurations with the best possible performance:

- Populate slots on each CPU module with pairs of identical DIMMs; mixing DIMM sizes within the same CPU module is not supported.
- Balance DIMMs across CPU modules. To help this effort, DIMM slots are color coded and the memory population order is presented on the service label inside the Sun Server X2-8 system chassis top cover.
- Installing a uniform amount of memory across CMODs is recommended. For example, all channels should be populated with a single DIMM before adding a second DIMM to a channel.
- Using higher rank DIMMs results in better performance; dual-rank DIMMs are better than single-rank DIMMs.
- If low-voltage DIMMs are mixed with standard 1.5V DIMMs, the low-voltage DIMMs will operate at 1.5V instead of 1.35V.

## PCI Express ExpressModule

The Sun Server X2-8 system supports up to eight PCI-SIG form-factor compliant PCI Express ExpressModules (EMs). This approach increases the flexibility of the Sun Server X2-8 system by allowing the use of numerous hot-pluggable expansion module options from multiple expansion module vendors. Sample images of commonly utilized EMs are provided in Figure 5.



**Figure 5. Examples of PCI-SIG form-factor EMs.**

Oracle supports a number of types of EMs for the Sun Server X2-8 system, including the following: !

- Quad-port Gigabit Ethernet, PCI Express x4 !
- 10 Gigabit Ethernet PCI Express !
- Dual-port quad data rate InfiniBand !
- Dual-port 8Gb/sec Fibre Channel plus dual-port Gigabit Ethernet !
- Eight-port SAS-2

The Sun Server X2-8 system passive midplane facilitates connectivity between the EMs and the CPU ! modules. Pairs of EMs are physically connected to individual CPU modules as follows: !

- Slots 0 and 1 are paired to IOH0 on CMOD0. !
- Slots 2 and 3 are paired to IOH2 on CMOD1. !
- Slots 4 and 5 are paired to IOH3 on CMOD2. !
- Slots 6 and 7 are paired to IOH1 on CMOD3. !

Figure 6 provides an illustration of the CMOD and EM pairings. For best performance, balance ! high-traffic EM cards among the available IOHs. In addition, the installation of a filler panel is required for configurations that do not utilize all eight EM slots. The use of a filler panel helps to maintain proper airflow, optimizing system cooling.




Figure 6. Each CPU module controls a pair of EMs.

The recommended population order for a four-socket Sun Server X2-8 system is listed in Table 3.

TABLE 3. RECCOMENDED EM POPULATION ORDER FOR FOUR-A SOCKET SUN SERVER X2-8 SYSTEM

POPULATION ORDER	SLOT NUMBER	DESCRIPTION
First	EM Slot 3.0	Assigned to IOH1 on CMOD3
Second	EM Slot 0.0	Assigned to IOH0 on CMOD0 Shared if REM (x4) is present
Third	EM Slot 3.1	Assigned to IOH1 on CMOD3
Fourth	EM Slot 0.1	Assigned to IOH 0 on CMOD0

Table 4 provides the recommended population order for an eight-socket Sun Server X2-8 system.

TABLE 4. RECCOMENDED EM POPULATION ORDER FOR AN EIGHT-SOCKET SUN SERVER X2-8 SYSTEM

POPULATION ORDER	SLOT NUMBER	DESCRIPTION
First	EM Slot 3.0	Assigned to IOH1 on CMOD3
Second	EM Slot 1.0	Assigned to IOH2 on CMOD1
Third	EM Slot 0.0	Assigned to IOH0 on CMOD0 Shared if REM (x4) is present
Fourth	EM Slot 2.0	Assigned to IOH3 on CMOD2
Fifth	EM Slot 3.1	Assigned to IOH1 on CMOD3
Sixth	EM Slot 1.1	Assigned to IOH2 on CMOD1
Seventh	EM Slot 0.1	Assigned to IOH0 on CMOD0
Eighth	EM Slot 2.1	Assigned to IOH3 on CMOD2

## Network Express Module

A Network Express Module (NEM) is a hot-swappable component that provides multiple 10 GbE, single GbE, and SAS connections to each CMOD in the Sun Server X2-8 system chassis. Although similar in principle to the NEMs within Sun Blade 6000 Modular Systems, the form-factor of this specific type of NEM is unique to the Sun Server X2-8 system chassis. Sun Server X2-8 systems come standard with one NEM. As an option, a second NEM can be added to the configuration. Up to 88 Gb/sec of network bandwidth is available when both NEMs are installed.

Each Sun Server X2-8 system NEM provides the following interfaces:

- Four 10 GbE ports with SFP+ connectors
- Two x4 mini SAS-2 ports
- Four Gigabit Ethernet ports with RJ-45 connectors

To support these ports, each NEM provides two dual 10 Gigabit Ethernet PHYs and one SAS-2 Expander ASIC. The logic for the four Gigabit Ethernet ports actually resides on the CPU modules (two per CPU module). The Sun Server X2-8 system's midplane provides access and connectivity from the Gigabit Ethernet ASIC on the CPU modules through to the NEM interfaces. Figure 7 illustrates the architecture of the Sun Server X2-8 NEM.

Each of the two Sun Server X2-8 NEMs supports two miniSAS form-factor connectors on the NEM face plate, each with 4x SAS-2/SATA-2 links capable of running at 6.0 Gb/sec, 3.0 Gb/sec, and 1.5-Gb/sec rates. The miniSAS connector links connect to an expander on the NEM, allowing the port to be shared by any of the four CPU modules with a RAID Express Module attached (REM).

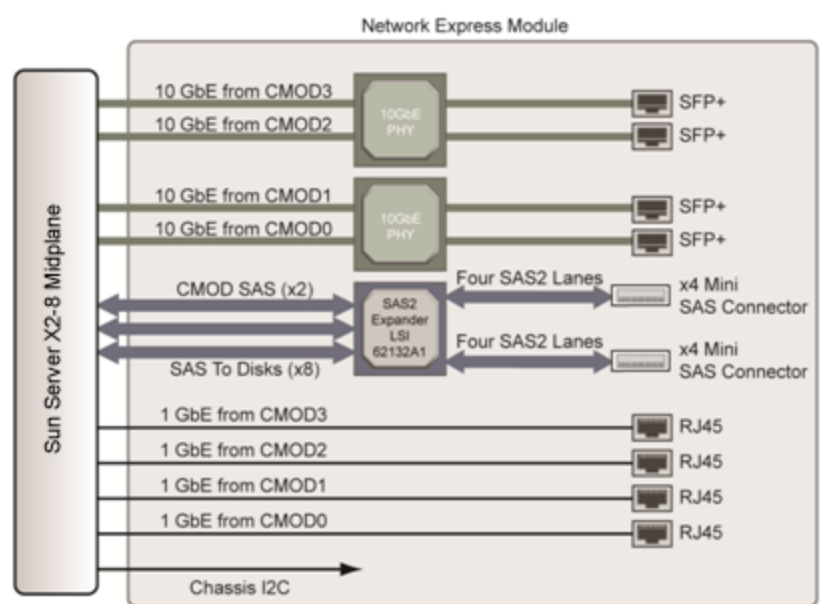


Figure 7. The Sun Server X2-8 NEM provides multiple Ethernet and SAS-2 interfaces.

## Internal Drives

The Sun Server X2-8 system provides eight internal disk drive slots. These slots can be populated with traditional hard disk drives or solid state drives. At product release, 300 GB and 600 GB 10000-rpm 2.5-inch SAS-2 hard disk drives are supported. 300 GB 2.5-inch SATA-1 SLC SSD drives are also supported. Using the eight internal hot-swappable drive bays, the Sun Server X2-8 system can support up to a maximum of 4.8 TB (using 600 GB drives).

## Chassis and External Interfaces

The Sun Server X2-8 system is a 5 rack unit (5 U) chassis. Figure 8 provides a front view of the Sun Server X2-8 system and Figure 9 provides a rear view.

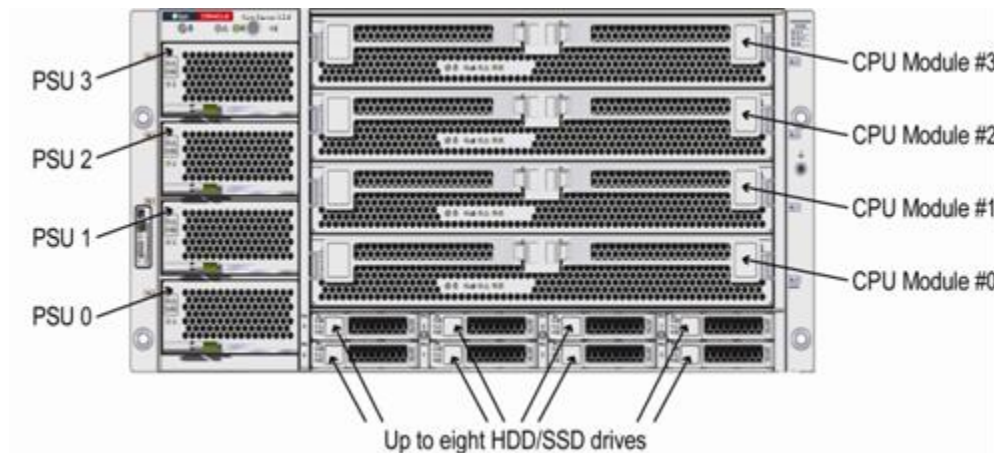


Figure 8. Front view of the Sun Server X2-8 system.

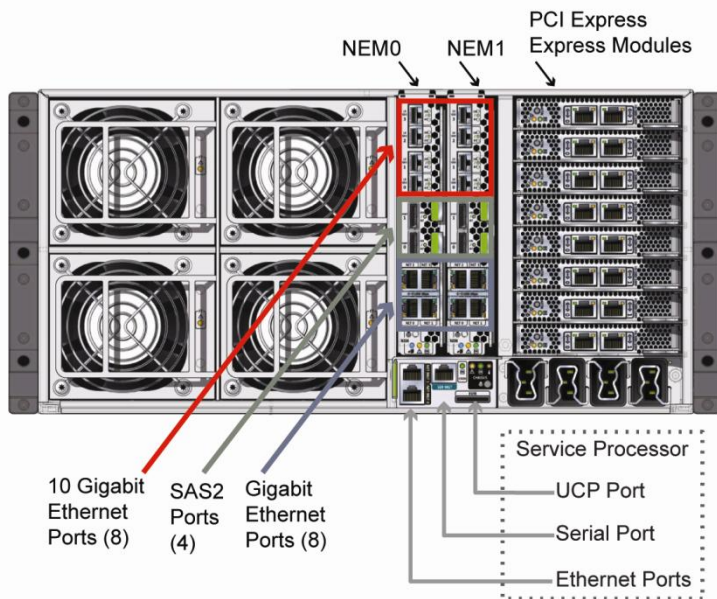


Figure 9. Rear view of the Sun Server X2-8 system.

## Front Panel

The front panel of the Sun Server X2-8 system offers the following indicators:

- Locate LED (white)
- Service action required LED (amber)
- Power OK LED (green): Blinks when in standby power mode. Always on when in full power mode.
- Power button
- Over Ambient Temperature fault LED (temperature)

## Fault LEDs

Fault-remind buttons help identify or locate a faulty component even after the system has been disconnected from the power source or a module has been removed from the chassis. The CPU module, NEM, and service processor components provide an amber fault LED that activates when the corresponding in-module fault-remind button is pressed. For a DIMM fault, pressing the system fault-remind button lights the fault LED on the module where the faulty component resides.

## Cooling

For increased cooling efficiency, the Sun Server X2-8 system chassis implements front-to-back cooling with two independent cooling zones, as shown in Figure 10. Within Zone 1, EMs and NEMs are cooled by airflow from the AC-DC PSU fans. In Zone 2, CPU modules are cooled by four 92-mm fans controlled by the ILOM service processor. Temperature sensors monitor each zone and the ILOM service processor controls the fan speeds to keep cooling levels optimized. The four hot-swappable fans for Zone 2 are arranged in two rows of two fans for redundancy. Air is pressurized in the main compartment for front-to-back cooling. Airflow is also unrestricted by this design, resulting in less noise. To more fully optimize cooling, filler panels must be installed for empty NEM, EM, DIMM, and disk drive slots.

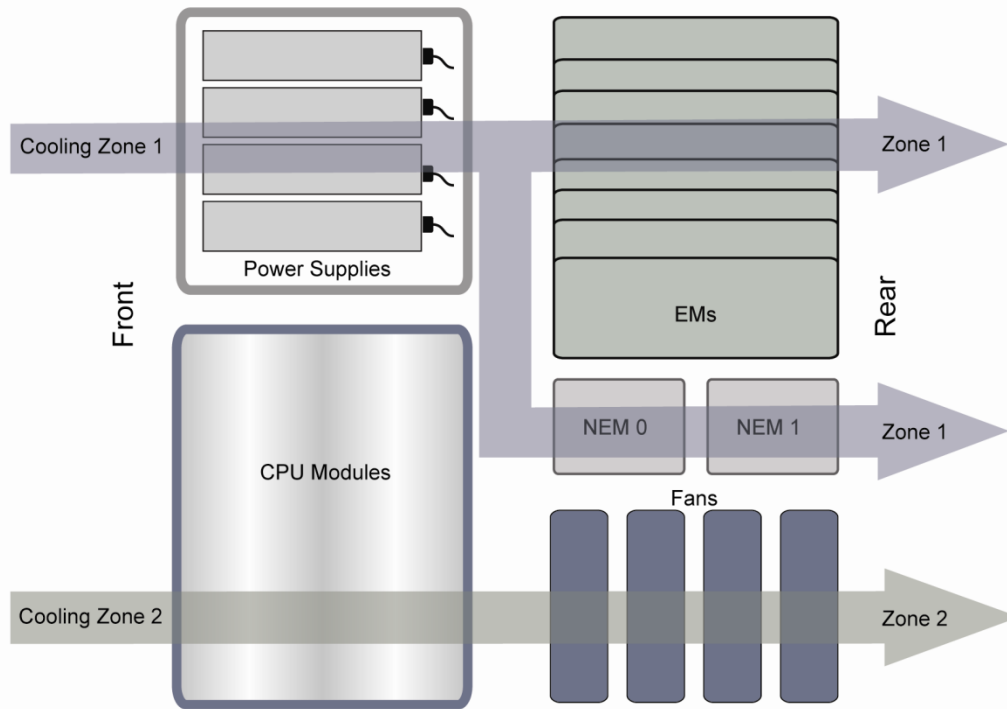


Figure 10. The use of multiple cooling zones helps optimize the cooling efficiency of the Sun Server X2-8 system.

## Power

The Sun Server X2-8 system features redundant power supplies that are rated at 12 V at 167 Amps. Power output is rated 2060 W at 200–240 V. Powering on the Sun Server X2-8 system requires that two PSUs are connected into a high-line AC (200–240 Volt) power source.

If needed, the Sun Server X2-8 system PSUs can support power over-subscription. This means that the PSUs in the Sun Server X2-8 system are designed to support 50% power overload for at least 10 minutes. In order to return to normal operation, the ILOM service processor invokes software power-capping algorithms to reduce system power consumption until a level below the intended maximum rating for the PSU is reached.

PSUs are known to operate at peak efficiency at near 50% of the maximum load rating. To increase the likelihood of operating within this range, the PSUs within the Sun Server X2-8 systems support Light Load Efficiency Mode (LLEM). If the power load falls below a set threshold, the ILOM Service Processor enables LLEM. In LLEM mode, a power supply can be set to warm-standby mode. While on standby, the power supply does not provide power to the system, yet remains ready to take over in the event of a fault within one of the other power supplies. When the power load is above a given threshold, power supplies share the load equally. For deployments that prioritize redundancy over efficiency, LLEM mode can be disabled from the ILOM Service Processor configuration.

## System Management

To address the need for remote system monitoring and management, the Sun Server X2-8 system provides an integrated ILOM service processor. The Sun Server X2-8 system also supports open standards to simplify integration with add-on Oracle and third-party enterprise management software tools.

### ILOM Service Processor

Lights-out management features provided by the Sun Server X2-8 system ILOM service processor can help organizations simplify system management tasks. Powered by an ASPEED AST2100 integrated communications processor, the service processor runs independent of the host platform, executing a robust, security-hardened operating system.

Capabilities of the ILOM service processor include the following: !

- Full local and remote keyboard, video, mouse, and storage (KVMS) access via redirection over IP, ! eliminating the need for KVM switches
- Monitoring and reporting of environmental, power, hardware, BIOS, and operating system events
- Remote power control, diagnostics, media attachment, and upgrades of the system BIOS and service processor software
- System configuration information retrieval
- User-configurable serial console access through a physical port or redirected through the management network
- Java™ technology-enabled remote console access across a secure Web connection
- Multilevel role-based access with support for RADIUS, LDAP, LDAP/SSL, and Microsoft Active Directory Service lookup of authentication data
- User SSH key authentication and predefined root and user accounts
- Ability to disable the network management port, as well as individual services, such as IPMI, SSH, and KVMS
- Simple Network Management Protocol (SNMP) V1, V2c, and V3 support

For system management operations, the AST2100 uses the following connections:

- UCP port, supporting one VGA interface, two USB ports, and one serial console connection
- One RJ-45 RS-232 serial interface for console redirection
- Two RJ-45 10/100 MB/sec Ethernet network interfaces for IP-based management connections

Management functions provided by the service processor are implemented by ILOM 3.0 system management software. This software provides an Intelligent Platform Management Interface (IPMI) Baseboard Management Controller, platform control agents, diagnostics software, and remote KVMS

drivers. Many other Sun servers from Oracle incorporate this same firmware, providing organizations with a single, consistent, and standards-based management interface.

Secure access to the service processor and associated ILOM software functions is available in several ways, including the following:

- Intuitive browser-based user interface (BUI) over SSL
- Distributed Management Task Force (DMTF) command line interface over Secure Shell (SSH)
- Redirection of the platform console, keyboard, mouse, and video to the Sun ILOM Remote Console application
- SNMP v3 interfaces, which provide easy integration with Oracle Enterprise Manager Ops Center, or third-party applications from companies such as Hewlett-Packard and IBM
- IPMI 2.0 command interface for remote management using IPMI-based tools, such as IPMItool

### Maximizing Management Flexibility

While system management tools can play an important role in streamlining operations, organizations must consider the best approach for each environment. Executing management software directly on the host with or without use of a service processor is known as in-band management. Using a dedicated Ethernet or serial port to execute administrative tasks independent of the host is known as out-of-band management. The Sun Server X2-8 system supports both in-band and out-of-band management. Table 5 offers a comparison of in-band and out-of-band management strategies.

STRATEGY	CHARACTERISTICS	BENEFITS	IDEAL USE CASES
In-band	<ul style="list-style-type: none"> <li>• Relies upon operating system-resident software.</li> <li>• Management tasks utilize platform compute resources.</li> </ul>	<ul style="list-style-type: none"> <li>• Use of a single network connection and switch port minimizes cost and complexity.</li> </ul>	<ul style="list-style-type: none"> <li>• Heterogeneous environments with requirements to provide a common administrative tool across all platforms.</li> </ul>
Out-of-band	<ul style="list-style-type: none"> <li>• Utilizes a dedicated Ethernet or serial port for administrative traffic.</li> <li>• Management tasks execute on an independent service processor.</li> </ul>	<ul style="list-style-type: none"> <li>• Provides continuous access to management capabilities even when host is disabled.</li> <li>• Management tasks do not consume host resources.</li> <li>• Provides increased security by physically separating management traffic and server data.</li> </ul>	<ul style="list-style-type: none"> <li>• Environments with compute or bandwidth-intensive applications.</li> <li>• Projects with complex management requirements or high levels of administrative burden.</li> </ul>

### In-Band Server Management

In-band server management offers the opportunity for organizations to take advantage of industry-standard protocols and applications across all datacenter platforms. The Sun Server X2-8

system facilitates in-band server management by supporting the IPMI 2.0 and SNMP v1, v2c, and v3 standards. Operating system-resident platform management functions for the Sun Server X2-8 system are accomplished using one of the following two options:

- IPMI with a Keyboard Controller Style (KCS) interface and an IPMI kernel driver
- SNMP agents

#### **Out-of-Band Server Management**

An out-of-band management approach supports the completion of administrative tasks without placing unnecessary burden on the host. This administrative strategy is desirable for performance-intensive environments. While in-band management works only as long as the host operating system is up and running, out-of-band management is fully functional even while the host is powered off. Out-of-band management offers the following capabilities and benefits:

- Based on serial port redirection (serial-over-LAN), serial port connectivity to the ILOM service processor provides direct console access to the command line interface (CLI) and to the system console stream. The CLI is designed to follow the Distributed Management Task Force (DMTF) Command Line Protocol (CLP).
- Utilizing Ethernet connectivity, administrators can access a Web interface or command-line interface (CLI). A secure shell (SSH) session is required for access to the CLI and the Web interface supports both secure (HTTPS) and non-secure (HTTP) access. Secure access is the default configuration for Web-based access.

#### **Remote Keyboard, Video, Mouse, and Storage**

The ILOM service processor provides access to KVMs remotely over IP. Remote video display is accomplished through the Java Web Start software known as Sun ILOM Remote Console. To set up a system as a remote console, ILOM Remote Console software is downloaded from the ILOM service processor to the target machine<sup>1</sup>. From this point onward, the ILOM Remote Console executes locally. Since ILOM Remote Console does not run on the server, it does not put overhead on the host.

ILOM Remote Console software can be used to redirect the BIOS and setup screens as well as all other platform video output. A true remote video console to the management console is provided by handling the input and output to and from virtual devices and the Sun Server X2-8 system. With 8-bit and 16-bit support and 8 MB of video memory, the SVGA display provides resolutions up to 1024 x 768. A single instance of ILOM Remote Console can open multiple sessions to manage several remote servers simultaneously.

---

<sup>1</sup> ILOM Remote Console requires the installation of a Web browser with Java Runtime Environment 5.0 plugins on the management console.

The two USB 2.0 ports connected to the ILOM service processor enable remote keyboard, mouse, and storage functions. The ILOM Remote Console software captures keyboard, mouse, floppy, CD, and DVD input on the management console and redirects it over IP to the ILOM service processor. Keyboard, mouse, and storage inputs are then transmitted over the USB ports to the server. The Sun Server X2-8 system interprets these inputs as originating from locally connected USB devices. These devices are referred to as virtual devices.

ILOM Remote Console can also be used to boot the remote server from a local device. The virtual device can be a local physical device or an image file. Several types of devices can be redirected as virtual devices by ILOM Remote Console, including the following:

- CD-ROM/DVD-ROM
- Floppy
- USB Flash disk drives
- CD-ROM/DVD-ROM image (.iso files)
- Floppy image (.img files)

#### **Intelligent Platform Management Interface**

Intelligent Platform Management Interface (IPMI) refers to the autonomous monitoring, logging, recovery, and inventory control features that are implemented in hardware and firmware. The key differentiation of Intelligent Platform Management is that these functions are independent of the main CPU, BIOS, and OS. There are two major components of platform management: the system management controller (or BMC) and System Management Software (SMS). Intelligent Platform Management provides key capabilities that allow enterprise-class management for high-availability systems.

The ILOM service processor provides autonomous sensor monitoring and event logging. Typical sensor-related events are out-of-range temperature or voltage and fan failure. When an event occurs, it is noted in the system event log and made available to the system management controller. The system management controller is powered by the power supply stand-by voltage and functions even when the server is powered down or the operating system has crashed. As a result, platform status can be obtained and recovery initiated under situations where in-band delivery mechanisms are unavailable.

In modern systems, the Intelligent Platform Management Interface (IPMI) provides a hardware-level interface specification for monitoring and control functions. It defines a standard, abstract, message-based interface between the BMC and SMS and a common set of commands for operations such as accessing sensor values, setting thresholds, logging events, and controlling a watchdog timer. IPMI messages can be used to communicate with the BMC over serial and LAN interfaces, so software designed for in-band (local) management can be reused for out-of-band (remote) management simply by changing the low-level communications layer.

## IPMItool

IPMItool is a simple command-line interface to systems that support the IPMI v2.0 specification. IPMItool provides the ability to remotely read the sensor data repository and print sensor values, display the contents of the system event log, print field-replaceable unit information, read and set LAN configuration parameters, and perform remote chassis power control. IPMItool was originally written to take advantage of IPMI-over-LAN interfaces. This tool is also capable of interfacing with the system through a Linux kernel device driver, such as OpenIPMI or the Oracle Solaris BMC driver within the Oracle Solaris 10 IPMItool, and it is available under a BSD-compatible license.

System Management Software (SMS) is generally complex and is only part of a much larger management picture. However, system administrators and developers can use command-line tools and scripting to manage these complex systems. IPMItool takes a different approach to SMS and provides a completely command line-oriented tool.

IPMItool is not designed to replace the OpenIPMI library. Where possible, IPMItool supports printing comma-separated values for output to facilitate parsing by other scripts or programs. IPMItool is designed to run quick command-response functions that can be as simple as turning the system on or off, or as complex as reading in the sensor data records and extracting and printing detailed sensor information for each record.

## SNMP Support

The Simple Network Management Protocol (SNMP) provides remote access to monitor and control network devices and to manage configurations, statistics collection, performance, and security on a network. SNMP is a network management protocol used almost exclusively in TCP/IP networks. The Sun Server X2-8 system provides SNMP MIBs (Management Information Base) to manage and monitor the servers using any SNMP-capable network management system, such as HP OpenView Network Node Manager (NNM), Tivoli, CA Unicenter, or IBM Director. The MIB data describes the information being managed, reflects current and recent server status, and provides server statistics.

The ILOM service processor supports SNMP v1, v2c, and v3. SNMP v3 is enabled by default; v1 and v2c are disabled by default. SNMP sets may be enabled or disabled, and they are disabled by default. SNMP traps can be generated from within the service processor. An IPMI-specific trap, called a Platform Event Trap (PET), may also be generated. The following SNMP MIBs are supported:

- The system group and SNMP group from the RFC1213 MIB
- SNMP-FRAMEWORK-MIB
- SNMP-USER-BASED-SM-MIB
- SNMP-MPD-MIB
- ENTITY-MIB
- SUN-PLATFORM-MIB

## Oracle® Enterprise Manager Ops Center Software

Oracle Enterprise Manager Ops Center software is a highly scalable datacenter management platform that provides organizations with systems lifecycle management and process automation capabilities to help simplify consolidated platform management, compliance reporting, and system provisioning tasks. Oracle Enterprise Manager Ops Center provides a single console to help discover, provision, update, and manage globally dispersed heterogeneous IT environments, which may include Oracle and non-Oracle hardware running Windows, Linux, and Oracle Solaris operating systems. The capabilities of Oracle Enterprise Manager Ops Center can help organizations to provision and administer both physical and virtual datacenter assets.

When used with Sun Server X2-8 systems, Oracle Enterprise Manager Ops Center can automate patch lifecycle management and maintenance. Oracle Enterprise Manager Ops Center can help system administrators automate software installations, simulation, rollback, compliance checking, reporting, and many other related activities. The framework of Solaris JumpStart software is utilized by Oracle Enterprise Manager Ops Center to effectively provision Oracle Solaris onto individual domains. Oracle Enterprise Manager Ops Center also helps facilitate and control administrative actions from a central location to ensure accountability and auditing. These automation capabilities can be used for knowledge-based change management with existing configuration management investments. Taking advantage of Oracle Enterprise Manager Ops Center can help organizations create a more reliable environment that offers considerable cost savings through maintenance reduction and the ability to rapidly rebuild systems as needed.

## A Choice of Operating Systems

Compatibility with a variety of 32-bit and 64-bit operating systems can help protect investments in server hardware. The Sun Server X2-8 system supports multiple operating systems, letting organizations deploy a choice of application environments without having to shift hardware platforms when software requirements change. This added flexibility can reduce cost and complexity and help organizations to increase return on investment while lowering risk.

The Sun Server X2-8 system supports the following operating systems and versions:

- Oracle Solaris 11
- Oracle Solaris 10 09/10 (Update 9)
- Oracle Solaris 11 Express
- Oracle VM 2.2.1, 2.2.2, 3.0.2, 3.0.3
- Oracle Linux 5.6, 6.0, 5.7, 5.8, 6.1, 6.2
- SUSE Linux Enterprise Server 11 SP1, SP2
- Red Hat Enterprise Linux 5.6, 6.0, 5.7, 5.8, 6.1, 6.2
- Microsoft Windows 2008 R2 Standard/Enterprise/Data Center Edition (64-bit)

- VMware ESX 4.1 Update 1 and VMware ESXi 4.1 Update 1, 5.0

The latest operating system support information can be found at

## Oracle Solaris Operating System

As a part of a broad strategic alliance, Oracle and Intel work together—from design and architecture through implementation—to optimize Oracle Solaris and unleash the power and capabilities of current and future Intel Xeon processors. Oracle Solaris and the Intel Xeon processor E7 family-based platform work together to maximize the performance and efficiency of the Sun Server X2-8 system, providing the following:

- **Improved performance.** Oracle Solaris takes advantage of unique Intel Xeon Processor E7-8800 product family features, including Intel HyperThreading Technology, Intel Turbo Boost Technology, and Intel QuickPath Technology, to drive exceptional performance. Oracle Solaris provides an outstanding threading model for commercial workloads, outperforming the competition on customer applications as well as in industry-standard benchmarks. With specific optimizations for Intel Xeon Processor E7-8800 product family CPUs, Oracle Solaris enables new levels of performance as applications incorporate multithreaded design, increasing throughput, responsiveness, efficiency, scalability, and overall performance.
- **Automated power efficiency and utilization.** Oracle Solaris has been optimized to take advantage of Intel's power management functions, improving energy efficiency and performance-per-watt. Intel Intelligent Power Technology enables policy-based control that allows processors to operate at optimal frequency and power. Oracle Solaris can make this determination automatically, or administrators can designate which applications require high-frequency processing and which should be executed at lower frequencies to conserve power. Automated low-power states automatically put processor and memory into the lowest available power states that will meet the requirement of the current workload, greatly enhancing power efficiency.
- **Increased reliability, availability, and serviceability (RAS).** The Oracle Solaris Fault Management Architecture (FMA) infrastructure is enhanced to take advantage of the Intel Xeon Processor E7-8800 product family RAS features to provide a dependable enterprise compute platform. For example, Intel MCA recovery enables the system to detect and correct errors in memory and cache that were previously “uncorrectable” through ECC or other means. MCA accomplishes this task by first detecting and containing errors before the data is consumed by an application, and then working with Oracle Solaris to determine the best course of action to keep the system and application running. This advanced recovery capability means that systems equipped with Intel Xeon Processor E7-8800 product family CPUs and Oracle Solaris can recover and continue to operate in situations where other x86-based systems might fail to continue operating.
- **Virtualization enhancements.** Oracle Solaris takes advantage of Intel virtualization technology features to offer cost-effective virtualization options. In fact, many of the unique, built-in features of Oracle Solaris contribute to the value of virtualization running on Intel Xeon Processor E7-8800 product family CPUs. Oracle Solaris offers a top-to-bottom engineered approach to virtualization where the hardware, the hypervisor, the operating system, networking, and the Oracle Solaris

Zetabyte File System (ZFS) are all designed to deliver optimal performance and manageability. In addition, Oracle Solaris offers features that help improve overall IT efficiency. For example, Oracle Solaris Containers safely isolate resources, allowing enterprise administrators to maintain the one-application-per-server deployment model when consolidating many single servers into a larger system equipped with Intel Xeon Processor E7-8800 product family CPUs.

Additional details at [Oracle Solaris optimizations for Intel architectures](#)

## Linux and Microsoft Windows Operating Environments

Sun Server X2-8 systems support Oracle Linux and SUSE Linux operating environments. Organizations that choose to utilize Oracle Linux can take advantage of the Oracle Unbreakable Linux support program—affordable enterprise-class Linux support that includes premier backports, comprehensive management, cluster software, and indemnification.

Oracle also certifies the Microsoft Windows operating environment on Sun Server X2-8 systems. Support for Microsoft Windows increases the flexibility of Sun Server X2-8 systems and provides organizations the freedom to choose solutions that best meet business needs.

## Virtualization Platforms

Oracle VM is a no-cost, next-generation server virtualization and management solution that makes enterprise applications easier to deploy, manage, and support. Backed worldwide by affordable enterprise-quality support, Oracle VM facilitates the virtualization of enterprise application workloads on Sun Server X2-8 systems. Adopting Oracle VM for Sun Server X2-8 system deployments can help reduce operations and support costs while simultaneously increasing IT efficiency and agility.

## Conclusion

The Sun Server X2-8 system redefines the x86 marketplace with leading performance, enterprise-level RAS features, and unmatched scalability. By surpassing today's expectations for x86 system scalability and expansion limits, the design of the Sun Server X2-8 system continues Oracle's tradition as a leader in the enterprise x86 market. Built to accelerate application performance and maximize compute capacity, the Sun Server X2-8 system supports up to eight Intel Xeon E7-8800 product family CPUs, up to 4 TB of memory, two NEMs, and up to eight PCI Express EMs. Ideal for rack-dense system environments, Sun Server X2-8 systems also offer a compact 5 U design and features that enhance power and cooling efficiency.

## For More Information

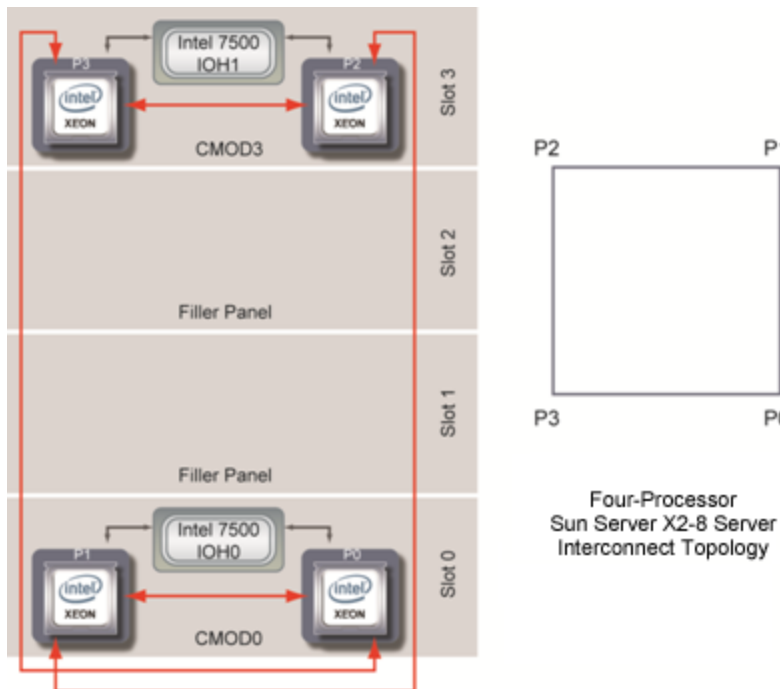
To learn more about the capabilities and benefits of Oracle's Sun Server X2-8 systems, please contact an Oracle sales representative, or consult the related Web sites listed in Table 6.

**TABLE 6. WEB SITES FOR MORE INFORMATION**

DESCRIPTION	URL
Sun x86 systems from Oracle	<a href="http://www.oracle.com/goto/x86">http://www.oracle.com/goto/x86</a>
Sun Server X2-8 system	<a href="http://www.oracle.com/goto/x2-8">http://www.oracle.com/goto/x2-8</a>
Oracle Solaris	<a href="http://www.oracle.com/us/products/servers-storage/solaris">http://www.oracle.com/us/products/servers-storage/solaris</a>
Oracle VM	<a href="http://www.oracle.com/us/technologies/virtualization/oraclevm/index.html">http://www.oracle.com/us/technologies/virtualization/oraclevm/index.html</a>
Oracle Linux	<a href="http://www.oracle.com/us/technologies/linux/index.html">http://www.oracle.com/us/technologies/linux/index.html</a>
Oracle Enterprise Manager Ops Center software	<a href="http://www.oracle.com/us/products/enterprise-manager/opscenter">http://www.oracle.com/us/products/enterprise-manager/opscenter</a>

## Appendix A—Intel® Technology

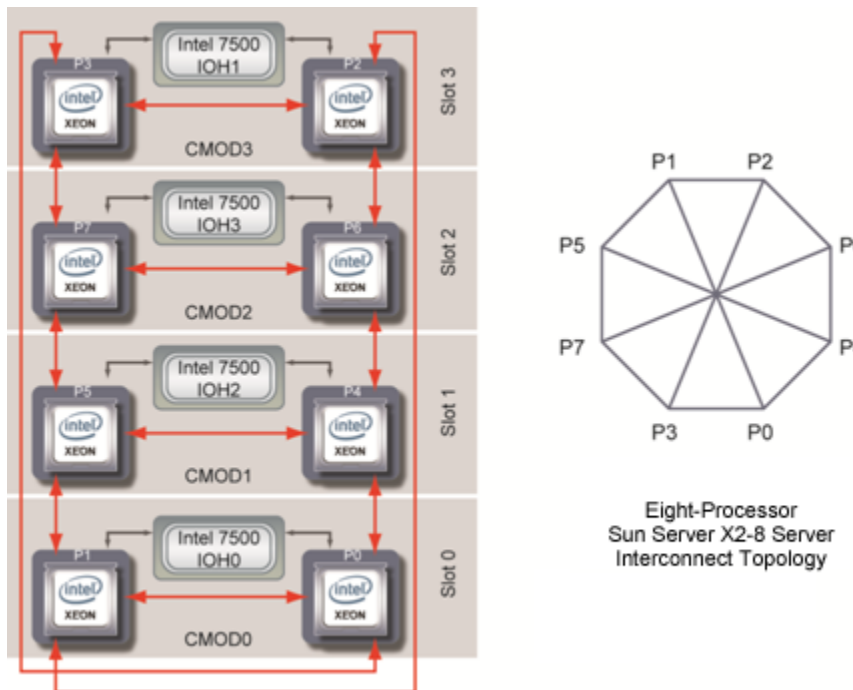
The Sun Server X2-8 system supports four or eight Intel Xeon Processor E7-8800 product family CPUs interfacing to each other over Intel QPI links. Four-socket systems offer two processor-to-processor QPI connections per CPU and one processor-to-IOH QPI link per CPU. As shown in Figure 11, each processor in a four-socket system provides a QPI link to the other processor on the local CPU module. In addition, a second QPI link is established between each CPU and the CPU that is seated in the opposite socket position on the remote CPU module.



**Figure 11.** The four-socket version of the Sun Server X2-8 system provides redundant QPI links between processors.

The QPI connections between processors in eight-socket systems form a spoke-and-wheel formation, as shown in Figure 12. Each processor within a fully populated system links to three other processors:

- The processor on the same CPU module
- A processor on the adjacent CPU module
- One additional processor on an alternate CPU module



**Figure 12.** The eight-socket version of the Sun Server X2-8 system offers redundancy through a spoke-and-wheel connectivity pattern.

In both the four-processor and eight-processor configurations, the QPI links that exist between processors allow the CPUs to access data from the local and remote memory and IOH resources. This design is considered a Non-Uniform Memory Access (NUMA)-style memory architecture, since each processor in these multisoeketed systems can access local memory (connected to the local memory controller) as well as remote memory that is connected to another processor.

The Intel Xeon 7500 IOHs essentially serve as bridges, providing 36 PCI Express (PCIe) lanes capable of operating at PIC-SIG 2.06-compliant Gen2 rates (5 Gb/sec.) In addition, the IOH on CMOD0 provides a set of dedicated x4 PCI Express links (ESI) to the Intel 82801JB (ICH10) Controller Hub. The Intel ICH10 chip provides bus support for connectivity to the ILOM service processor and the internal USB port.

### Intel Xeon Processor E7-8800 Product Family

The Intel Xeon Processor E7-8800 product family CPU illustrated in Figure 13 features the following:

- Eight or ten cores per die
- Up to 30 MB Level-3 cache
- Two Integrated Memory Controllers (IMCs) with two Intel Scalable Memory Interfaces (SMIs) each
- Four full-width, bidirectional Intel QPI buses

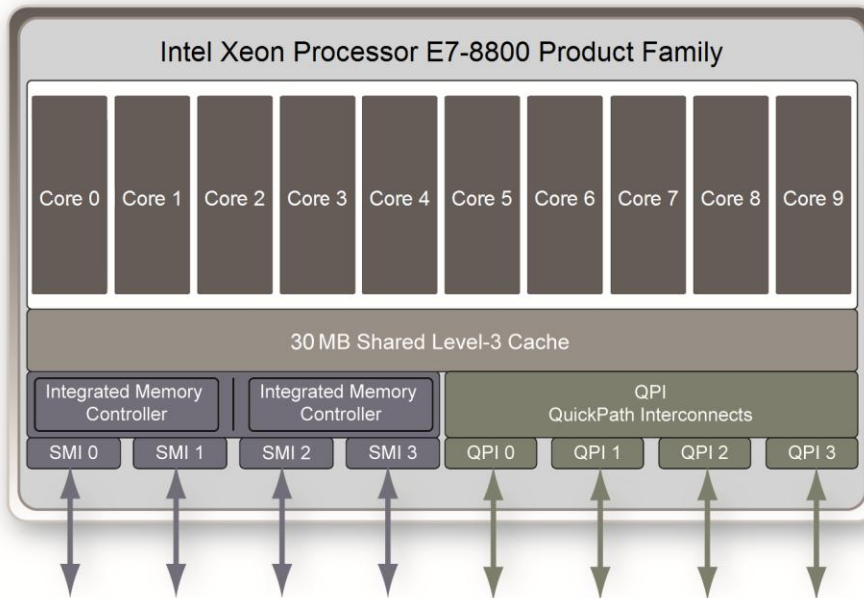


Figure 13. The Intel Xeon Processor E7-8800 product family CPU supports high performance for virtualized workloads.

Within the Sun Server X2-8 system, Intel QPI technology provides high-speed, low-latency point-to-point interconnects between processors and I/O. Because QPI provides a point-to-point interconnect, processors do not contend for a single bus when accessing I/O, and they do not compete for bus bandwidth, which enhances scalability. Each QPI port includes two unidirectional links that support up to 6.4 GigaTransfer/second (GT/sec) per link, offering up to 12.8-GB/sec bandwidth in each direction for a total bandwidth of 25.6-GB/sec per port—significantly higher than previous bus designs.

Additional features of the Intel Xeon Processor E7-8800 product family that help drive optimal performance for bandwidth-intensive, threaded applications include the following:

- **Intel HyperThreading (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 and thermal headroom. At times when a system is operating below maximum power and thermal ratings, Turbo Boost pushes the processors to operate cores at a higher frequency than normal. The frequency increase realized is up to three speed bins depending on the CPU type and the headroom available. Table 7 lists the number of additional speed bins possible based on the CPU type and number of active cores.

TABLE 7. TURBO BOOST TECHNOLOGY CAPABILITIES BY CPU TYPE AND NUMBER OF ACTIVE CORES

NUMBER OF ACTIVE CORES	INTEL XEON PROCESSOR E7-8870 AND INTEL XEON PROCESSOR E7-8860	INTEL XEON PROCESSOR E7-8830
9 or 10	1 step (+133 MHz)	Not Applicable (8-core CPU)
7 or 8	1 step (+133 MHz)	1 step (+133 MHz)
5 or 6	2 steps (+266 MHz)	1 step (+133 MHz)
3 or 4	3 steps (+400 MHz)	1 step (+133 MHz)
1 or 2	3 steps (+400 MHz)	2 steps (+266 MHz)

Turbo Boost is enabled and disabled in the BIOS and is controlled by the operating system. The amount of time the system spends in Turbo Boost depends on workload, operating environment, and system design. More information on Turbo Boost is available at

<http://www.intel.com/technology/turboboost>.

- **Intel Intelligent Power Technology.** When a processor workload decreases, unneeded components, such as cores, cache, and memory, are put into sleep mode to reduce power consumption.
- **Intel Hardware-Assisted Virtualization.** This technology improves virtualization efficiency, and enhances connectivity within a virtualized server. Intel Hardware-Assisted Virtualization can reduce virtual machine monitor (VMM) interventions by eliminating the need for the VMM to listen, trap, and execute certain instructions on behalf of the guest operating system as is required in software-only virtualization. The technology also provides hardware support for transferring platform control between the VMM and guest operating systems. As a result, handoffs are faster, more reliable, and more secure when VMM intervention is required.

### Intel 7500 Scalable Memory Buffer

Each Intel Xeon Processor E7-8800 product family CPU provides two integrated memory controllers. By default, memory is interleaved between the two memory controllers belonging to each processor. A pair of SMI links connects each integrated memory controller to the memory subsystem. To increase reliability, the SMI links that originate from the same memory controller operate in lock-step fashion to access memory DIMMs. Scalable Memory Buffers (MBs) control SMI link access to the memory DIMMs. Each MB connects to one SMI link and up to four DIMMs using two DDR3 channels. Each DDR3 channel supports up to two x4 and x8 single-rank, dual-rank, or quad-rank DIMMs and offers error detection and correction features. Figure 14 provides an illustration of the memory architecture utilized by Sun Server X2-8 systems.

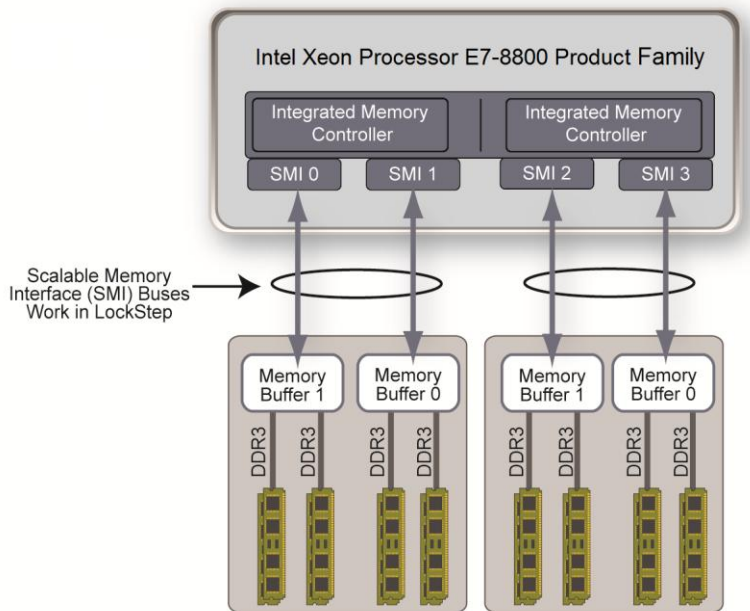


Figure 14. The Intel Xeon 7500 Processor Series CPU includes two integrated memory controllers and two pairs of scalable memory interface buses.

The memory controller supports several advanced RAS features, including the following:

- **x4 and x8 SDDC**—As an advanced form of ECC technology, x4 and x8 SDDC protects computer memory systems from any single memory chip failure. This technology can detect and correct 1-bit to 4-bit internal data and data pin failures within one DDR memory device, and detect up to 8-bit internal data and data pin failures within two DDR memory devices. SDDC performs this function by scattering the bits of an ECC word across multiple memory chips, such that the failure of any one memory chip affects only one ECC bit.
- **x4 DDDC**—Similar to the SDDC, DDDC increases reliability in the case of a hard failure of any two x4 DRAM devices.
- **Demand and patrol scrubbing**—This technology proactively searches system memory, repairing correctable errors. In the case of uncorrectable errors, the algorithm permanently marks the memory location as unreadable.
- **Data Protection**—The memory controller performs replays on errors to recover from transient errors and it supports lane failover and spare lanes to recover from single SMI channel lane failures. Along with each 64 bytes of cache line stored in memory, there are 16 bits for CRC.



Oracle's Sun Server X2-8 System Architecture  
July 2012

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 © 2012, 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. 0310

**SOFTWARE. HARDWARE. COMPLETE.**