An Oracle Technical White Paper
September 2013

Oracle's Sun Blade 6000 I/O and Management Architecture

Introduction 2
Flexible and Expandable Sun Blade 6000 Modular System 2
Focus of This Paper 2
Complete Separation Between CPU and I/O 3
Leading I/O Throughput 3
PCIe ExpressModules 4
Network Express Modules (NEMs) 6
Fabric Expansion Modules (FEMs) 6
Sun Blade 6000 GbE 10-Port Pass-Through NEM 8
Sun Blade 6000 Virtualized 40 GbE Networking Express Module 8
Sun Blade 6000 Ethernet Switched NEM 24p 10 GbE 10
Extending 10 GbE Server Networks Beyond a Single Shelf 12
Storage Expansion 13
Server Module Hard Disk Drives 14
Innovative Enterprise Flash Technology 14
Solid-State Drives (SSDs) 15
Oracle's Sun Storage 6 Gb SAS REM HBA 15
Transparent and Open Chassis and System Management 15
Chassis Monitoring Module 15
Oracle Enterprise Manager Ops Center 17
Conclusion 17
For More Information 17

Introduction

Blade systems play a key role in allowing organizations to derive maximum benefit from their infrastructure, even as their needs change. Many of today's blade systems, however, have limitations that restrict customer choice and flexibility in I/O, networking, or management. Oracle's Sun Blade 6000 modular system is designed to work with open and multivendor industry standards without dictating components, topologies, or management scenarios. This enables organizations to acquire and configure systems that closely match their application environments while avoiding vendor lock-in.

Flexible and Expandable Sun Blade 6000 Modular System

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

- Oracle's SPARC T5-1B and SPARC T4-1B server modules, and Sun Blade X4-2B and X3-2B server modules in any combination up to 10 modules per chassis
- Blade-dedicated PCIe ExpressModules (EMs), which support industry-standard PCI Express interfaces and provide each blade with its own unique I/O configuration (just like rack servers)
- Network express modules (NEMs), which provide access and an aggregated "chassis-wide" common I/O interface to all of the server modules in the Sun Blade 6000 chassis
- Integral chassis monitoring module for transparent management access to server 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. The Sun Blade 6000 modular system also gives organizations a choice of operating systems and virtualization environments. They are optimized to run Oracle Solaris, Oracle Linux, 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 the I/O and management components of Oracle's Sun Blade 6000 modular system. It covers PCI Express ExpressModules (EMs), network express modules (NEMs), and the chassis monitoring module.

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

- Oracle's Sun Blade 6000 Modular System—An overview of the Sun Blade 6000 modular system's architecture
- Oracle's Sun Blade 6000 Server Module Architecture—An in-depth review of the architecture of the SPARC T5-1B, SPARC T4-1B, Sun Blade X4-2B and Sun Blade X3-2B server modules.

Complete Separation Between CPU and I/O

The Sun Blade 6000 modular system design avoids compromises because it provides complete separation between server modules and I/O modules. The following types of I/O modules are supported:

- Up to two industry-standard PCIe EMs can be dedicated to each individual server module.
- Up to two NEMs provide bulk or consolidated "chassis-wide" common I/O for all of the server modules installed in the system.

Through this flexible approach, each server module can be configured with different I/O options depending on the applications hosted. I/O modules are hot-plug and hot-swap capable, and organizations can choose from Sun-branded or third-party adapters for networking, storage, clustering, and other I/O requirements.

Leading I/O Throughput

Sun Blade 6000 server modules provide extensive I/O capabilities and a wealth of I/O options, allowing modular servers to be used for applications that require significant I/O throughput:

- Up to 302 Gb/sec of I/O throughput is provided per server module (when using the latest blade server modules). I/O is delivered through 32 lanes of PCIe 2.0 I/O and PCIe 3.0*****, as well as multiple Gigabit Ethernet (1 GbE) and SAS-2 links. Each server module delivers its I/O to the passive midplane and the I/O devices connected to it in the rear of the Sun Blade 6000 chassis.

  ***** Cards that are categorized as PCIe 3.0 are qualified to run with the SPARC T5-1B; however, due to midplane restriction, will only run at 2.0 speeds. Between switch and CPU, the blade runs at PCIe 3.0. Due to midplane connector capabilities, the speed from switch to the EM/NEM is limited to PCIe 2.0.

- 2.5-inch SAS-2 disk drives and solid state drives (SSDs) are supported on some server modules, while others provide Oracle's Sun Flash Module or USB support.
- Two hot-plug PCIe (EM slots are dedicated to each server module; 20 per Sun Blade 6000 chassis for granular (per blade) I/O configuration support.
- NEMs provide bulk, common I/O across multiple server modules and aggregate I/O functions. Up to two NEMs are supported per Sun Blade 6000 chassis. Each NEM slot features a PCIe x8 or high bandwidth XAUI connection, a Gigabit Ethernet connection, and two SAS-2 link connections to each server module.

Table 1 shows an overview of the I/O modules and total bandwidth that can be configured on the available Sun Blade 6000 server modules.

TABLE 1. MIDPLANE THROUGHPUT FOR SUN BLADE 6000 SERVER MODULES

LINKS	PCI EXPRESS	PCI EXPRESS LINKS	GIGABIT	OTHER LINKS	TOTAL BANDWIDTH

	LINKS TO EMS	TO NEMS	ETHERNET LINKS		(Gb/SEC)
SPARC T5-1B Server Module	2 x 8 PCIe 3.0 links @ 64 Gb/sec	2x 8 PCIe 3.0 links @ 64 Gb/sec	2 @ 1 Gb/sec		258 Gb/sec
SPARC T4-1B Server Module	2 x 8 PCIe 2.0 links @ 64 Gb/sec	2x 8 PCIe 2.0 links @ 64 Gb/sec	2 @ 1 Gb/sec	4 SAS-2 links @ 6 Gb/sec	282Gb/sec
Sun Blade X4-2B Server Module	2 x8 PCIe 2.0 links, @ 64 Gb/sec	2 x8 PCIe 2.0 links, @ 64 Gb/sec	2 @ 1 Gb/sec	4 SAS-2 links, @ 6 Gb/sec	282 Gb/sec
Sun Blade X3-2B Server Module	2 x8 PCIe 2.0 links, @ 64 Gb/sec	2 x8 PCIe 2.0 links, @ 64 Gb/sec	2 @ 1 Gb/sec	4 SAS-2 links, @ 6 Gb/sec	282 Gb/sec

PCIe ExpressModules

Industry-standard I/O, long a staple of rack-mounted and vertically scalable servers, has been elusive in legacy blade platforms. Unfortunately, the lack of industry-standard I/O has meant that customers often paid more for fewer (proprietary) options, and were ultimately limited by a single vendor's innovation. Unlike legacy blade platforms, the Sun Blade 6000 modular system utilizes PCIe EMs, which are a PCI SIG industry-standard form factor. This approach allows a wealth of expansion module options from multiple expansion module vendors and avoids a single-vendor lock on innovation.

The passive midplane implements connectivity between the EMs and the server modules, and physically assigns pairs of EMs to individual server modules. As shown in Figure 1, EMs 0 and 1 (from right to left) are connected to server module 0, EMs 2 and 3 are connected to server module 1, EMs 4 and 5 are connected to server module 3, and so on. Each EM is supplied with an x8 PCIe 2.0 or 3.0 link back to its associated server module, providing up to 64 Gb/sec of I/O throughput. EMs are hot-plug capable according to the standard defined by the PCI SIG and fully customer replaceable without opening either the chassis or removing the server module.

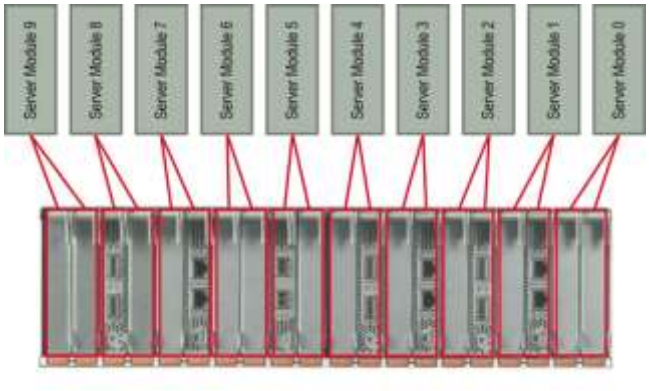

Figure 1. A pair of 8-lane (x8) PCI Express slots allow up to two PCI Express ExpressModules per server module in the Sun Blade 6000 chassis.

With the industry-standard PCIe ExpressModule form factor (Figure 2), EMs are available for multiple types of connectivity, including:

- Gigabit Ethernet
- Configurable as either 16 Gb FC or 10 GbE FCoE
- Combination dual-port GbE, dual-port Fibre Channel
- 10 Gigabit Ethernet (10 GbE)
- Fibre Channel Over Ethernet (FCoE)
- Fibre Channel
- InfiniBand
- SAS

For a list of the latest available PCIe ExpressModules, please see:
https://wikis.oracle.com/display/SystemsComm/Sun+Blade+Systems+Products

Figure 2. Several PCI Express ExpressModules are available for the Sun Blade 6000 chassis.

Network Express Modules (NEMs)

Providing a wide array of I/O access to individual blade servers has always been challenging. Many legacy blade platforms were restrictive in their available options, and many of the options dictated topology and management choices. As a result, data centers often found legacy blade server platforms difficult to integrate into their existing networks, or were resistant to admitting new switch hardware into their chosen network fabrics.

The Sun Blade 6000 modular system addresses this problem through a specific NEM form factor that provides configurable network I/O for all of the server modules in the system. Connecting to all of the installed server modules through the passive midplane, NEMs represent a space-efficient mechanism for deploying high-density configurable I/O, while providing bulk and consolidated I/O options for the entire chassis.

A selection of NEMs are available for configuration with the Sun Blade 6000 modular system, providing pass-through access to the Gigabit Ethernet interfaces located on the server modules as well as other I/O protocol capabilities.

Fabric Expansion Modules (FEMs)

A variety of FEMs are offered to provide access to interfaces on Sun Blade server modules. In some cases these FEMs simply convey access to interfaces that reside either on the processors or the server modules themselves. FEMs can also provide access to additional functionality such as alternative I/O via 10 GbE networking complementing NEM functionality and connectivity. Through the flexibility of FEMs, different network and I/O fabric interfaces can be exposed to the passive midplane, and on to available Network Express Modules.

As an example of FEM connectivity, the Sun Blade PCIe pass-through FEM inserts directly into the FEM socket on the SPARC T5-1B, SPARC T4-1B, or Sun Blade X4-2B, X3-2B server module (Figure 3). The FEM connects the two x8 PCIe links from the onboard PCIe switches to the passive midplane. Available network express modules can then expose these ports as PCIe interfaces.

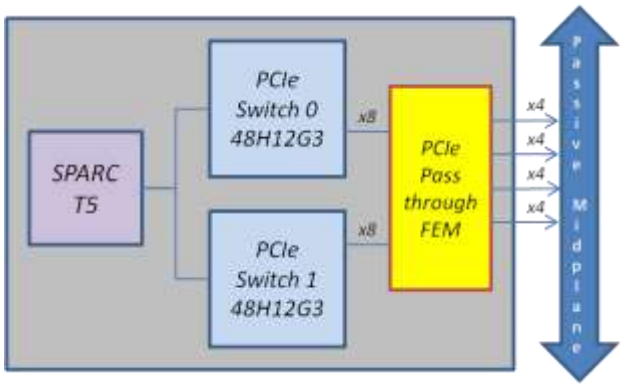

Figure 3. The Sun Blade PCIe pass-through fabric expansion module connects 2 x8 PCIe links to NEMs via the passive midplane.

Network express modules available as of this writing for the Sun Blade 6000 chassis are described in the sections that follow, along with FEMs required for specific server modules. For the latest network express module information, please refer to the following link: Oracle's Blade Systems Wiki.

A variety of NEMs are provided for the Sun Blade 6000 chassis, offering pass-through access to 1 GbE and 10 GbE interfaces on the various server modules. As of this writing, available Sun Blade 6000 NEMs include:

- Sun Blade 6000 GbE 10-port pass-through NEM
- Sun Blade 6000 Ethernet Switched NEM 24p 10 GbE
- Sun Blade 6000 Virtualized 40 GbE Networking Express Module
- To access advanced networking functionality, some NEMS require that fabric expansion modules (FEMs) be installed in the server modules. Part numbers for FEMs for connecting the various server modules to the Sun Blade 6000 NEMs are listed in Table 2.

To access the SAS-2 storage fabric zoning capabilities RAID expansion modules (REMs) may also be required as shown in Table 3. For the latest NEM and FEM information, please refer to: *Oracle's Blade Systems Wiki*.

Sun Blade 6000 chassis RAID expansion module (REM) SAS connectivity using NEMs is shown in Table 3.

SERVER MODULE	SUN BLADE 6000 ETHERNET SWITCHED NEM 24P 10 GbE	SUN BLADE 6000 VIRTUALIZED 40 GbE NEM
SPARC T5-1B	X4871A-Z-N	7100633
SPARC T4-1B	X4871A-Z-N	7100633
Sun Blade X4-2B Server Module	X4871A-Z-N	7100633
Sun Blade X3-2B Server Module	X4871A-Z-N	7100633

SERVER MODULE	SUN BLADE 6000 ETHERNET SWITCHED NEM 24P 10 GbE	SUN BLADE 6000 VIRTUALIZED 40 GbE NEM
SPARC T4-1B	SGX-SAS6-REM-Z	SGX-SAS6-REM-Z
Sun Blade X4-2B	SGX-SAS6-REM-Z	SGX-SAS6-REM-Z
Sun Blade X3-2B	SGX-SAS6-REM-Z	SGX-SAS6-REM-Z

Sun Blade 6000 GbE 10-Port Pass-Through NEM

For applications that don't require 10 GbE networking, the Sun Blade 6000 GbE 10–port pass-through NEM offers an effective solution. A single Sun Blade 6000 GbE 10-port pass-through NEM (Figure 4) provides ten 1 GbE pass-through ports that provide access to one of two 1 GbE interfaces on all server modules. In order to support both GbE links that are integrated in the respective server modules, two of the 10-port pass-through NEMs need to be installed.

Figure 4. The Sun Blade 6000 GbE 10-port pass-through NEM provides ten Gigabit Ethernet pass-through ports.

Sun Blade 6000 Virtualized 40 GbE Networking Express Module

The Sun Blade 6000 Virtualized 40 GbE Networking Express Module delivers virtualized 40 GbE or 10 GbE access with near-zero management. Designed to seamlessly integrate into the Sun Blade 6000 chassis, Oracle's Sun Blade 6000 Virtualized 40 GbE Networking Express Module is the industry's first fully-virtualized 40 GbE network aggregation for all 10 blade server modules in the Sun Blade 6000 chassis. It is an ideal network interface for connecting multiple virtualized blades that require scalable I/O throughput for workloads such as Web servers, application servers, and database servers.

The Sun Blade 6000 Virtualized 40 GbE Networking Express Module is a multifunction PCIe 2.0 connectivity module that combines virtualized 40 GbE or 10 GbE network connectivity and 1 GbE pass through. This NEM is a unique platform for simplifying data center networks without adding the extra cost of switches to manage. It utilizes an Oracle-designed 40 GbE NIC ASIC that virtualizes the 40 GbE or 2x10 GbE network connectivity across 10 server modules in the Sun Blade 6000 chassis, thus simplifying the networks by reducing cables by 10:1—and without adding the extra switching layer. Figure 5 illustrates the hardware architecture of the NEM.

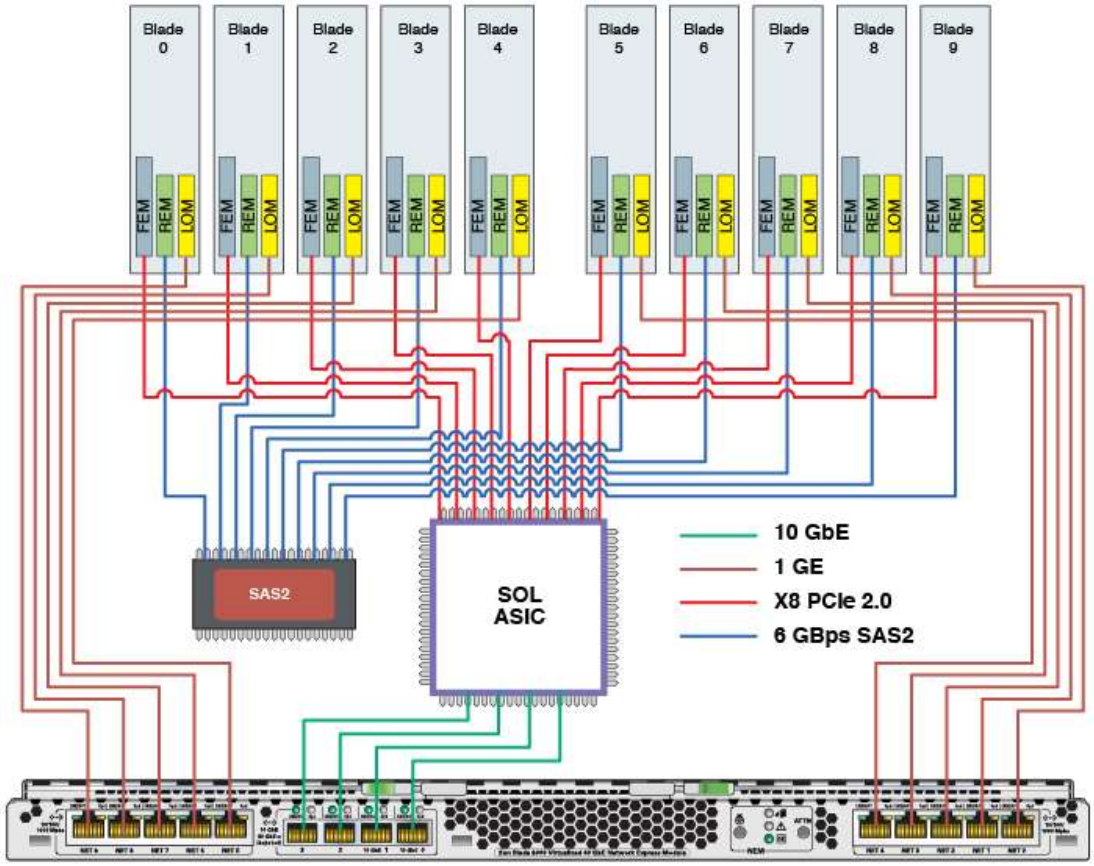

Figure 5. The Sun Blade 6000 Virtualized 40 GbE Networking Express Module provides aggregation and optional redundancy at the chassis level.

The Sun Blade 6000 Virtualized 40 GbE Networking Express Module offers configurable options, making it ideally suited for different types of workloads ranging from applications just needing high intrablade communications to those needing higher virtualized bandwidth. Examples of some of these workloads include those that move smaller amounts of data and require high-speed network for lower latency such as Web servers and application servers, as well as workloads that move a lot of data intermittently such as enterprise applications and databases. By default, the blades share the available bandwidth in equal amount and are in privacy mode so that unicast traffic from one blade is not seen by other blades. More deterministic bandwidth can be configured for any or all blades in increments of 1 percent of the total bandwidth. Ingress filtering and VLANs supported by the NEM provide security options for both the unicast and broadcast traffic from the blades.

The Sun Blade 6000 Virtualized 40 GbE Networking Express Module supports connection to external devices either through small form-factor pluggable (SFP)+ ports, or 10/100/1000 twisted-pair Ethernet (TPE) ports. Up to 2 SFP+ ports can be used as 2x10 GbE. A 40 GbE link is provided when 4 SFP+ ports are used (with QSFP to 4 SFP+ splitter cables).

The Sun Blade 6000 Virtualized 40 GbE Networking Express Module form factor provides a method of deploying bulk remote I/O that allows tool-less installation/removal and packs more performance and functionality in a smaller space, while delivering higher network throughput. It makes efficient use of data center real estate by reducing the number of cables.

This NEM is easy to install and manage. The flexible architecture of the Sun Blade 6000 modular system is based entirely on the hot-pluggable components for I/O, processing, system management, and chassis infrastructure. All critical components, including the NEM modules, are hot swappable and redundant, providing enterprise-class reliability, availability, and serviceability (RAS) features. The Sun Blade 6000 Virtualized 40 GbE Networking Express Module helps to boost data center efficiency and uptime and lower total cost of ownership (TCO).

Pass-through PCIe 2.0 Fabric Expansion Module (Marketing Part # 7100633) is required on the server blade for its connectivity to the SFP+ ports on the NEM.

Sun Blade 6000 Ethernet Switched NEM 24p 10 GbE

Beyond providing aggregation of 10 Gigabit Ethernet, the Sun Blade 6000 Ethernet Switched NEM 24p 10 GbE offers an active non-blocking, low-latency 10 GbE switch compatible with the Sun Blade 6000 chassis (shown in Figure 7). Based on a 24-port 10 GbE switch chip, the NEM is a full-featured Layer 2 / Layer 3 edge switch, allowing dense, non-blocking 10 GbE fabrics to be constructed in combination with suitable rack and enterprise-level switches. Providing a full switch within a standard single-height NEM form factor allows considerable consolidation and savings. A block-level diagram of the NEM aligned with a photo of the back panel is provided in Figure 7.

The switch is designed to support Sun Blade 6000 server modules via 10 GbE connections that attach through the Sun Dual 10 GbE PCIe 2.0 Fabric Expansion Module.

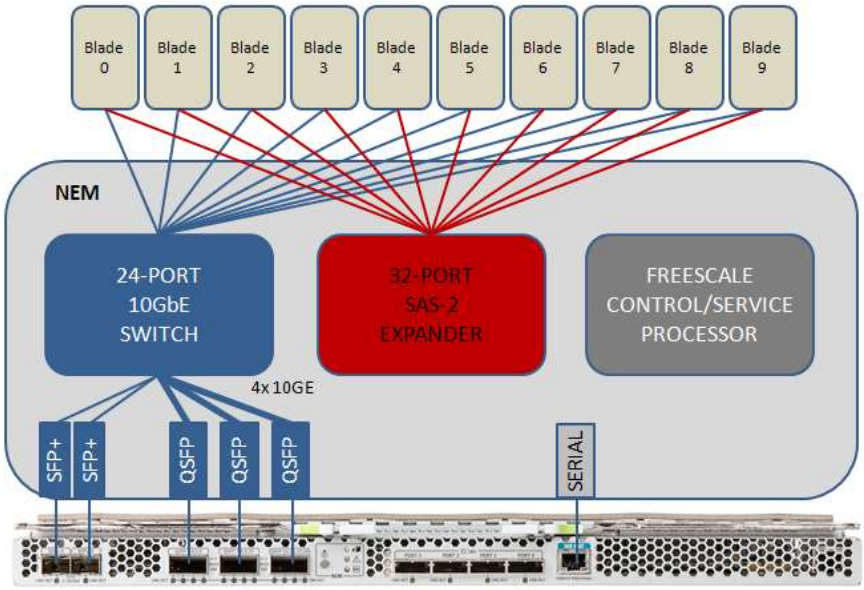

Figure 7. The Sun Blade 6000 Ethernet Switched NEM 24p 10 GbE provides a 24-port 10 GbE switch for non-blocking 10 Gb access to each server module.

Each Sun Blade 6000 Ethernet Switched NEM 24p 10 GbE delivers a 10 GbE connection to each server module installed in the chassis. Redundant 10 GbE connections to each server module can be provided by installing two NEMS into the Sun Blade 6000 chassis.

In order to offer nonblocking throughput, and scale-out connectivity to other chassis or to existing data center networks, each NEM provides a total of 14 external 10 GbE connections through the back panel:

- Two SFP+ connectors
- Three 4x 10 GbE QSFP (quad SFP) connectors

In addition to switch consolidation offered by the NEM, the three 4x 10 GbE QSFP connectors can be split for connectivity to 4 SFP+ or used as an aggregated 40 GbE interconnect, providing considerable cable consolidation when connecting to a compatible rack or enterprise switch. In addition, the two SFP+ ports can be configured to operate at either 1 GbE or 10 GbE for cable aggregation and legacy interconnect. The external SAS-2 ports are not supported as of this writing.

Consistent with Oracle's modular design principals, the NEM is easy to manage, and offers standard interfaces and network protocols, including:

- Unified chassis management
- A Web browser interface and a standard command-line interface (Oracle Integrated Lights Out Manager or Oracle ILOM shell)
- Multiple user privileges
- Single sign on

- Oracle ILOM support via the chassis monitoring module
- Environmental monitoring
- An industry-standard compatible L2 / L3 network stack
- CLI and command set

Several server modules utilize the Sun Dual 10 GbE PCIe 2.0 Fabric Expansion Module to connect to the NEM (see Table 2). The FEM provides dual 10 GbE interfaces to the server module, and is pictured in Figure 8.

Figure 8. The Sun Dual 10 GbE PCIe 2.0 Fabric Expansion Module provides two 10 GbE interfaces.

Extending 10 GbE Server Networks Beyond a Single Shelf

The Sun Blade 6000 Ethernet Switched NEM 24p 10 GE provides 10 GbE connectivity within a single chassis of server modules but can be further linked to other Sun Blade 6000 chassis to significantly increase the size of the server module deployment. Multiple chassis of Sun Blade 6000 Ethernet Switched NEM 24p 10 GbE-equipped server modules can be interconnected using only a few QSFP-to-QSFP cables to provide high bandwidth low latency, high node count configurations of servers. To extend even further, to very large node counts, many 10 GbE-equipped chassis of server modules can be connected using the Sun Network 10 GbE Switch 72p top-of-rack (ToR) 1U switch. The Sun Network 10 GbE Switch 72p ToR switch is equipped with 16 (4x10 GbE) QSFP ports and 8 (1/10 GbE) SFP+ switch ports providing wire rate, low latency, cut through switching on all 72 ports and delivers full, nonblocking bandwidth to each blade or rack server. The Sun ToR has enough density to build a sizable fabric that operates at full bandwidth without the need for spanning tree protocol. The Sun ToR supply of 8 SFP+ uplinks can cost effectively join 64 servers to a tier 2 switch, obviating the need for complex tier 1 interconnectivity. Some simple, example network configurations are illustrated in Figure 9.

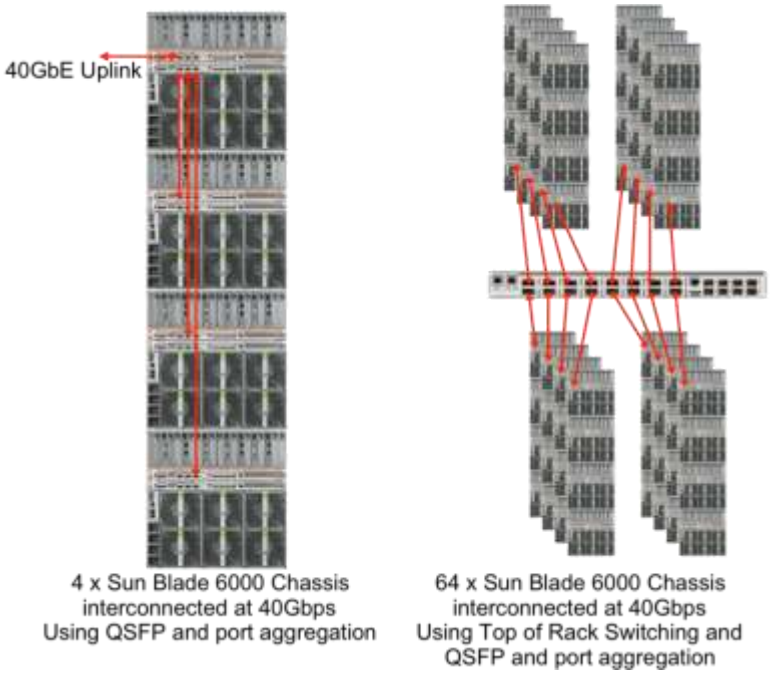

Figure 9. A rack full of Sun Blade 6000 chassis using the 24p NEM switch interconnected using a single QSFP uplink to one NEM, and utilizing a 72p Switch to connect 16 racks of servers from 24p NEMs to a single 72p Switch. Further bandwidth improvement could be achieved through port aggregation configurations.

Storage Expansion

A number of storage expansion options are available for server modules. Individual server modules have their own internal expansion options, including hard disk drives (HDDs) and flash memory devices, including solid-state drives (SSDs and USB. Supported storage options for the various server modules are shown in Table 4. Note that an appropriate REM and NEM combination is required in some circumstances.

TABLE 4. SERVER MODULE STORAGE OPTIONS		
SERVER MODULE ARCHITECTURE	**SERVER MODULE**	**HDDS/SSDS**
SPARC T5	SPARC T5-1B server module	2
SPARC T4	SPARC T4-1B server module	2
Intel Xeon Processor E5-2600 v2 Product Family	Sun Blade X4-2B server module	4
Intel Xeon Processor E5-2600 Product Family	Sun Blade X3-2B server module	4

Server Module Hard Disk Drives

A choice of (optional) hot-swappable 2.5-inch SAS are supported with all server modules.

- Serial Attached SCSI (SAS) drives provide high performance and high density. Drives are 10,000 rpm and available in various capacities. These drives provide enterprise-class reliability with up to 1.6 million hours mean time between failures (MTBF).

Innovative Enterprise Flash Technology

Modern servers are driving throughput levels that can rapidly outpace the capabilities of traditional HDD storage solutions. While many servers can achieve processing capabilities in excess of one million I/O operations per second (IOPS), today's fastest HDDs are only capable of about 300 to 400 IOPS. To match throughput more closely to server performance and to address the challenging demands of data-intensive applications, many data centers implement large pools of high-speed disk drives, often utilizing only the fastest outer portions of their storage media, wasting capacity. In some cases, a large buffer of expensive DRAM is also deployed so that the application's working set can be stored in memory to reduce latency. Flash technology provides a more economical alternative that can dramatically enhance application I/O performance while also operating with significantly better energy efficiency than conventional HDDs. Recent advances in the flash technology have made enterprise levels of reliability, life expectancy, and manageability available from flash-based storage devices, in the form of SSDs. These now become an effective and reliable solution for enterprise storage. Flash technology contains no moving parts, avoiding the seek times and rotational latencies inherent with traditional hard disk drive technology and removing the single largest bottleneck to disk drive performance. Oracle flash technology can offer instant availability of increased performance, power and cost savings, and reliability benefits over rotational media, including:

- 97 percent lower cost per I/O operation per second (IOPS)

- 99 percent lower watts/IOPS
- 66 percent lower watts/GB

Solid-State Drives (SSDs)

Because SSDs offer low latency and are significantly less expensive than an equivalent amount of DRAM storage, they balance cost and performance in a manner that can provide significant value for I/O-intensive workloads. SSDs are offered in conventional disk drive form factor (Figure 10).

Figure 10. **Solid-st**ate drives (left) provide enterprise flash technology in a standard disk drive form factor. **Oracle's Sun** FlashFire technology (right) provides optimized performance in a considerably smaller form factor.

Oracle's Sun Storage 6 Gb SAS REM HBA

The Sun Storage 6 Gb SAS REM HBA is a mezzanine card that can be fitted on the rear of the motherboard for mirroring or striping of the disk storage devices as well as offering access for simple JBOD arrays. The REM currently provides four SAS-2 (6 Gb/sec) links to the four internal disks. Integrated striping can be used in conjunction with integrated mirroring or integrated mirroring enhanced. The RAID configuration is self sufficient and does not require intervention from the host CPU or OS. All standard SPARC T4- server modules support one SAS REM and the card is also supported on Sun Blade X3-2B server modules.

Transparent and Open Chassis and System Management

Management in legacy blade platforms has typically either been lacking, or administrators have been forced into adopting unique blade or platform-specific management infrastructure. To address this issue, the Sun Blade 6000 modular system provides a wide range of flexible management options.

Chassis Monitoring Module

The chassis monitoring module is the primary point of management for all shared chassis components and functions, providing a set of management interfaces. Each server module contains its own service processor, giving it similar remote management capabilities to other Sun servers. Through their respective Oracle ILOM service processors, individual server modules provide IPMI, HTTPs, CLI (SSH), SNMP, and file transfer interfaces that are directly accessible from the Ethernet management

port on the chassis monitoring module. Each server module is assigned an IP address (either manually, or via DHCP) that is used for the management network.

Chassis Monitoring Module Network Functionality

A single chassis monitoring module is provided with each Sun Blade 6000 modular system, and is configured with an individual IP address assigned either statically or dynamically via DHCP. The chassis monitoring module provides complete monitoring and management functionality for the chassis (or shelf) while providing access to server module management functions. In addition, the chassis monitoring module supports HTTP and CLI "pass-through" interfaces that provide transparent access to each server module. The chassis monitoring module also provides access to each server module via a single serial port through which any of the various lights-out management (LOM) interfaces can be configured. The chassis monitoring module's management functions include:

- Implementation of an IPMI satellite controller, making the chassis environmental sensors visible to the server module's BMC functions
- Direct environmental and inventory management via CLI and IPMI interfaces
- Chassis monitoring module, Oracle ILOM, and NEM firmware management
- Pass-through management of blades using IPMI, SNMP, and HTTP links along with command-line interface (CLI) SSH contexts

The management network internal to the chassis monitoring module joins the local management processor on each server module to the external management network through the passive midplane.

Chassis Monitoring Module Architecture

A portion of the chassis monitoring module functions as an unmanaged switch dedicated exclusively to remote management network traffic, letting administrators access the remote management functions of the server modules. The switch in the chassis monitoring module provides a single network interface to each of the server modules and to each of the NEMs, as well as to the service processor located on the chassis monitoring module itself. Figure 11 provides an illustration and a block-level diagram of the Sun Blade 6000 chassis monitoring module.

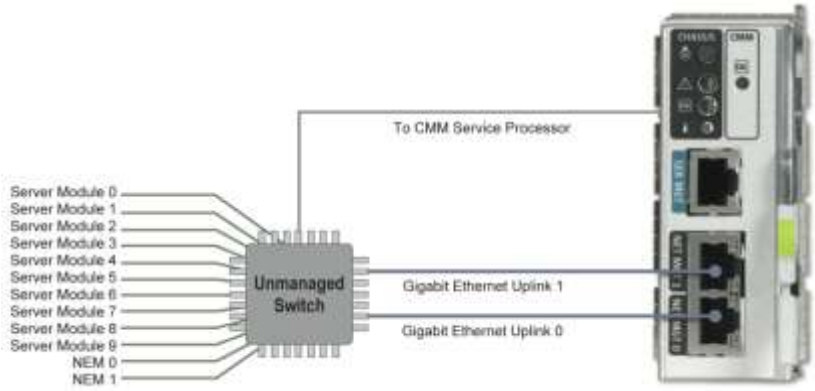

Figure 11. The chassis monitoring module provides a management network that connects to each server module, the two NEMS, and the chassis monitoring module itself.

The chassis monitoring module's functionality provides various management functions, including power control of the chassis as well as hot-plug operations of infrastructure components such as power supply modules, fan modules, server modules, and NEMs. The chassis monitoring module acts as a conduit to server module LOM configuration, allowing settings such as network addresses and administrative users to be configured or viewed.

Oracle Enterprise Manager Ops Center

Beyond local and remote management capabilities, data center infrastructure needs to be agile and flexible, allowing not only fast deployment, but also streamlined redeployment of resources as required. Oracle Enterprise Manager Ops Center technology provides an IT infrastructure management platform for integrating and automating management of thousands of heterogeneous systems. To improve lifecycle and change management, Oracle Enterprise Manager Ops Center supports the management of applications and the physical and virtual servers on which they run, including the Sun Blade 6000 modular system.

Further information about Oracle Enterprise Manager Ops Center is provided in the Sun Blade 6000 modular system architecture overview white paper and is available on the Web at http://www.oracle.com/us/products/servers-storage/servers/blades/sun-blade-6000-chassis/resources/index.html .

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.

Oracle's standard- and open-systems-based approach yields choice and avoids compromise—providing a platform that benefits from widespread industry innovation. Industry-standard I/O and multiple storage options provide flexibility and leading throughput for individual server modules. Transparent networking and management means that the Sun Blade 6000 modular system fits easily into an existing network and management infrastructure.

For More Information

For more information on Oracle's Sun Blade 6000 modular system, 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 Oracle Technology Network 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/technetwork/oem/ops-center/index.html
- Oracle Solaris operating system*:* http://www.oracle.com/technetwork/server-storage/solaris/overview/

Oracle's Sun Blade 6000 I/O and Management Architecture
March 2013

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

Hardware and Software, Engineered to Work Together