An Oracle White Paper
April 2012

Oracle's Sun Blade 6000 Modular Systems

ORACLE

Executive Overview

Growth in devices connected to the network is driving new IT demands for agility and scalability. Organizations must be able to respond quickly to unpredictable needs for capacity — adding compute power or growing services on demand. At the same time, most data centers are rapidly running out of space, power and cooling while energy costs continue to rise. Rapid growth must be met with a consolidated and converged infrastructure, controlled and predictable costs, and efficient management practices. Simply adding more low-density power-consumptive servers only exacerbates the problem.

Blade server architecture offers considerable promise toward addressing these issues through increased compute density, improved serviceability, and reduced levels of complexity. Unfortunately, most legacy blade platforms don't provide the necessary flexibility required by applications such as Oracle E-business Suite, Oracle Fusion Middleware, and the Oracle database itself. Complicating matters, many legacy blade server platforms lock customers into a proprietary and vendor-specific infrastructure that often requires redesign of existing network, management, and storage environments. These legacy chassis designs also often artificially constrain expansion capabilities. As a result, traditional blade architectures have been largely restricted to low-end Web and IT services.

Responding to these challenges, Oracle's Sun Blade 6000 modular systems provide open modular architectures that deliver on the promises of blades without compromising features or functions. The Sun Blade 6000 modular systems present a comprehensive multitier blade portfolio that lets organizations deploy the broadest range of applications on the most ideal platforms. Oracle blades offer a choice of server modules based on the latest Oracle SPARC and Intel Xeon processors and are optimized to run Oracle Solaris, Oracle Linux, and Oracle VM. They are also certified to run Red Hat Enterprise Linux, SUSE Linux Enterprise Server, Windows Server, and VMware. Oracle's open systems approach enables organizations to select the platforms that best satisfy their requirements. The result is a modular IT infrastructure architecture that serves the needs of the data center and the goals of the business while protecting existing investments into the future.

This document provides an overview of the Sun Blade 6000 modular systems architecture. It is complemented by two additional white papers that provide additional component architecture details:

- Sun Blade 6000 Server Module Architecture white paper
- Sun Blade 6000 I/O and Management Architecture white paper

An Open Systems Approach to Modular Architecture

Organizations operating traditional IT infrastructure, business processing, and back office applications are always looking for ways to cut costs and safely consolidate infrastructure. For many, large numbers of older and less efficient systems constrain the ability to grow and adapt, both physically and computationally. Oracle E-business Suite, Oracle Fusion Middleware, and the Oracle database all require computational performance, density, and scalability. Enterprise environments too need dense and capable platforms as they grapple with growing needs for consolidation through virtualization for systems ranging from databases to common applications such as file, print, and mail servers. With most data centers constrained by space, heat, or power issues, the challenge of scaling within fixed physical, thermal, and electrical boundaries is now very real. Successful solutions must be power- and cooling-efficient, cost-effective, and reliable — with investment protection factored into fundamental design considerations.

Fortunately, new technology is yielding opportunities for increased efficiency and flexibility in the data center. Multicore processor technologies are continually increasing compute density. Virtualization technologies and more powerful servers are making it possible to consolidate widely distributed data centers using smaller numbers of more powerful servers. Standard high-bandwidth networking and interconnect technologies are becoming more affordable. Modern provisioning technology makes it possible to dynamically deploy and readjust workloads on the fly. Lastly, advances in data center cooling designs are providing new efficiencies that can dramatically reduce energy costs.

Regrettably, most current server form-factors have failed to take full advantage of these trends. For instance, most traditional rackmount servers require a box swap in order to allow an organization to deploy new CPU and I/O technologies. Modular architectures offer the opportunity to rapidly harvest the returns of new technology advances, while serving the constantly evolving needs of the enterprise.

The Promise of Blade Architecture

At its best, modular or blade server architectures blend the enterprise availability and management features of vertically scalable platforms with the scalability and economic advantages of horizontally scalable systems. In general, modular architectures offer considerable promise, and can contribute to:

- *Higher compute density* — Providing more processing power per rack unit (RU) than with rackmount systems.
- *Increased serviceability, availability and power efficiency* — Featuring shared common system components such as power, cooling, and I/O interconnects.
- *Reduced complexity* — Through fewer required components, cable and component aggregation, and consolidated management.
- *Faster service expansion and bulk deployment* — Letting organizations expand or scale existing services and flexibly pre-provision chassis and I/O components.
- *Lowered costs* —Modular servers can be less expensive to acquire, service, and manage, and are also faster to deploy as well as more power-efficient and more reliable.

While some organizations adopted first-generation blade technology for Web servers or simple low-end IT infrastructure, many legacy blade platforms have not been able to deliver results for a broader set of applications. Part of the problem is that most legacy blade systems are based on proprietary mechanical and electrical architectures that lock adopters into an extensive infrastructure that constrains deployment. In addition, though vendors typically try to price server modules economically, they often charge a premium for required proprietary I/O and switching infrastructure. Availability of suitable computational platforms has also been problematic.

Together, these constraints caused compromises in terms of features, functionality, and performance that had to be weighed when considering blade technology for individual applications:

- Power and cooling limitations often meant that processor choices were limited to less powerful versions.
- Limited processing power, memory capacity, and I/O bandwidth severely constrained the applications that could be deployed on blade server platforms.
- Proprietary tie-ins and other constraints in chassis design dictated networking topology, and limited I/O expansion possibilities to a small number of proprietary modules.

These compromises in chassis design were largely the result of a primary focus on density — with smaller chassis requiring reduced-footprint form factors. Ultimately these designs limited the broad application of blade technology.

Oracle's Sun Blade 6000 Modular Systems

To address the compromises of earlier blade platforms, Oracle started with a "clean sheet" design point focused on the needs of the data center, rather than with preconceptions of chassis design. With this innovative and truly modular approach and a no-compromise feature set, Oracle's Sun Blade 6000 modular systems offer considerable advantages for a wide range of applications. Organizations gain the promised benefits of blades, and can save more by deploying a broader range of their applications on the Sun Blade 6000 modular system.

- **Scalable, Expandable, and Serviceable Multitier Architecture**. Sun Blade 6000 modular systems let organizations deploy multitier applications on a single unified modular architecture. These systems support Intel Xeon and SPARC T-Series processors. The Oracle VM virtualization platform and the Oracle Solaris and Oracle Linux operating systems are supported uniformly on all x86 platforms. Oracle's x86 platforms are also certified to run Red Hat Enterprise Linux, SUSE Linux Enterprise Server, VMware and Windows Server.

  By offering the latest Intel and SPARC T-Series processors, large memory capacity, high I/O bandwidth, and integrated storage, these systems support a very broad range of applications. In addition, the Sun Blade 6000 modular systems achieve better power efficiency than rackmount servers by consolidating power and cooling infrastructure for multiple systems into highly efficient modular system chassis. The result is a high-performance IT infrastructure that packs more capacity and functionality into a smaller package — in terms of both real estate as well as power envelope.

With innovative chassis design, Sun Blade modular systems allow organizations to take full advantage of future technology without "forklift upgrades". All major components are hot pluggable and hot swappable, including I/O modules. Compute, I/O, storage, power, cooling and management modules can all be independently serviced, upgraded, and expanded.

- **Sun Blade Transparent Management.** Many, if not most, blade vendors sell proprietary blade management platforms specifically for their blade infrastructure. With the Sun Blade 6000 modular systems, customers have the choice of using their existing management tools, and just as importantly, existing practices and approaches. Each individual blade has its own Oracle Integrated Lights Out Manager (Oracle ILOM) Service Processor, giving operators and administrators full control at a blade level. Individual blades can be completely managed, upgraded, configured, and monitored without impact to other blades within the chassis. In fact, there is no difference in management approach compared to a rackmount server. In addition, each Sun Blade 6000 modular system chassis provides its own dedicated chassis monitoring module (CMM), enabling administrators to manage, configure, and monitor at a chassis level. Oracle ILOM has strong support for major standards (e.g. IPMI, WS-MAN, ssh to CLI, HTTPS), which enables very flexible integration and allows administrators to use their favorite management tools from Oracle (such as Oracle Enterprise Manager Ops Center) or from third parties.
- **Open and Independent Industry-Standard I/O.** The Sun Blade 6000 modular systems provide a cable-once architecture with complete hardware isolation of compute and I/O modules. Sun also supports true industry-standard I/O on its modular system platforms. Sun Blade modular systems utilize standard PCI Express (PCIe) I/O architecture and adapters — the same technology that dominates the rackmount server industry — allowing organizations to pick and choose I/O options on a per-blade basis rather than being forced into chassis-wide only options. I/O adapters from multiple vendors are available to work with Sun Blade modular systems. A truly modular design based on industry-standard, hot-pluggable I/O means that systems are easier to install, upgrade, and service — providing simpler administration, higher reliability and up-time opportunities as well as better compatibility with existing network and storage environments. For instance, replacing an I/O module in a Sun Blade modular system can be performed in as little as 30 seconds, and with no downtime.
- **Standard and Expandable Storage Options.** Sun Blade 6000 modular systems include a range of flexible and expandable storage options. Many server modules contain hot-pluggable and hot-swappable drive slots. Enterprise flash technology in the form of solid-state drives (SSDs) can be employed instead of hard disk drives (HDDs). Other server modules feature Sun Flash Modules and USB slots to facilitate additional storage requirements. The Sun Blade Storage Module M2 provides considerable "in-chassis" storage expansion and flexibility beyond the individual server modules.
- **Highly Efficient Cooling.** Traditional blade platforms have a reputation for being hot and unreliable — a reputation earned by systems with insufficient cooling and chassis airflow. Not only do higher temperatures negatively impact electronic reliability, but also hot and inefficient systems require more data center cooling infrastructure, with its associated footprint and power draw. In response, the Sun Blade 6000 modular systems provide optimized cooling and airflow that facilitates reliable system operation and efficient data center cooling. In fact, Sun Blade modular systems

deliver similar cooling and airflow capacity as do Sun rackmount systems — for both SPARC and x86 server modules. This results in reliable system operation and requires less cooling infrastructure even while accommodating the high performance processors and memory footprints that are found in equivalent rackmount servers.

Open and Modular System Architecture

The Sun Blade 6000 modular systems provide a new approach to modular system architecture. This approach combines careful long-term chassis design with open and industry-standard systems architecture.

Innovative Industry-Standard Design

Providing choice in modular system platforms is essential, both to help enable the broadest set of applications, and to provide the best investment protection for a range of different organizations and their requirements. Sun Blade 6000 modular systems offer flexible choice and key innovations for modular computing.

A Choice of Operating Systems and Virtualization

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, and SUSE
- Microsoft Windows Server

Virtualization software support includes:

- Oracle VM
- VMware

A Choice of Intel Xeon and Oracle SPARC Processor-Based Blades

Legacy blade platforms were often restrictive in the processor architectures they supported, limiting innovation for modular systems and forcing difficult architectural choices for adopters. In contrast, Sun Blade 6000 modular systems offer a choice of server modules based on Intel Xeon or Oracle's SPARC processors. In addition, Sun Blade 6000 server modules provide large memory capacities, while the individual chassis provide significant power and cooling capacity.

The Sun Blade 6000 server module family is described below.

SPARC T4-1B Server Module

The high density SPARC T4-1B server module is based on the latest generation SPARC T4 processor, which offers a 5X single-thread performance improvement over the previous generation T-Series processors in addition to high performance for multithreaded application workloads. Designed for a

broad range of requirements, the SPARC T4-1B server module now provides single-thread performance to excel at database, batch, and OLTP workloads while also delivering high multithreaded performance for workloads such as middleware applications.

The SPARC T4-1B server module is configured with one SPARC T4 processor (8 cores, 32 threads, 2.85 GHz), 16 DIMM slots for up to 256 GB memory, and two drive slots for hot-pluggable 2.5 inch drives. The SPARC T4-1B server module provides multiple network interfaces, including two 10/100/1000 Base-T Ethernet ports, one dedicated 10/100 Base-T Ethernet port, and two optional 10Gb XAUI Ethernet ports through fabric expansion modules.

Up to 10 SPARC T4-1B server modules can be configured in a single Sun Blade 6000 chassis, providing the highest-density, lowest-cost SPARC T4 solution and offering high reliability and integrated networking in a compact 10U form factor. Both x86 and SPARC blade server modules can be configured together in the Sun Blade 6000 chassis, enabling support for multiple workloads with unique processor performance profiles.

SPARC T3-1B Server Module

The highest density SPARC T3-1B server module is based on the previous generation SPARC T3 chip multithreading technology (CMT) processor, delivering massive throughput and scalability optimized for multithreaded enterprise application workloads. The SPARC T3 processor architecture is uniquely designed with on-chip dual 10 Gigabit Ethernet (10GbE) for high-speed networking, and on-chip cryptographic acceleration for data center security.

The SPARC T3-1B offers a high degree of scalability, density and deployment flexibility. The SPARC T3-1B is available in two base server options, an eight core (64 thread) or 16 core (128 thread) single SPARC T3 1.65 GHz processor. The SPARC T3-1B contains 16 DIMM slots supporting up to 128 GB of memory using DDR3 DIMMs, providing significant memory bandwidth for multithreaded workloads. The T3-1B provides network interfaces including two 10/100/1000 Base-T Ethernet ports, one dedicated 10/100 Base-T Ethernet port, and two optional 10 Gb XAUI Ethernet ports through fabric expansion modules.

Sun Blade X6270 M2, X6270 M3, and X6275 M2 Server Modules

Sun Blade X6270 M2 server modules are ideal for virtualized business applications and enterprise collaboration workloads that require a high-performance x86 server module. Each Sun Blade X6270 M2 server module features two sockets for four or six core Intel Xeon Processor 5600 Series (Westmere-EP) CPUs. The processor sockets connect to an Intel 5520 chipset using the Intel QuickPath Architecture. Up to 18 DDR3 (low voltage) DIMMs are supported.

The Sun Blade 6270 M3 server module is a follow on to the Sun Fire X6270 M2 server module, implementing the newest Intel architecture. Each Sun Blade X6270 M3 server module features two sockets for four-, six-, or eight-core Intel Xeon Processor E5-2600 product family CPUs, and 24 DDR3 (low voltage) DIMMs.

The Sun Blade X6275 M2 server module, available in two models: 10GbE or standard 1GbE provides new levels of density by featuring two separate compute nodes within a single server module form

factor. Ideal for environments where high-performance and computational density are paramount, each of the compute nodes support two sockets for Intel Xeon Processor 5600 Series CPUs, and 12 DDR3 (low voltage) DIMMs. Each compute node is provided with a connection to a 10GbE port or 1GbE port (model dependent). Additional I/O expansion is available via a PCIe ExpressModule (one per compute node).

Designed for Balance and Versatility

Sun Blade server module computational power is balanced by significant I/O capacity, with considerable bandwidth delivered from server modules to the multiple available I/O expansion modules (a total of up to 302 Gb/sec on the latest server modules). To enhance availability, server modules have no power supplies or fans, but rely on redundant, large and efficient chassis-level power and cooling subsystems. Organizations can deploy individual server modules with their own processor, I/O, and operating system configurations that best serve the requirements of their applications or workloads. Different server modules (SPARC and x86) can co-reside within a single chassis, and be deployed and redeployed as needs dictate.

Complete Separation Between CPU, I/O, and Storage Modules

The Sun Blade 6000 modular system design avoids compromises because it provides a complete separation between server modules and I/O modules. Two 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 Network Express Modules (NEMs) provide bulk or consolidated "chassis-wide" common I/O for all of the server modules installed in the system.
- Sun Blade Storage Module M2 provides SAS-2 zone-able storage that can be flexibly allocated amongst the supported server modules in a chassis when connected via appropriate NEMs to expand storage capacity beyond the individual server modules.

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.

Innovative and Highly-Reliable Chassis Design

Organizations need a modular chassis that allows them to deploy exactly the amount of processing and types of I/O that they require, while scaling effectively to meet their growing needs. A blade system chassis should also be designed for a long life so that ongoing enhancements in technology do not force a chassis upgrade.

Oracle's focus on the requirements of the data center has resulted in a chassis design that does not force compromises in the performance, configurations or capabilities delivered by the server modules. For example, in addition to offering a choice of server modules supporting the latest volume processors, these systems deliver 100% of system I/O to the disaggregated I/O modules through a

passive midplane. The Sun Blade 6000 modular system is provided in a 10 rack unit (10U) chassis with up to four chassis supported in a single 42U rack (Figure 1).

Figure 1. Sun Blade 6000 modular systems support Sun Blade 6000 server and storage modules.

Each Sun Blade 6000 chassis can house any combination of server and storage modules up to 10 total, providing support for up to 40 modules per (42 RU) rack while still providing space for additional network aggregation switches such as the Sun Network 10GbE Switch 72p, a 72 port (top of rack) 10 gigabit network switch capable of easily aggregating up to 4 racks full of blades in a dense and efficient network fabric.

The Sun Blade 6000 chassis integrates AC power supplies and cooling fans for all of the server and I/O modules. This approach keeps these components off of the server modules, making them more efficient and more reliable. Power supplies and fan modules in the chassis are designed for ease-of-service, hot-swappability, and redundancy. The chassis provides power and cooling infrastructure to support current and future CPU, memory, and I/O configurations, helping to ensure that the chassis lifecycle will span generations of upgrades. All modular components such as the CMM, server modules, EMs, and NEMs are hot-plug capable. In addition, I/O paths can be configured in a redundant fashion.

Sun Blade 6000 Modular Systems Overview

Sun Blade 6000 modular systems bring significant advancements to deploying modular systems across the organization. The Sun Blade 6000 modular system is ideal for delivering maximum price/performance with superior features as compared to traditional rackmount servers. Supporting a choice of x86 or SPARC platforms, Sun Blade 6000 modular systems are ideal for a variety of applications and markets (Figure 2).

Figure 2. Sun Blade 6000 modular systems can deliver extremely dense data center configurations.

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

- Sun Blade 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 PCIe EMs, which support industry-standard PCIe interfaces and provide each blade with its own unique I/O configuration (just like rack servers)
- Network Express Modules (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 with either fine or coarse granularity, as their needs dictate. Table 1 lists the capacities of the Sun Blade 6000 modular systems. Maximum numbers of sockets, cores, and threads are listed for Intel Xeon, and SPARC processors.

TABLE 1. SUN BLADE 6000 MODULAR SYSTEM CAPACITIES

CATEGORY	SUN BLADE 6000 MODULAR SYSTEM
Sun Blade 6000 server modules	10
Maximum compute nodes per chassis	Up to 20[a]
PCIe EMs	Up to 20
Network Express Modules	Up to 2

Chassis monitoring modules (CMM)	1
Hot-swap power supplies (N+N)	2 @ 5740 Watt ea.
Redundant (rear) cooling fans (N+1)	6
Maximum Intel Xeon sockets/cores/threads	40[a]/240[a]/480[a]
Maximum SPARC T4 sockets/cores/threads	10/80/640
Maximum SPARC T3 sockets/cores/threads	10/160/1280

[a] Sun Blade X6275 M2 server module

Chassis Front Perspectives

The Sun Blade 6000 chassis houses the server modules and I/O modules, connecting the two through a passive mid-plane. Redundant and hot-swappable power supplies and fan modules are also hosted in the chassis. All slots are externally accessible from either the front or the rear of the chassis for easy upgrade or servicing. Server modules, I/O modules, power supplies, and fan modules can all be added and removed while the chassis and other elements in the enclosure are powered on. This capability yields great real-time expansion opportunity and provides considerable flexibility. The front view of the Sun Blade 6000 chassis is shown in Figure 3, with components described in the sections that follow.

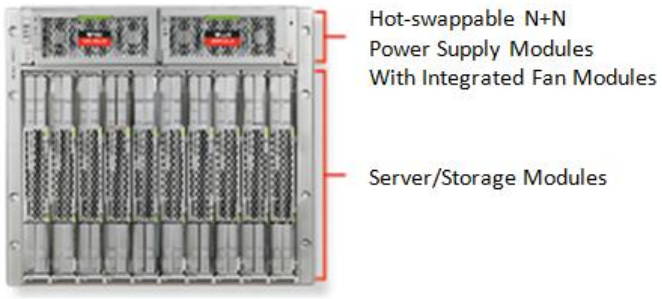

Figure 3. The Sun Blade 6000 chassis is shown from the front perspective.

Operator Panel

An easy-to-read operator panel is located at the top left of the chassis, providing status on the overall condition of the system. Indicators show whether the chassis is on standby or operational mode, and if an over-temperature condition is occurring. A push-button indicator acts as a locator button for the chassis in case there is a need to remotely identify a chassis within a rack, or in a crowded data center. If any of the components in the chassis should present a problem or a failure, the operator panel reflects that issue as well.

Power Supply Modules and Front Fan Modules

Two power supply modules are inserted from the front of the chassis. Each module contains two power supply cores enclosed within a single unit. Each module requires a corresponding number of power inlets. Power supply modules are hot-swap capable and contain a replaceable fan module that

helps cool both the power supplies as well as the PCIe EMs located in the rear of the chassis. In case of a power supply failure, the integrated fan modules will continue to function because they are energized directly from the chassis power grid, independently from the power supply modules that contain them.

The two power supply modules provide the total power required by the chassis. The power supply modules can be configured redundantly in an N+N configuration, with a single power supply module able to power the entire chassis at full load (assuming all installed components do not exceed the power capacity of a single power supply module). In order to provide N+N redundancy, all four power cords must be energized. If both power supply modules are energized, all of the systems in the chassis are protected from power supply failure. A power supply module can fail or be disconnected without affecting the server modules and components running inside the chassis. To further enhance this protection, power grid redundancy for all of the systems and components in the chassis can be easily achieved by connecting each of the two power supply modules to different power grids within the data center.

Sun Blade 6000 power supply modules have a high 90% efficiency rating and an output voltage of 12 v DC. The high efficiency rating indicates that there are fewer power inefficiencies within the power supply itself, therefore wasting less power in the energy conversion stage from alternating current (AC) to direct current (DC). Also, by feeding 12V DC directly to the mid-plane, fewer conversion stages are required in the individual server modules. This strategy yields less power conversion energy waste, and generates less waste heat within the server module, improving overall system efficiency and reliability.

Provisioned power for rackmounted configurations depends on the number of chassis deployed per rack. A 42U rack with four installed Sun Blade 6000 chassis should be provisioned with 24 kilowatts of available power, depending on the ongoing load of the systems. For a more in-depth analysis of day-to-day power consumption of the system please visit the power calculator located at: *http://www.oracle.com/us/products/servers-storage/sun-power-calculators/calc/6000chassis-power-calculator.html*

Server Modules

Any combination of up to 10 Sun Blade 6000 family server or storage 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 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.

Chassis Rear Perspective

The rear of the Sun Blade 6000 chassis (Figure 4) provides access to the I/O modules. Slots for PCIe EMs and Network Express Modules (NEMs) are provided. I/O modules are all hot-swap capable and provide I/O support to server modules.

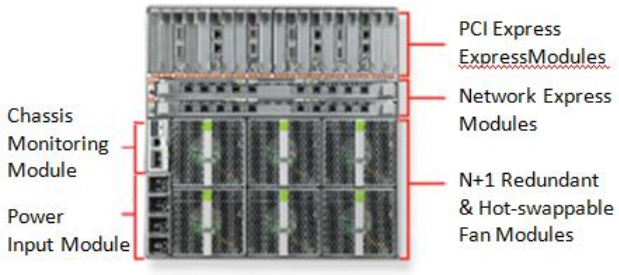

Figure 4. The Sun Blade 6000 chassis is shown from the rear perspective.

PCIe ExpressModules (EMs)

Twenty hot-plug and hot-swap capable PCIe EMs slots are accessible at the rear top of the Sun Blade 6000 chassis. EMs provide a variety of choices for communications, including SAS, Gigabit Ethernet, 10 Gigabit Ethernet, Fibre Channel, Ethernet/Fibre Channel combined, FCoE, and InfiniBand. Different EMs can be selected for every server module in order to provide each server module with the type of I/O connectivity to match the workload it is running. Two PCIe EMs slots are directly available to each server module through the passive midplane. Slots 0 and 1 from right to left are connected to server module 0, slots 2 and 3 are connected to server module 1, and so on continuing across the rear of the chassis.

Network Express Modules (NEMs)

Space is provided for up to two NEMs in the rear of the Sun Blade 6000 chassis. NEMs provide the same I/O capabilities across all of the server modules installed in the chassis, simplifying connectivity and usually offering a low-cost I/O solution by providing common I/O to all of the server modules. All server modules are directly connected to each of the configured NEMs via high-speed point-to-point connections across the passive chassis midplane. More details on available NEMs are provided later in this document.

Chassis Monitoring Module (CMM)

A chassis monitoring module (CMM) is located to the left of the Network Express Module slots on the rear left-hand side of the Sun Blade 6000 chassis. The CMM provides remote monitoring capability and a central access point to the chassis. The CMM includes an integrated network switch that gives LAN access to the CMM's Ethernet ports and to the individual server module management ports. Individual server module management is completely transparent and independent from the CMM.

Power Supply Inlets

Four power supply inlets (plugs) are available from the rear of the Sun Blade 6000 chassis corresponding to the number of power supply cores in the two front-loaded power supply modules.

Integral cable holders prevent accidental loss of power from inadvertent cable removal. Each of the cables requires a 220V, 20A circuit, and a minimum of two circuits are required to power each chassis. For full N+N redundancy, four circuits are required.

Fans and Airflow

Chassis airflow is entirely front-to-back. Airflow is powered by rear fan modules and by the front fan modules mounted in the power supply modules. All rear fan modules are hot-swap capable and N+1 redundant, with six fan modules provided for each Sun Blade 6000 chassis. Each rear fan module is comprised of two redundant in-line fans. The front fan modules pull cool air in from the front of the chassis and blow it across the power supplies and out through the EM and NEM spaces. The rear fan modules pull cool air from the front of the chassis (cool aisle) and out through the rear (hot aisle). When all of the fans in the chassis are running at full speed, the chassis can provide up to 1,000 cubic feet per minute (CFM) of airflow through the chassis.

Passive Midplane

In essence, the passive midplane in the Sun Blade 6000 chassis is a collection of wires and connectors between different modules in the chassis (Figure 5). Since there are no active components, the reliability of these printed circuit boards is extremely high — in the millions of hours, or hundreds of years. The passive midplane provides electrical connectivity between the server modules and the I/O modules.

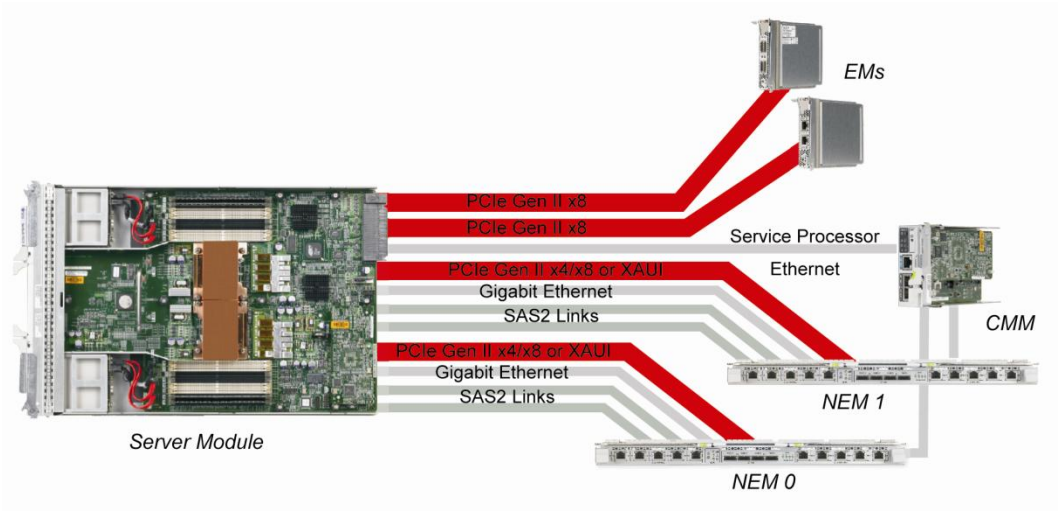

Figure 5. Server modules communicate with PCIe EMs and Network Express Modules via high-speed point-to-point links across the passive chassis midplane.

All modules, front and rear, with the exception of the power supplies and the fan modules, connect directly to the passive midplane. The power supplies connect to the midplane through a bus bar and to the AC inputs via a cable harness. The redundant fan modules plug individually into a set of three fan boards, where fan speed control and other chassis-level functions are implemented. The front fan modules that cool the PCIe EMs each connect to the chassis via blind-mate connections.

The main functions of the mid-plane include:

- Providing a mechanical connection point for all of the server modules
- Providing 12 VDC from the power supplies to each customer-replaceable module
- Providing 3.3 VDC power used to power the System Management Bus devices on each module, and to power the CMM
- Providing a PCIe 2.0 interconnect between the PCIe root complexes on each server module to the EMs and NEMs installed in the chassis
- Connecting the server modules, CMMs, and NEMs to the chassis management network

Each server module is energized through the midplane from the redundant chassis power grid. The midplane also provides connectivity to the I2C network in the chassis, letting each server module directly monitor the chassis environment, including fan and power supply status as well as various temperature sensors. A number of I/O links are also routed through the midplane for each server module.

The capabilities and bandwidth of each server module connection typically includes:

- Two midplane x8 PCIe 2.0 links connect from each server module to each of its dedicated EMs.
- Two midplane x8 PCIe 2.0 links connect from each server module, one to each of the NEMs.
- Two Gigabit Ethernet links are provided, each connecting to one of the NEMs.
- Four x1 Serial Attached SCSI (SAS-2) links are also provided, with two connecting to each NEM slot through the passive midplane.

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.

Oracle VM Server Support in SPARC T-Series Server Modules

Supported in all Oracle servers that utilize SPARC T-Series processors, Oracle VM Server for SPARC (formerly Sun Logical Domains) provides a full virtualization layer, allowing multiple independent virtual machines, each with their own operating system instance, and providing 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 server module with

a single SPARC T4 processor supports up to 64 logical domains, and each individual logical domain can run a unique instance of the operating system[1].

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 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.

Scalability and Support for Oracle Solaris and SPARC T-Series Processor Technology

Oracle Solaris 10 and Oracle Solaris 11 operating systems (OS) are specifically designed to deliver the available capacity of the considerable resources provided by SPARC T4 processor-based systems, such as the SPARC T4-1B server module. In fact, the Oracle Solaris OS provides new functionality for optimal utilization, availability, security, and performance of these systems:

- *Core awareness* — The Oracle Solaris OS is aware of the SPARC processor hierarchies so that the scheduler can effectively balance the load across all the available pipelines. For instance, even though it exposes the SPARC T4 processors as 128 logical processors, Oracle Solaris understands the correlation between individual cores and the threads they support.
- *Fine-granularity manageability* — The Oracle Solaris OS has the ability to dedicate, enable or disable individual processors and threads. In the case of the SPARC processors, this ability extends to individual cores and logical processors (hardware thread contexts). In addition, standard Oracle Solaris OS features such as processor sets provide the ability to define a group of logical processors and schedule processes or threads on them.
- *Binding interfaces* — Oracle Solaris allows considerable flexibility in that processes and individual threads can be bound to either a processor or a processor set, if required or desired.
- *Support for virtualized networking and I/O with hardware accelerated cryptography* — Oracle Solaris contains technology to support and virtualize components and subsystems on the SPARC processor, including support for the dual on-chip 10GbE ports and PCIe interface. As part of a high-performance network architecture, device drivers are provided so that applications running within virtualization frameworks can effectively share I/O and network devices. Hardware

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

accelerated cryptography is supported through the Oracle Solaris Cryptographic framework and the in-core hardware accelerators present in all SPARC processors.

The following subsections provide additional information about important features of Oracle Solaris. For more details about the Oracle Solaris operating system, visit *http://www.oracle.com/technetwork/server-storage/solaris/overview/*.

Oracle Solaris Zones for Consolidation, Secure Partitioning, and Virtualization

Oracle Solaris Zones comprise a group of technologies that work together to efficiently manage system resources, virtualize the system, and provide a complete, isolated, and secure runtime environment for applications. Oracle Solaris Zones can be used to partition and allocate the considerable computational resources of the server modules. Oracle Solaris Zones and Oracle Solaris Resource Management work together with the Oracle Solaris fair-share scheduler on both SPARC processor-based and x86 processor -based server modules.

- *Oracle Solaris Zones* — Oracle Solaris Zones can be used to create an isolated and secure environment for running applications. A zone is a virtualized operating system environment created within a single instance of Oracle Solaris. Oracle Solaris Zones can be used to isolate applications and processes from the rest of the system. This isolation helps enhance security and reliability since processes in one zone are prevented from interfering with processes running in another.
- *Resource Management* — Resource management tools provided with Oracle Solaris let administrators dedicate resources such as CPU cycles to specific applications. CPUs in a multicore multiprocessor system — such as those provided by Sun Blade 6000 family server modules — can be logically partitioned into processor sets and bound to a resource pool, and can ultimately be assigned to an Oracle Solaris Zone. Resource pools provide the capability to separate workloads so that consumption of CPU resources does not overlap. Resource pools also provide a persistent configuration mechanism for processor sets and scheduling class assignment. In addition, the dynamic features of resource pools let administrators adjust system resources in response to changing workload demands.

Oracle Solaris DTrace to Instrument and Tune Live Software Environments

When production systems exhibit nonfatal errors or sub-par performance, the sheer complexity of modern distributed software environments can make accurate root-cause diagnosis extremely difficult. Unfortunately, most traditional approaches to solving this problem have proved time-consuming and inadequate, leaving many applications languishing far from their potential performance levels.

The Oracle Solaris DTrace facility on both SPARC and x86 platforms provides dynamic instrumentation and tracing for both application and kernel activities — even allowing tracing of

application components running in a Java Virtual Machine (JVM)[2]. Oracle Solaris DTrace lets developers and administrators explore the entire system to understand how it works, track down performance problems across many layers of software, or locate the cause of aberrant behavior. Tracing is accomplished by dynamically modifying the operating system kernel to record additional data at locations of interest. Best of all, although Oracle Solaris DTrace is always available and ready to use, it has no impact on system performance when not in use, making it particularly effective for monitoring and analyzing production systems.

NUMA Optimization in Oracle Solaris

With memory managed by each processor on Sun Blade X6270 M2, X6270 M3, and X6275 M2 server modules, the implementations represent non-uniform memory access (NUMA) architectures. Namely, the speed with which a processor can access its own local memory is higher than that required to access memory managed by the other processor. Oracle Solaris provides technology that can specifically help applications improve performance on NUMA architectures.

- *Memory Placement Optimization (MPO)* — The Oracle Solaris OS uses MPO to improve the placement of memory across the physical memory of a server, resulting in increased performance. Through MPO, the Oracle Solaris OS works to help ensure that memory is as close as possible to the processors that access it, while still maintaining enough balance within the system. As a result, many database and enterprise applications are able to run considerably faster with MPO.
- *Hierarchical lgroup support (HLS)* — HLS improves the MPO feature in Oracle Solaris. HLS helps Oracle Solaris optimize performance for systems with more complex memory latency hierarchies. HLS lets Oracle Solaris distinguish between the degrees of memory remoteness, allocating resources with the lowest possible latency for applications. If local resources are not available by default for a given application, HLS helps Oracle Solaris allocate the nearest remote resources.

Oracle Solaris ZFS File System

The Solaris ZFS file system offers a dramatic advance in data management, automating and consolidating complicated storage administration concepts and providing unlimited scalability with the world's first 128-bit file system. ZFS is based on a transactional object model that removes most of the traditional constraints on I/O issue order, resulting in dramatic performance gains. ZFS also provides data integrity, protecting all data with 64-bit checksums that detect and correct silent data corruption.

A Secure and Robust Enterprise-Class Environment

Best of all, Oracle Solaris doesn't require arbitrary sacrifices. The Oracle Solaris Binary Application Guarantee Program helps ensure that existing applications continue to run unchanged, protecting

[2] The terms "Java Virtual Machine" and "JVM" mean a Virtual Machine for the Java platform.

investments. Certified multilevel security protects Oracle Solaris environments from intrusion. Oracle's comprehensive Fault Management Architecture means that elements such as Oracle Solaris Predictive Self Healing can communicate directly with the hardware to help reduce both planned and unplanned downtime.

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, Sun Blade 6000 modular systems provide a robust and wide range of flexible management options. Key components of the management architecture include:

- A dedicated Oracle ILOM service processor on each server module for blade-level management granularity
- A chassis monitoring module (with its own Oracle ILOM service processor) for direct access to server and storage module management features
- Optional Oracle Enterprise Manager Ops Center for server module discovery and OS provisioning as well as bulk application-level provisioning

Within the Sun Blade 6000 modular systems, a chassis monitoring module (CMM) works in conjunction with the service processor on each server module to form a complete and transparent management solution. Each server module contains its own directly addressable management service processor that is also accessible through the CMM. Each server module incorporates Oracle Integrated Lights Out Manager (Oracle ILOM), providing support for IPMI, SNMP, CLI (through serial console or SSH), and HTTP(S) management methods. In addition, Oracle Enterprise Manager Ops Center can be used for discovery, aggregated management, and bulk deployment for multiple systems.

A more detailed discussion of the architecture of these system management components can be found in the Oracle white paper entitled, "Sun Blade 6000 I/O and Management Architecture," which is available for download from the Sun Blade Systems page on OTN at http://www.oracle.com/technetwork/server-storage/sun-blade/documentation/index.html.

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 servers on which they run, including the Sun Blade 6000 modular systems.

Oracle Enterprise Manager Ops Center simplifies infrastructure life-cycle management by letting administrators perform standardized actions across logical groups of systems. Oracle Enterprise Manager Ops Center can automatically discover and group bare-metal systems, performing actions on

the entire group as easily as operating on a single system. Oracle Enterprise Manager Ops Center remotely installs and updates firmware and operating systems, including virtual OS instances running within a virtualized server module. It includes support for:

- Solaris 10 and 11 on x86 platforms
- Oracle Enterprise Linux
- Solaris 8, 9, 10, and 11 on SPARC systems
- Red Hat and SUSE Linux distributions

In addition, the software provides considerable lights-out monitoring of both hardware and software, including fans, temperature, disk and voltage levels — as well as swap space, CPU utilization, memory capacity, and file systems. Role-based access control (RBAC) lets IT staff grant specific management permissions to specific users. A convenient hybrid user interface integrates both a command-line interface (CLI) and an easy-to-use graphical user interface (GUI), providing remote access to manage systems from virtually anywhere.

Oracle Enterprise Manager Ops Center provides advanced management and monitoring features to the Sun Blade 6000 modular systems. The remote management interface discovers and presents the server modules in the chassis as if they were individual servers. In this fashion, the server modules appear in exactly the same way as individual rack mount servers, making the same operations, detailed inventory, and status pages available to administrators. The server modules are discovered and organized into logical groups for easy identification of individual modules, and the system chassis and racks that contain them. Organizing servers into groups also allows features such as OS deployment across multiple server modules. At the same time, individual server modules can also be managed independently from the rest of the chassis. This flexibility allows the management of server modules that may have different requirements than the other modules deployed in the same chassis.

Some of the functions available through Oracle Enterprise Manager Ops Center software include operating system provisioning, firmware updates (for both the BIOS and ILOM service processor firmware), and health monitoring. In addition, Oracle Enterprise Manager Ops Center includes a framework allowing administrators to easily access inventory information, simplifying the task of running jobs on multiple servers with server grouping functionality.

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. 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 makes it easy to select the right platform for Oracle and non-Oracle enterprise applications. Industry-standard I/O and multiple storage options also provide flexibility and leading throughput for individual server modules. Transparent networking and management means that the Sun Blade 6000 modular systems fit easily into an existing network and management infrastructure.

Oracle's Sun Blade 6000 modular systems get blade architecture right. With these systems, Oracle now has one of the most comprehensive modular system families in the industry. This breadth of coverage translates directly to savings in terms of administration and management. For example, unified support for Oracle Solaris across all server modules means that the same features and functionality are available on all processor platforms. This approach saves time in both training and administration — even as the system delivers agile infrastructure for the organization's most critical 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 systems white papers. The Web sites listed below also provide more specific references:

- Blades power calculator*: http://www.oracle.com/us/products/servers-storage/sun-power-calculators/calc/6000chassis-power-calculator.html*
- Sun Blade systems OTN page: *http://www.oracle.com/technetwork/server-storage/sun-blade/documentation/index.html*
- 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/*.

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

Oracle's Sun Blade 6000 Modular Systems
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

Hardware and Software, Engineered to Work Together