An Oracle White Paper
April 2011

High Availability for Business-Critical IT Services Using Oracle's SPARC Enterprise M-Series Servers

ORACLE

Executive Overview

Unavailable IT systems can compromise business results in the form of decreased profitability, reduced employee productivity, revised customer loyalties, and even the inability to continue as a viable company. As organizations strive to minimize planned and unplanned downtime, the importance of IT infrastructure choices, administrative skills, and operational processes increases. Oracle's SPARC Enterprise M-Series servers, Oracle® Solaris, Oracle Solaris Cluster, and educational and professional services from Oracle can help organizations maximize the availability of business-critical IT services. By taking a comprehensive approach to availability, organizations also gain many business benefits, including increased revenue opportunities, lower costs, greater flexibility to respond to changing business conditions, and better ROI.

Introduction

A system fault that prevents customers from accessing critical applications, often equates to closing down a revenue stream for that period of time. Even relatively brief unscheduled outages can have a huge financial impact. For example, when applications in a credit card processing center experience downtime, costs can run as high as $2.6 million per hour. A typical financial broker misses out on up to $6.5 million for every hour a critical system isn't running and downtime for an airline reservation system can result in costs of up to $90,000 per hour[1].

Achieving high levels of availability relies on an interdependent relationship between products, people, and processes. While utilizing highly reliable hardware and software products is

[1] The financial impact of system downtime varies based on the business model and characteristics of the outage. All numbers presented are approximations and not necessarily based on actual events.

important to creating a dependable IT infrastructure, organizations must also take measures to prevent system failures that can result from human errors and flawed procedures. Product and service offerings from Oracle work to address each of these aspects of availability. Oracle's highly-reliable SPARC Enterprise M3000, M4000, M5000, M8000, and M9000 servers combine with Oracle Solaris and Oracle Solaris Cluster to create a dependable hardware and software infrastructure. A broad portfolio of professional and education services from Oracle can help contribute to the creation of a stronger more capable IT organization. In addition, Oracle offers system management tools that can automate procedures, helping organizations increase adherence to best practices and working to prevent system failures that can result from simple administrative mistakes.

Beyond just increasing uptime, adopting Oracle's comprehensive approach to maximizing availability can result in many business benefits, including the following:

- **Higher Revenue.** Oracle's products, training classes, and services help organizations minimize planned and unplanned system downtime. As a result, IT services can achieve greater levels of availability, supporting revenue generation during more hours of the day and more days of the week.

- **Lower Costs.** Management tools from Oracle support automation and improvements to administrative processes. In addition to reducing the opportunity for downtime caused by human error, these tools can help increase operational efficiency, enabling system administrators to manage a greater number of servers with less effort.

- **Greater Business Agility.** Oracle's SPARC Enterprise M3000, M4000, M5000, M8000, and M9000 servers can help organizations respond faster to changing business demands, allowing adjustments to compute capacity — without imposing downtime.

- **Better ROI.** SPARC Enterprise M3000, M4000, M5000, M8000, and M9000 servers support the creation of multiple, fault-isolated hardware partitions with granular resource control, helping simplify efforts to maximize system utilization levels.

Oracle's Comprehensive Approach to Availability

Today, organizational success depends on the availability of IT services. System downtime — both planned and unplanned — can impact schedules, delay revenue, and decrease customer satisfaction. Oracle's comprehensive approach to maximizing availability addresses the role of products, people, and processes on uptime. The following sections take a closer look at key technologies and services from Oracle that can help organizations maximize uptime for business-critical IT services.

Systems with Mainframe-Class Reliability

Specifically designed to support enterprise computing solutions with stringent high availability requirements, Oracle's SPARC Enterprise M3000, M4000, M5000, M8000, and M9000 servers (Figure 1) scale from one to 64 processors and can incorporate up to 4 TB of memory. These highly reliable, vertically-scalable systems are ideal for server consolidation, mainframe rehosting, large-scale database deployments, and mission-critical enterprise application deployments.

Figure 1. SPARC Enterprise M-Series servers are a scalable family of reliable and high-performance symmetric multiprocessing (SMP) servers.

The advanced architecture of these reliable servers fosters high levels of individual system uptime and rapid recovery from many types of hardware faults. SPARC Enterprise M3000, M4000, M5000, M8000, and M9000 servers include redundant and hot-swappable system components, fault resilient power options, diagnostic and error recovery features throughout the design, and built-in remote management features. To support even greater levels of availability, Oracle Solaris Cluster software can be used to assemble clusters that include a mix of up to sixteen SPARC Enterprise M-Series servers and other Sun systems.

SPARC Enterprise M4000, M5000, M8000, and M9000 servers offer Dynamic Domains and Dynamic Reconfiguration (DR) virtualization technologies to enable completion of reconfiguration tasks without

taking the system offline. Administrators can dynamically add or reassign hardware resources to systems in need of additional compute capacity. As a result, organizations can rapidly respond to changing business conditions, simplify resource sharing, and create efficiencies that help improve ROI.

Oracle Solaris further enhances the availability of SPARC Enterprise M-Series servers. Built with a small, compact kernel, the design of Oracle Solaris limits exposure to errors that can bring the operating system down. In addition, Oracle Solaris offers features that help to proactively diagnose, isolate, and recover from both hardware and software failures.

Training and Professional Services

Administrative expertise is critical to achieving the highest levels of system availability. Oracle offers many programs that help businesses to leverage the knowledge and skills of Oracle technical consultants, learning specialists, and support professionals to architect, deploy, and manage a successful implementation. Professional services help to assess overall availability requirements and create methodologies and strategies that can help organizations achieve the required level of uptime. Oracle also offers a full portfolio of education services to help organizations develop the skills needed to build and maintain secure, scalable, and highly-available business solutions.

Process Automation

Organizations can increase availability levels by ensuring that known best practices are followed consistently. By providing the ability to automate administrative tasks, Oracle Solaris and Oracle's portfolio of powerful management applications can help enforce the use of proper procedures — preventing errors and simplifying the management of complex environments. The following tools for SPARC Enterprise M-Series servers help reduce the degree of human interaction and improve procedural accuracy, resulting in higher application availability and lower costs through more efficient operations.

- **The eXtended System Control** Facility on SPARC Enterprise M-Series servers provides built-in remote management and monitoring capabilities and facilitates completion of a broad range of administrative tasks.
- **Oracle Enterprise Manager Ops Center software** helps organizations simplify management of Sun systems, automating installation and administrative tasks, facilitating compliance reporting, and supporting management of thousands of systems simultaneously.
- **Sun Management Center software** from Oracle provides a flexible graphical user interface (GUI) for monitoring, administration, and management of SPARC Enterprise M-Series servers, supporting true remote system management and lights out operation.

Minimizing Planned Downtime

Today's IT organizations are challenged by the need to support non-stop business operations, forcing planned downtime windows to shrink and in some cases disappear entirely. SPARC Enterprise M-Series servers can help decrease planned downtime through features that allow applications to continue

running during the execution of many maintenance and system configuration tasks. By allowing dynamic adjustments to compute capacity, these same features also simplify efforts to maximize resource utilization levels, and can lead to improved ROI.

Dynamic Domains

The Dynamic Domains capability offered by SPARC Enterprise M4000, M5000, M8000, and M9000 servers supports the division of a single system into separate, virtual servers, each running an independent copy of Oracle Solaris and possessing use of designated I/O devices. When system components are exclusively dedicated to individual Dynamic Domains, hardware or software faults in one domain remain isolated and unable to impact the operation of other domains, providing increased application availability.

Dynamic Domains allow organizations to customize and adjust the compute capacity of SPARC Enterprise M4000, M5000, M8000, and M9000 servers to meet specific enterprise needs. For instance, a SPARC Enterprise M9000 server can be configured as a single domain with up to 64 processors in order to host an exceptionally compute-intensive application. Alternatively, an organization with multiple databases that require isolation from one another might divide a single SPARC Enterprise M9000 server into as many as 24 domains. More typical configurations involve three or four domains in a single server.

Domains, used with great success in the mainframe world for a number of years, prove equally useful in the world of open systems. This virtualization approach supports many efficiencies, including the following:

Using Dynamic Domains to virtualize SPARC Enterprise M-Series servers can help organizations deploy fewer physical systems while at the same time providing needed privacy in support of service-level agreements. With fewer servers to manage, Dynamic Domains simplify and unify administration efforts, leading to lower TCO.

Each domain within a single server platform can run a different version of Oracle Solaris, making this technology extremely useful for pre-production testing of new or modified applications or system software.

eXtended System Boards

In order to achieve complete hardware isolation, previous generations of Sun servers designated entire system boards as the smallest unit assignable to a domain. However, some organizations can benefit from more granular resource control and the ability to create a higher number of domains. To meet these needs, the SPARC Enterprise M4000, M5000, M8000, and M9000 servers support eXtended System Boards (XSBs).

There are two types of eXtended System Boards. A Uni-XSB consists of an entire system board or logical system board. Alternatively, a system board or motherboard that is logically divided into four parts is called a Quad-XSB. A Dynamic Domain can consist of any combination of Uni-XSBs and Quad-XSBs, facilitating sophisticated asset allocation. For example, a SPARC Enterprise M4000 server supports up to two Dynamic Domains and consists of a single system board that can be configured as

a Uni-XSB with full control of all resources, or it can be divided into a four Quad-XSB configuration. Logical division lines for the creation of Quad-XSBs in the SPARC Enterprise M4000 server are shown in Figure 2.

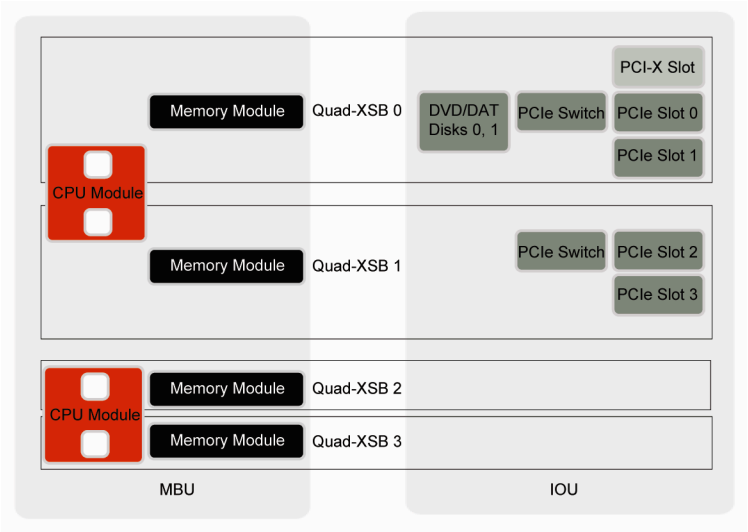

Figure 2. A SPARC Enterprise M4000 server can be divided into Quad-XSBs to allow more granular assignment of resources.

A SPARC Enterprise M5000 server is built as a single motherboard, featuring two logical system boards. Each logical system board can be configured as a Uni-XSB or Quad-XSB. As such, there are three options available for partitioning the resources of a SPARC Enterprise M5000 server. As shown in Figure 3, Figure 4, and Figure 5, a SPARC Enterprise M5000 server can be configured as two Uni-XSBs, eight Quad-XSBs, or a combination of one Uni-XSB and four Quad-XSBs. Within the SPARC Enterprise M5000 server, individual XSBs can be assigned to up to four domains.

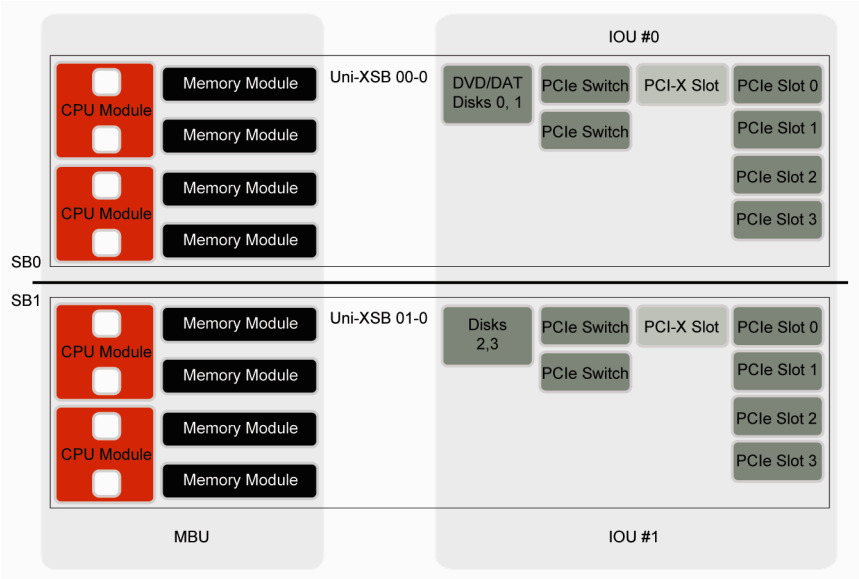

Figure 3. A SPARC Enterprise M5000 can be configured as two Uni-XSBs.

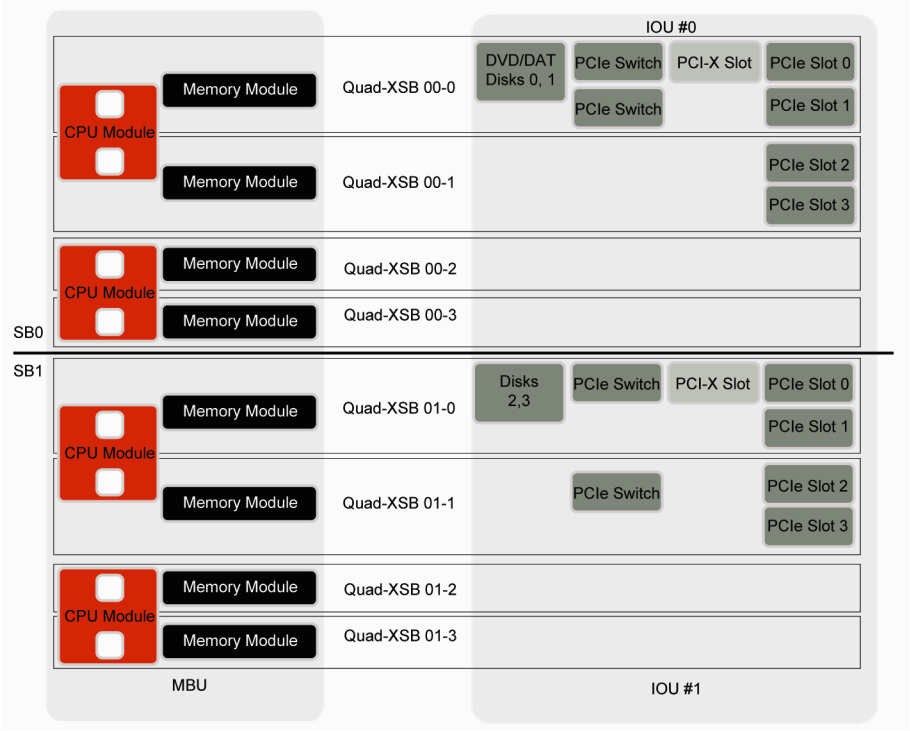

Figure 4. A SPARC Enterprise M5000 can be configured as eight Quad-XSBs.

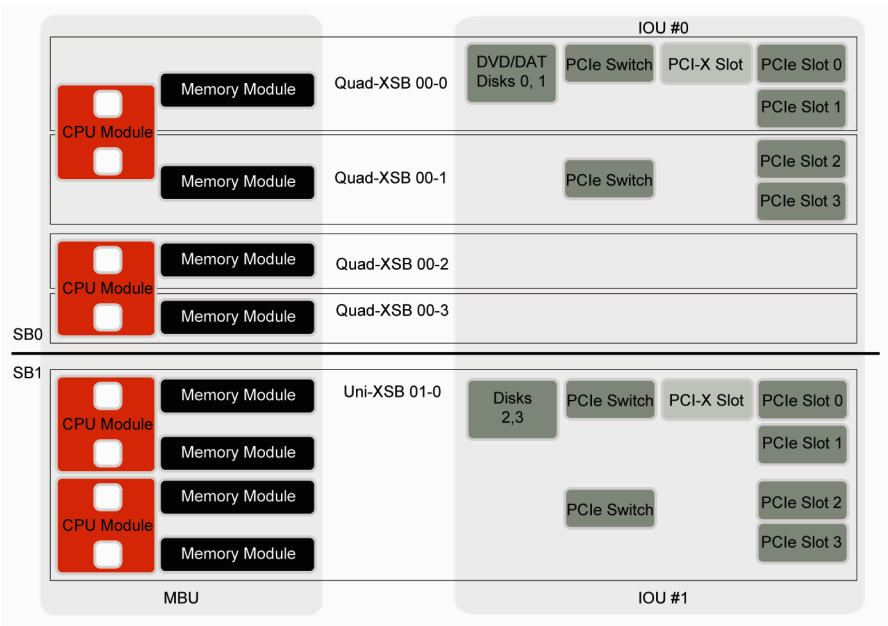

Figure 5. A SPARC Enterprise M5000 can be configured as a combination of one Uni-XSB and four Quad-XSBs.

SPARC Enterprise M8000 servers can support up to four system boards and 16 Dynamic Domains. SPARC Enterprise M9000 servers can support up to eight system boards and 24 Dynamic Domains.

Each system board in these servers can be configured as a Uni-XSB or divided into Quad-XSBs. A Quad-XSB configuration is shown in Figure 6.

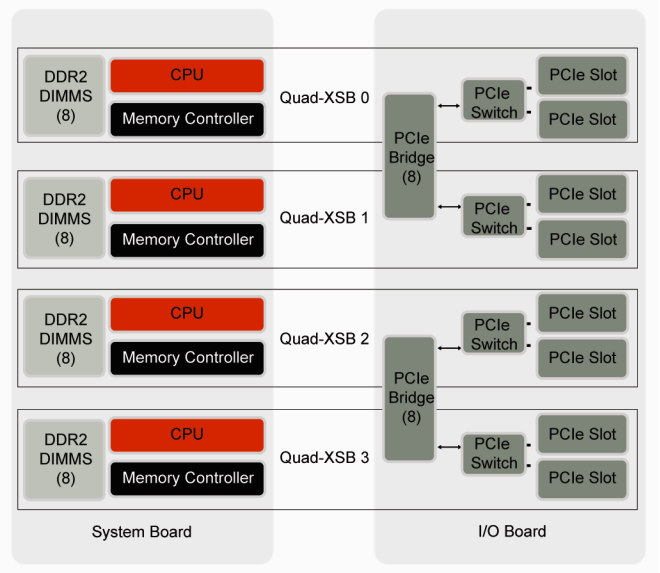

Figure 6. A SPARC Enterprise M8000 or M9000 system board can be configured as four Quad-XSBs.

Determining the exact number and type of XSBs for inclusion in a domain requires balancing the need for fault isolation against the desire to maximize resource utilization. Table 1 helps to differentiate the characteristics of Uni-XSB and Quad-XSB configurations. In general, Uni-XSB's offer the highest level of fault isolation while Quad-XSB's offer greater granularity in regards to the division of resources.

TABLE 1. HIGH-LEVEL COMPARISON OF THE CHARACTERISTICS OF UNI-XSB AND QUAD-XSB CONFIGURATIONS

UNI-XSB	QUAD-XSB
• Physical domaining — boundaries at board level	• Logical division — boundaries at sub-system board level
• Best fault isolation	• Best resource utilization
• Suitable for configuration of a large quantity of resources	• Suitable for configuration of smaller scale domains
• Single Uni-XSB fault incurs greater performance impact	• Single Quad-XSB fault may affect multiple domains

Enforcing Domain Isolation

Domains run applications in strict isolation from applications running in other domains. Isolation and security between domains is maintained by an important ASIC called the System Controller (SC). The SC ensures one domain cannot access data packets from another domain. The SC ASICs on each system board are responsible for the data pathway access. In the event the system board is placed in quad mode, the SC further restricts the data pathways. This applies to other XSBs on the same board, as well as addressing from other XSB on other boards, with data passed via the Crossbar Unit.

Online Reconfiguration, Maintenance, and Upgrades

Given the constant changes in business priorities, IT managers often need the ability to adjust the amount of compute power assigned to each application workload. Dynamic Reconfiguration (DR)

technology provides added value to Dynamic Domains by allowing resources to shift without taking the system offline. Using Dynamic Reconfiguration technology, administrators can perform physical or logical changes to system hardware resources while the server continues to execute applications.

The ability to remove or add components such as CPUs, memory, and I/O subsystems while a system continues to run can help reduce planned system downtime for capacity reconfiguration, maintenance tasks, and upgrades. In fact, Dynamic Reconfiguration simplifies these tasks by eliminating the need for system reboots after hardware configuration changes. Multiple DR operations can also execute simultaneously for efficient management of resources. This capability helps independent domain administrators to perform DR operations simultaneously without concern for the status of DR requests or executions in other domains.

The following examples demonstrate how Dynamic Reconfiguration can help organizations dramatically reduce planned downtime for maintenance activities and support adjustments to the compute capacity of a Dynamic Domain.

- **Online Upgrades.** Many essential compute systems need the capability to respond quickly to unpredictable loads while continuing to deliver critical services. Online upgrades facilitate rapid installation of additional capacity within SPARC Enterprise M4000, M5000, M8000, and M9000 systems — sometimes instantly — without interruption. Using this feature, additional processors, memory, I/O, and mass storage devices can be added to a Dynamic Domain without shutting down application processing or requiring the system to reboot.
- **Concurrent Maintenance.** Administrators are increasingly challenged to carve out planned downtime windows to perform necessary service on essential systems. Since SPARC Enterprise M4000, M5000, M8000, and M9000 servers are equipped with redundant components, failures no longer necessarily result in server downtime. However, the replacement of failed or degraded components must occur at some point. Dynamic Reconfiguration simplifies maintenance operations by supporting the completion of component removal and installation tasks while the system continues to operate.
- **Automated Adjustments to Compute Capacity.** Automatic Dynamic Reconfiguration (ADR) supports the execution of Dynamic Reconfiguration operations without interaction from a user. ADR activities are triggered by pre-defined system events set by a system administrator. For example, an ADR configuration can be created to allow the automatic addition of a system board when a business-critical application reaches full CPU utilization. Implementations can include application-specific preparatory tasks before a DR operation, execution of error recovery actions during DR, and clean-up procedures after DR completion.
- **Capacity on Demand.** The Capacity on Demand (COD) option helps SPARC Enterprise M4000, M5000, M8000, and M9000 servers gain immediate access to extra capacity in the event of a resource shortage. For example, if a spike in demand creates a drop in service level performance, configuration of additional compute power can occur very rapidly by moving a COD system board into the target domain. COD resources remain offline at no charge until the resources are actually needed. When additional processing power is required, COD system boards can be immediately configured into the system by purchasing a COD hardware activation option for each extra CPU.

Reducing Unplanned Downtime

Unplanned downtime can be extremely costly to organizations. In addition to missed business transactions, the true cost of unplanned downtime also includes expenses related to lost employee productivity and decreased customer satisfaction. To help organizations avoid unplanned downtime, SPARC Enterprise M-Series servers include built-in and configurable redundancy, hot-swap components, and advanced reliability features that facilitate rapid recovery or continuous operation despite error conditions. Table 2 lists the redundant and hot-swappable components available for SPARC Enterprise M-Series servers.

TABLE 2. REDUNDANT AND HOT-SWAPPABLE COMPONENTS HELP INCREASE SYSTEM AVAILABILITY AND EASE MAINTENANCE

	REDUNDANT COMPONENTS	HOT-SWAPPABLE COMPONENTS
SPARC ENTERPRISE M3000 SERVERS	• Disk Drives • Power Supply Units • Fan Units	• Disk Drives • Power Supply Units • Fan Units
SPARC ENTERPRISE M4000 SERVERS AND SPARC ENTERPRISE M5000 SERVERS	• Internal System Controllers • Disk Drives • Power Supply Units • Fan Units • External I/O Expansion Unit (optional)	• Disk Drives • Power Supply Units • Fan Units • External I/O Expansion Unit (optional) • PCI Cards
SPARC ENTERPRISE M8000 SERVERS AND SPARC ENTERPRISE M9000 SERVERS	• Internal System Controllers • Disk Drives • Power Supply Units • Fans in Fan Tray Unit • Dual Power Feed (optional) • CPU Memory Board Units • I/O Units • External I/O Expansion Unit (optional) • eXtended System Control Facility Units • Crossbar Units • Clock Board (M9000 only)	• Disk Drives • Power Supply Units • Fan Units • CPU Memory Board Units • I/O Units • External I/O Expansion Unit (optional) • PCI Cards • DVD-ROM drive • Tape drive • eXtended System Control Facility Units

Built-in Hardware Redundancy

Key components of SPARC Enterprise M-Series servers that offer built-in redundancy to help avoid unplanned system downtime are described in the following sections.

Multi-Path System Bus Architecture

SPARC Enterprise M4000, M5000, M8000, and M9000 servers each include a level of redundancy within the system bus. These interconnects vary in implementation but share many of the same

characteristics — providing combined address, data, and control transport. System buses are uni-directional paths with multiplexed address and data, plus control and ECC in each direction. The motherboard in the SPARC Enterprise M4000 and M5000 servers, and the CPU/Memory board unit in the SPARC Enterprise M8000 and M9000 servers contain multiple system controllers (SC), enhancing reliability and performance.

The architecture of the system interconnect in these servers delivers more than just outstanding reliability. By implementing a point-to-point bus and using packet-switched technology, multiple data streams can be transmitted, improving performance and response time. For example, the server interconnect implemented in the SPARC Enterprise M9000 server delivers as much as 304.2 GB/sec. of peak bandwidth, offering 7.5 times more system throughput than Sun Fire E20K and Sun Fire E25K servers while maintaining similar latency.

SPARC Enterprise M4000 Server Architecture

The SPARC Enterprise M4000 server is implemented within a single motherboard and features one logical system board with two system controllers. Both system controllers connect to each other, as well as CPU modules, memory address controllers, and the IOU (Figure 7).

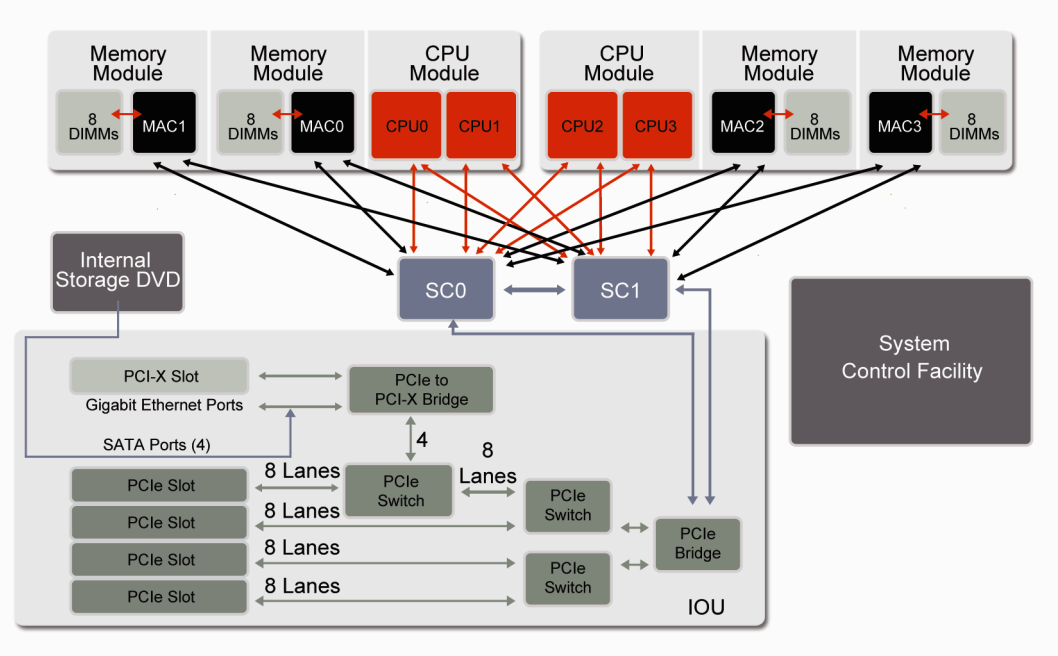

Figure 7. The SPARC Enterprise M4000 server system interconnect features two system controllers.

SPARC Enterprise M5000 Server Architecture

The SPARC Enterprise M5000 system is implemented within a single motherboard but features two logical system boards. Similar to the SPARC Enterprise M4000 server design, each logical system board contains two system controllers that connect to each other, as well as CPU modules, memory access controllers, and an IOU. In addition, each system controller connects to a corresponding system controller on the other logical system board (Figure 8).

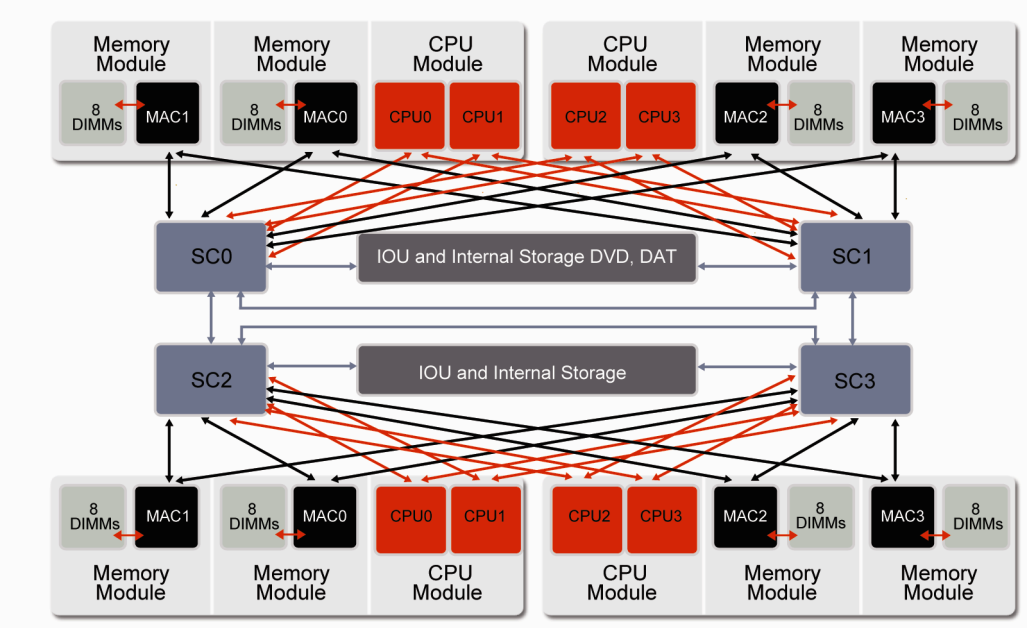

Figure 8 The SPARC Enterprise M5000 server system interconnect includes four system controllers.

SPARC Enterprise M8000 and M9000 Server Architecture

SPARC Enterprise M8000 and M9000 servers feature multiple system boards that connect to a common crossbar. Each system board contains four system controllers and each system controller connects to every CPU module. For improved bandwidth, every memory access controller connects to two system controllers, and each system controller connects to every other system controller within the system board. The system controllers also provide a connection to each crossbar unit, supporting data transfer to other system boards (Figure 9).

In addition to redundancy in the system interconnect, most of the internal parts of SPARC Enterprise M9000 server clock chip are redundant. There are two sources of clock signal and a dual signal line is implemented between the clock chip and the system boards. By implementing redundant bus routes, the SPARC Enterprise M9000 server can automatically restart in the event one route fails. The surviving route is used to continue operation.

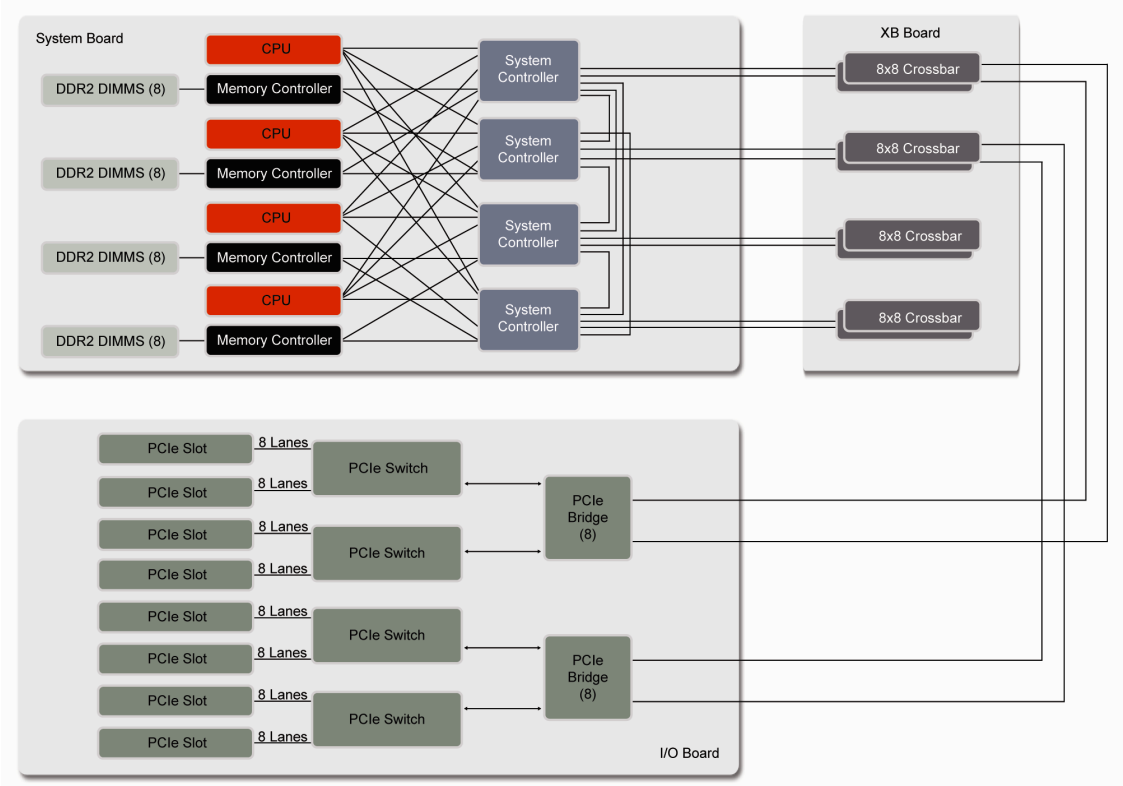

Figure 9. The architecture of SPARC Enterprise M8000 and M9000 servers provides redundant system controllers on each system board and redundant crossbar interconnects between system boards.

Service Processor — eXtended System Control Facility

SPARC Enterprise M3000, M4000, M5000, M8000, and M9000 servers use an eXtended System Control Facility (XSCF) for system operation, monitoring, and diagnostic and error notification functions. The XSCF is a service processor that operates independently from the main server system and increases manageability. The XSCF component is designed with features that enhance service processor availability. For example, the XSCF provides two Ethernet interfaces to support a redundant network connection to an administrative LAN, as well as an interface port for the attachment of an uninterruptable power supply.

The SPARC Enterprise M8000 and M9000 servers feature redundant XSCF boards. One XSCF board remains in standby mode ready to take over operations in the event of a fault to the primary XSCF. The XSCF in the SPARC Enterprise M8000 and M9000 servers is also hot-swappable, supporting maintenance of the XSCF without affecting server operation.

Redundant Power and Cooling

SPARC Enterprise M3000, M4000, M5000, M8000, and M9000 servers use separate modules for power and cooling. Redundancy in each of these subsystems combined with temperature sensors placed throughout the system keeps servers operating even under power or fan fault conditions. Sensors measure temperatures on processors and key ASICs as well as the ambient temperature at

several locations. Power and cooling specification for SPARC Enterprise M3000, M4000, and M5000 servers are included in Table 3 and power and cooling specifications for SPARC Enterprise M8000 and M9000 servers are included in Table 4.

TABLE 3. POWER AND COOLING FEATURES OF SPARC ENTERPRISE M3000, M4000, AND M5000 SERVERS

	SPARC ENTERPRISE M3000 SERVER	SPARC ENTERPRISE M4000 SERVER	SPARC ENTERPRISE M5000 SERVER
FAN UNITS	• 2 fan units • 2 80mm fans • 1+1 redundant	• 4 fan units • 2 172 mm fans • 2 60 mm fans • 1 of each type is redundant	• 4 fan units • 4 172 mm fans • 2 fan groups, each containing 2 fan units • 1 redundant fan per fan group
POWER SUPPLIES	• 505 watts • 2 units • 1+1 redundant • Single-phase • AC or DC Power Supplies	• 2350 watts • Maximum Power Consumption: 1390 watts • 1 unit • 1+1 redundant • Single-phase	• 4590 watts • Maximum Power Consumption: 2381 watts • 4 units • 2 + 2 redundant • Single-phase
POWER CORDS	• 2 power cables • 1+1 redundant	• 2 power cables • 1+1 redundant power cable	• 4 power cables • 2+2 redundant power cables

TABLE 4. POWER AND COOLING FEATURES OF SPARC ENTERPRISE M8000 AND M9000 SERVERS

	SPARC ENTERPRISE M8000 SERVER	SPARC ENTERPRISE M9000-32 SERVER	SPARC ENTERPRISE M9000-64 SERVER
FAN UNITS	• 12 fan units • Four 172 mm fans • Eight 60 mm fans • N+1 redundant	• 16 fan units • 16 172 mm fans • N+1 redundant	• 32 fan units • 32 172 mm fans • N+1 redundant
POWER SUPPLIES	• 10,500 watts • Maximum Power Consumption: 7.6 kW • 9 units • N+1 redundant	• 21,300 watts • Maximum Power Consumption: 14.9 kW • 15 units • N+1 redundant	• 42,600 watts • Maximum Power Consumption: 30.4 kW • 30 units • N+1 redundant
OPTIONS	• Single-phase • Three-phase • Dual-grid	• Single-phase • Three-phase • Dual-grid	• Single-phase • Three-phase • Dual-grid
POWER CORDS	• 3 power cables (single feed) • 6 power cables (dual feed) • 2 power cables (three-phase)	• 5 power cables (single feed) • 10 power cables (dual feed) • 2 power cables (three-phase)	• 10 power cables (single feed) • 20 power cables (dual feed) • 4 power cables (three-phase)

Fan Units

SPARC Enterprise M-Series servers use fully redundant, hot-swappable fans as the primary cooling system. If any one fans fails, the XSCF detects the failure and switches the remaining fans to high-speed operation to compensate for the reduced airflow. SPARC Enterprise M3000, M4000, M5000, M8000, and M9000 servers can operate normally under these conditions, providing ample time to service the failed unit. Furthermore, replacement of fans units can occur without interrupting application operation.

Power Supplies

The use of redundant power supplies and power cords adds to the fault resilience of SPARC Enterprise M3000, M4000, M5000, M8000, and M9000 servers. Power is supplied to these servers by redundant hot-swap power supplies, supporting continued server operation even if a power supply fails. Since the power units are hot-swappable, removal and replacement can occur while the system continues to operate.

As an option, SPARC Enterprise M8000 and M9000 servers can be ordered with a three-phase power supply unit and corresponding server cabinet. Models with a three-phase power supply permit two configurations, a star connection that connects a neutral line and each phase, and a delta connection that connects each phase.

Optional Dual Power Feed

The SPARC Enterprise M8000 and M9000 servers are dual power feed capable. The AC power subsystem in these servers is completely duplicated, providing the option for reception of power from two external AC power sources that are independent of each other. The use of a dual power feed and redundant power supplies increases system availability, as server operations can remain unaffected even after a single power grid failure.

Configuring Additional System Redundancy

The expansion capabilities of SPARC Enterprise M4000, M5000, M8000, and M9000 servers help support configuration of redundant numbers of processors, memory DIMMs, and I/O devices. Configuring a system with more processors, memory, or I/O components than required for a particular workload essentially creates a set of online hot spares that can help speed service resumption. For example, in the event of a processor fault, the SPARC Enterprise M-Series server can reboot, isolate and offline the failed processor and then run on the remaining good processors.

Organizations can create configurations with levels of system redundancy for each of the following components to match specific enterprise needs.

- **CPU and memory.** SPARC Enterprise M4000, M5000, M8000, and M9000 servers support the creation of redundant processing configurations with up to 64 processors and as much as 4 TB of memory. In addition, dynamically reconfigurable CPU and memory resources facilitate the active addition or removal of compute power without stopping system operation.

- **Built-in I/O.** Internal PCI expansion slots on SPARC Enterprise M-Series servers provide connectivity to external devices. Combining these expansion slots with optional disk management software, system configuration can easily include multiple connections to peripheral devices, providing redundant controllers and channels. Software tools can also maintain the multiple paths and facilitate a switch to an alternate path in the event the primary path fails. In addition, Dynamic Reconfiguration can add or remove an entire I/O unit and mounted PCI cards from the SPARC Enterprise M8000 or SPARC Enterprise M9000 system during active operation.
- **External I/O Expansion Unit.** SPARC Enterprise M4000, M5000, M8000, and M9000 servers support the connection of one or more optional, rackmount External I/O Expansion Units. These units are hot-swappable and can mount up to 12 PCI-X or 12 PCI Express cards. In addition, the External I/O Expansion Unit includes redundant hot-swappable power supplies and fan units.
- **PCI cards.** PCI cards that support PCI hot plug are hot-swapped into the system by placing the expansion card into one of the supplied cassettes before insertion into a SPARC Enterprise M4000, M5000, M8000, or M9000 server internal PCI slot or External I/O Expansion Unit. PCI cards are unmounted from the system by using this procedure in reverse. The ability to hot-swap PCI cards adds flexibility and eases serviceability.
- **Removable media.** While disk and tape devices are directly integrated into SPARC Enterprise midrange servers, an add-on base I/O card provides access to internal devices on high-end SPARC Enterprise servers. When combined with disk mirroring software, multiple hot-swappable, internal serial attached SCSI (SAS) hard disk drives provide support for boot disk redundancy. Also, hot-swappable DVD-ROM and tape drive units in the SPARC Enterprise M8000 and M9000 increase the flexibility of these systems.

Specific expansion capabilities for individual servers are found in Table 5 and Table 6.

TABLE 5. EXPANSION CAPABILITIES OF SPARC ENTERPRISE M3000, M4000, AND M5000 SERVERS

	SPARC ENTERPRISE M3000 SERVER	SPARC ENTERPRISE M4000 SERVER	SPARC ENTERPRISE M5000 SERVER
PROCESSORS	• 1 • 2.52 GHz SPARC64 VII • 2.75 GHz SPARC64 VII • 2.86 GHz SPARC64 VII+	• 2 or 4 • 2.15 GHz SPARC64 VI • 2.53 GHz SPARC64 VII • 2.66 GHz SPARC64 VII+	• 2, 4, 6, or 8 • 2.15 GHz SPARC64 VI • 2.53 GHz SPARC64 VII • 2.66 GHz SPARC64 VII+
MEMORY	• Up to 64GB maximum • 8 DIMM slots	• 256 GB maximum • 32 DIMM slots	• 512 GB maximum • 64 DIMM slots
EXTERNAL I/O CHASSIS	• None	• Up to 2 units	• Up to 4 units
INTERNAL I/O SLOTS	• 4 PCI Express	• 4 PCI Express • 1 PCI-X	• 8 PCI Express • 2 PCI-X
HARD DISK DRIVES	• Up to 4 drives	• Up to 2 drives	• Up to 4 drives

TABLE 6. EXPANSION CAPABILITIES OF SPARC ENTERPRISE M8000 AND M9000 SERVERS

	SPARC ENTERPRISE M8000 SERVER	SPARC ENTERPRISE M9000-32 SERVER	SPARC ENTERPRISE M9000-64 SERVER
PROCESSORS	• Up to 16 • 2.28 GHz or 2.4 GHz SPARC64 VI • 2.88 GHz SPARC64 VII • 3.0 GHz SPARC64 VII+	• Up to 32 • 2.28 GHz or 2.4 GHz SPARC64 VI • 2.88 GHz SPARC64 VII • 3.0 GHz SPARC64 VII+	• Up to 64 • 2.28 GHz or 2.4 GHz SPARC64 VI • 2.88 GHz SPARC64 VII • 3.0 GHz SPARC64 VII+
MEMORY	• 1 TB maximum • 128 DIMM slots	• 2 TB maximum • 256 DIMM slots	• 4 TB GB maximum • 512 DIMM slots
EXTERNAL I/O CHASSIS	• Up to 8 units	• Up to 16 units	• Up to 16 units
INTERNAL I/O SLOTS	• 16 PCI Express	• 32 PCI Express	• 96 PCI Express
HARD DISK DRIVES	• Up to 16 drives	• Up to 32 drives	• Up to 64 drives

Advanced Reliability Features

Within SPARC Enterprise M3000, M4000, M5000, M8000, and M9000 servers, architectural redundancies that inherently increase reliability are augmented by technology advances and error detection and recovery capabilities. SPARC Enterprise M3000, M4000, M5000, M8000, and M9000 servers feature the following important technologies that can help correct failures and keep marginal components from causing repeated downtime.

Automatic System Recovery

SPARC Enterprise M3000, M4000, M5000, M8000, and M9000 servers can recover quickly from many types of faults, including serious issues such as the failure of a CPU or a critical ASIC. In fact, no single hardware component failure prohibits SPARC Enterprise M8000 and M9000 servers from booting. When a fault occurs, the specific corrective action taken by the system depends upon the component and type of error. The system may continue to operate in a degraded mode or may start automatic system recovery (ASR) triggering an immediate system reboot and automatic configuration around a failed component. ASR capabilities help prevent faulty or marginal hardware from causing repeated failures or keeping an entire system down.

Reduced Component Count and Complexity

The basic architecture of SPARC Enterprise M3000, M4000, M5000, M8000, and M9000 servers provides the foundation for achieving outstanding system reliability. These server designs boast a lower component count and less complexity than the previous generation of servers, resulting in reduced Mean Time Between Failure (MTBF) ratings. The number of ASICs included in the design of these systems is as much as 60 percent lower than the previous generations of systems. For example, the ECC-protected interconnect features combined address, data, and control lines for all operations (i.e., Memory Read, Memory Write, Copy-Back, Interrupts, Peripheral Read), minimizing the signals required and maximizing both address and data bandwidth.

SPARC64 VI and SPARC64 VII/VII+ Processor

SPARC Enterprise M3000, M4000, M5000, M8000, and M9000 servers feature SPARC64 VI and SPARC64 VII/VII+ processors. The design of these processor modules increases system reliability by delivering improved fault avoidance and error correction capabilities. In fact, much of the area on the SPARC64 VI and SPARC64 VII/VII+ processor die is dedicated to error detection and correction of data within the CPU. Processor RAM units are ECC protected or duplicated and most latches and executive units are parity protected. Rather than force the loss of operation of the entire processor, a single bad core can be isolated and taken offline. In addition, the SPARC64 VI and SPARC64 VII/VII+ CPUs can autonomously offline bad cores without intervention from the service processor. Error data generated by the SPARC64 VI and SPARC64 VII/VII+ processors is sent to the service processor to support preventive maintenance, simplifying identification of the location and type of a failure, increasing serviceability.

Processor RAM

The SPARC64 VI and SPARC64 VII/VII+ processors offer reliability features that support high levels of data integrity. When memory read data has a multi-bit error, a special mark identifying the source of the error is written into the data and the ECC syndrome becomes a special value, providing valuable information for identifying the source of the fault. Table 7 highlights the error detection and correction capabilities of the SPARC64 VI and SPARC64 VII/VII+ processor.

TABLE 7. PROCESSOR DETECTION AND ERROR CORRECTION METHODS

TYPE	ERROR DETECTION OR PROTECTION METHOD	ERROR CORRECTION METHOD
L1 INSTRUCTION CACHE (DATA)	Parity	Invalidation and reread
L1 INSTRUCTION CACHE (TAG)	Parity plus duplication	Rewrite of duplicated data
L1 DATA CACHE (DATA)	Single Error Correction Double Error Detection (SECDED)	ECC One-bit error correction using ECC
L1 DATA CACHE (TAG)	Parity plus duplication	Rewrite of duplicated data
L2 CACHE (DATA)	SECDED ECC	One-bit error correction using ECC
L2 CACHE (TAG)	SECDED ECC	One-bit error correction using ECC
INSTRUCTION TLB PARITY INVALIDATION	Data TLB Parity Invalidation	Branch History Parity Recovery from branch predication failure

Since the SPARC64 VII/VII+ processor implements a set-associative scheme that divides the L1 cache, L2 cache, and TLB into way units, degradation can be performed separately in a granular manner. Error occurrences are counted for each unit. When an error occurrence count for a specified time interval exceeds the upper limit, degradation is performed and the relevant way unit is not used subsequently. For this type of event, hardware automatically performs degradation and initiates the following operations to assure the continuity of coherency of data stored in cache:

- **Write-back to L2 cache.** Write-back the dirty lines in the way unit of the L1D cache to be degraded.
- **Write-back to Memory.** Write-back the dirty lines in the way unit of the L2 cache to be degraded.

The degradation of a way unit is performed without adversely affecting software programs, and application operation is free from any effects except for a slowdown of processing speed.

Internal Registers and Execution Units

To further increase reliability, SPARC64 VI and SPARC64 VII/VII+ processors also provide error protection for registers and execution units. A summary of these capabilities is found in Table 8.

TABLE 8. INTERNAL REGISTERS AND EXECUTION UNIT ERROR DETECTION AND DATA PROTECTION METHODS

TYPE	ERROR DETECTION OR PROTECTION METHOD
Integer register	SECDED ECC
Floating-point register	Parity
PC, PSTATE register	Parity
Computation input-output register	Parity
Addition, subtraction, division, shift, and graphic execution unit	Parity prediction
Multiplication Execution Unit	Parity prediction plus residue check

Instruction Retry

SPARC64 VI and SPARC64 VII/VII+ processors employ a synchronous update method and include ECC protection for integer architecture registers. When an error occurs, the ECC circuit corrects the error. The floating point architecture registers and other registers are protected by parity bits. Also, the parity prediction circuit, residue check circuit, and other circuits are implemented within an execution unit to propagate parity information to output results.

In the unlikely event that a parity error is detected, instructions that reside at the point of execution are canceled. Intermediate results before commitment can be discarded, and only results updated by instructions that have been completed without encountering any errors remain in programmable resources. Therefore, not only can the destruction of programmable resources due to errors be prevented, hardware can also perform an instruction retry after error detection. Since stalled instructions can be discarded once and then retried from the beginning, there is a possibility of recovery in the case of a hang-up.

As shown in Figure 10, instruction retry is triggered by an error and is automatically started. A retry is performed instruction-by-instruction to increase the chance of normal execution. When the execution

completes normally, the state automatically returns to the normal execution state. During this period, no software intervention is required, and if the instruction retry succeeds, the error does not affect software. An instruction retry is repeated until the number of retry times reaches the threshold. If the threshold is exceeded, the processor logs the source of the error and reports the status to the operating system.

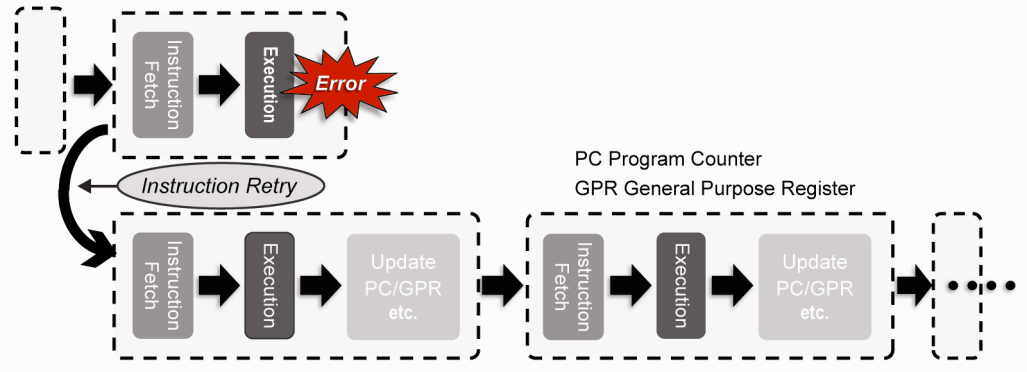

Figure 10. The SPARC64 VI and SPARC64 VII processors implement an automated instruction retry process to increase availability.

Memory Patrol, Extended-ECC, and Mirroring

The memory subsystem of the SPARC Enterprise M3000, M4000, M5000, M8000, and M9000 servers includes multiple technologies to increase stability[2]. The following features work toward early diagnosis and fault isolation within the memory subsystem to preserve system integrity and raise application availability.

- **Memory Patrol.** Memory patrol periodically performs a scan to detect memory errors. This function prevents use of faulty areas of memory, resulting in less opportunity for system or application errors and increases in system reliability.
- **Memory Extended-ECC.** The memory Extended-ECC function of these servers supports single-bit error correction, facilitating continuous processing despite events such as burst read errors that are sometimes caused by memory device failures.
- **Memory Mirroring.** When memory mirroring mode is enabled on SPARC Enterprise M4000, M5000, M8000, and M9000 servers, the memory subsystem duplicates the data on write and compares the data on read to each side of the memory mirror. In the event that errors occur at the bus or DIMM level, normal data processing continues through the other memory bus and alternate DIMM set. In SPARC Enterprise M4000 and M5000 servers, memory is mirrored within the same memory module, using the common memory address controller (MAC) ASIC (Figure 11 and Figure 12).

[2] Memory Mirroring is not supported on the SPARC Enterprise M3000 server.

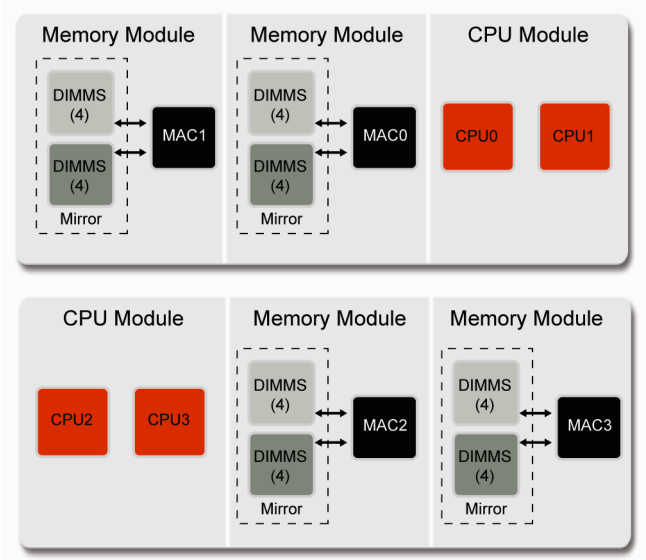

Figure 11. The SPARC Enterprise M4000 server memory architecture mirrors memory within the same memory module.

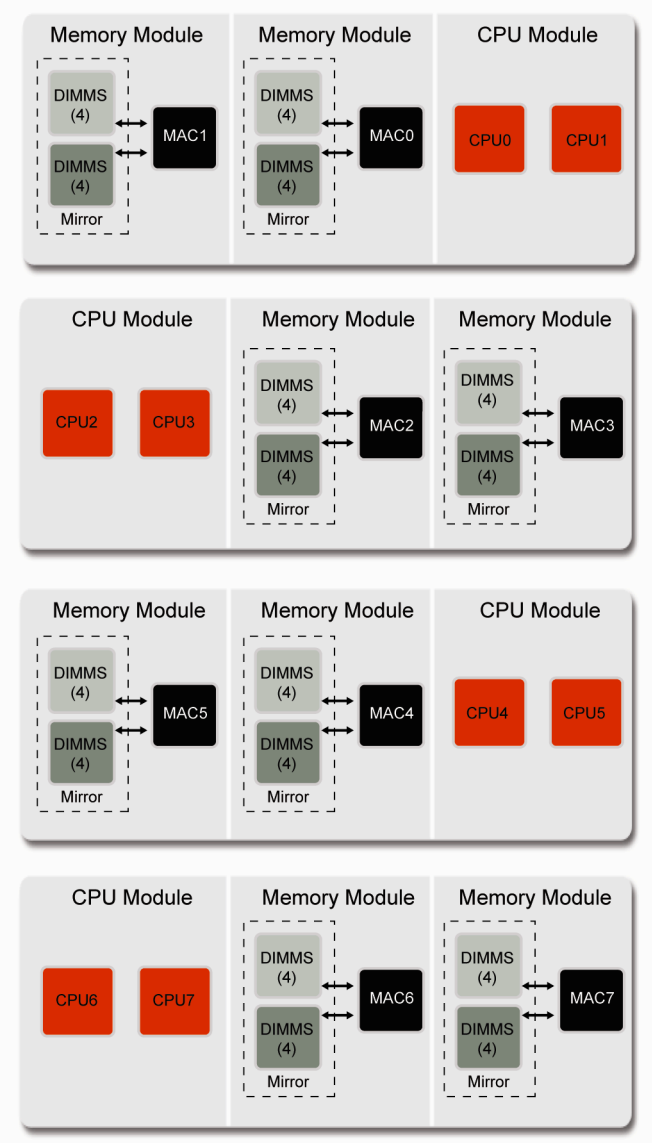

Figure 12. The SPARC Enterprise M5000 server memory architecture mirrors memory within the same memory module.

On SPARC Enterprise M8000 and M9000 servers, memory is mirrored across adjacent MAC ASICs to increase reliability (Figure 13). However, the configuration of Quad-XSBs in SPARC Enterprise high-end server system boards prevents the use of memory mirroring since each Quad-XSB only contains one MAC.

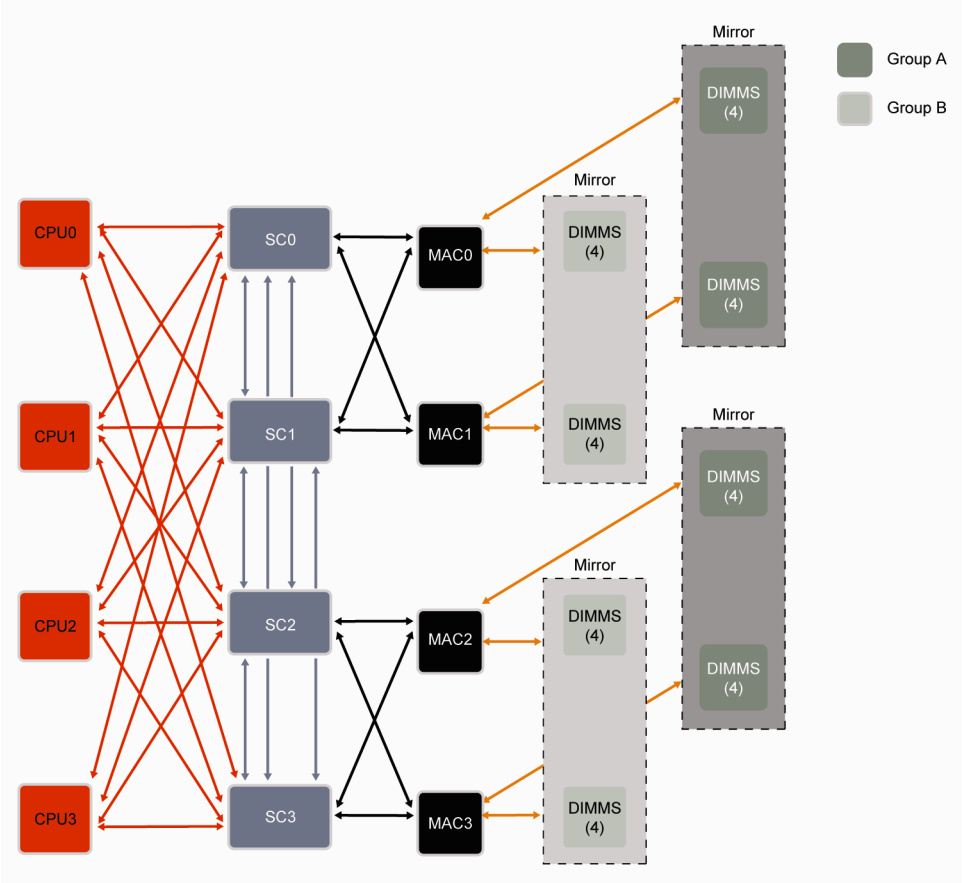

Figure 13. SPARC Enterprise M8000 and M9000 server memory architecture mirrors the memory of adjacent MAC ASICs

Additional Error Detection, Diagnosis, and Recovery Features

SPARC Enterprise M3000, M4000, M5000, M8000, and M9000 servers include features for error detection, diagnosis, and recovery throughout the system architecture. The following list highlights capabilities which not only identify problems, but work to mitigate consequences and contribute to rapid recovery of system operation.

- **End-to-end ECC protection.** SPARC Enterprise M-Series servers routinely store and move significant amounts of data. As such, these servers provide end-to-end error detection and correction on data paths within the system, including the SPARC64 VI and SPARC VII/VII+ processors, main memory, and data paths on the system interconnects.
- **State-of-the-art fault isolation.** While error conditions prompt previous generations of enterprise servers to offline entire hardware devices, the SPARC Enterprise M-Series servers isolate errors within component boundaries and offline only the relevant chips. Isolating errors down to the chip improves stability and provides continued availability of maximum compute power. This feature applies to CPUs, memory access controllers, crossbar ASICs, system controllers, and I/O ASICs.

- **Environmental monitoring.** Environmental monitoring provides a historical log of pertinent environmental and error conditions.
- **Host watchdog.** The host watchdog feature of SPARC Enterprise M4000, M5000, M8000, and M9000 servers periodically checks for operation of software, including the domain operating system. This feature also uses the XSCF firmware to trigger error notification and recovery functions.
- **Dynamic CPU resource deallocation.** SPARC Enterprise M4000, M5000, M8000, and M9000 servers support Dynamic CPU resource deallocation including processor fault detection, isolation, and recovery. This feature dynamically reallocates CPU resources into an operational system using Dynamic Reconfiguration without interrupting the applications that are running.
- **Component status checks.** Periodic checks are performed to determine the status of many system components and detect signs of an impending fault. Recovery mechanisms are triggered to prevent system and application failure.
- **Error identification methods.** Error logging, multistage alerts, electronic FRU identification information, and system fault LED indicators contribute to rapid problem resolution.
- **Power-on Self-test.** Under user control, a power-on self-test (POST) executes to test a wide range of system components. While not intended to perform as a comprehensive diagnostic, POST can quickly establish that no severe problems exist with the system. POST communicates system status through the System Controller and front panel display.
- **Sun Validation Test Suite.** The Sun Validation Test Suite (SunVTS) provides robust online diagnostic capabilities for SPARC Enterprise servers. The primary goal of the SunVTS software is to create an environment where Sun systems are thoroughly tested to support proper operation or to find elusive problems.

Maximizing IT Service Stability

In addition to delivering hardware features that help platforms withstand faults without incurring downtime, Oracle's approach to availability also focuses on creating software technologies that help prevent, predict, and rapidly recover from faults. The self-healing and diagnostic features of Oracle Solaris and failover mechanisms offered by Oracle Solaris Cluster can help further enhance the stability of IT services that run on SPARC Enterprise M-Series servers.

Oracle® Solaris

With over 20 years of engineering investment, Oracle Solaris is one of the most reliable operating systems available. In fact, many organizations can point to systems running Oracle Solaris that execute continuously for months or even years without a restart. Oracle Solaris offers many advanced technologies to help maximize the availability of systems, including the following.

- **Oracle Solaris Predictive Self Healing** software is an innovative feature of Oracle Solaris that proactively monitors and manages system components, facilitating capabilities such as automatic diagnoses, isolation, and recovery from many hardware and application faults.
- **Oracle Solaris Containers technology** supports virtualization of resources within a single image of Oracle Solaris, helping maximize the stability of consolidated environments by supporting secure isolation of independent software workloads.
- **Oracle Solaris Live Upgrade** increases the efficiency of software maintenance procedures, helping minimize planned downtime for operating system upgrade and installation tasks.
- **Oracle Solaris networking technology** automates network connectivity error detection and correction and IP traffic load spreading.
- **Advanced security features** of Oracle Solaris including authentication, data integrity, data privacy, and single sign-on capabilities help to prevent tampering, snooping, and eavesdropping from compromising data or associated transactions.

Oracle Solaris Predictive Self Healing

When failures occur on traditional systems, administrators generally try to make sense of the problem by looking at error messages and system and application log files. Unfortunately, these pre-programmed error messages tend to show only the symptoms and do not necessarily reveal the underlying problem. Administrators are often challenged to make sense of and interpret a stream of errors. Oracle Solaris Predictive Self Healing is an innovative capability in Oracle Solaris 10 that can automatically diagnose, isolate, and recover from many hardware and application faults, helping business-critical applications and essential system services continue operating without interruption. Oracle Solaris Fault Manager and Oracle Solaris Service Manager are the two main components of Oracle Solaris Predictive Self-Healing.

Oracle Solaris Fault Manager

Oracle Solaris Fault Manager is a Oracle Solaris Predictive Self Healing facility that increases system reliability by automatically diagnosing faults in the system and initiating self-healing actions to help prevent service interruptions. This software helps increase availability by configuring problem components out of a system before a failure occurs — and in the event of a failure, this feature initiates automatic recovery and application re-start. The Oracle Solaris Fault Manager diagnosis engine produces a fault diagnosis, once discernible patterns are observed from a stream of incoming errors. Following diagnosis, the Oracle Solaris Fault manager provides fault information to agents that know how to respond to specific faults. For example, an agent designed to respond to a memory error might determine the memory addresses affected by a specific chip failure and remove the affected locations from the available memory pool.

Oracle Solaris Service Manager

Oracle Solaris Service Manager converts the core set of services packaged with the operating system into first-class objects that administrators can manipulate with a consistent set of administration commands. Using Oracle Solaris Service Manager, administrators can take actions on services including start, stop, restart, enable, disable, view status, and snapshot. Service snapshots save a service's complete configuration, giving administrators a way to roll back any erroneous changes applied to a service. Snapshots are taken automatically whenever a service starts to help reduce risk by guarding against erroneous errors. Oracle Solaris Service Manager is integrated with Oracle Solaris Fault Manager. As a result, when a low-level fault is found to impact a higher-level component of a running service, Oracle Solaris Fault Manager can direct Oracle Solaris Service Manager to take appropriate action.

Oracle Solaris Containers Technology

Proper application resource isolation can be critical to realizing high levels of availability for consolidated workloads. Oracle Solaris Containers provides fine-grained resource management capabilities and allows virtualization of an Oracle Solaris image. Each Oracle Solaris Container is assigned a unique root file system, a shared set of system executables and libraries, and a set of compute resources as assigned by the root administrator. Oracle Solaris Containers boot and shut down just like any instance of Oracle Solaris, and reboot in only seconds if the need arises.

Using Oracle Solaris Containers to partition resources to particular tasks can further refine a Dynamic Domain on SPARC Enterprise M-Series servers. Oracle Solaris Containers provide a secure, virtualized environment for application execution while also giving administrators almost unlimited flexibility to assign and isolate resources to particular containers. In fact, Oracle Solaris Containers create a powerful security model enhancing an organization's ability to contain security breaches, limit the effectiveness of attacks, and minimize installation of rogue software such as Trojan horses. Oracle Solaris Containers can also help decrease the administration and maintenance requirements in a consolidated environment by reducing the number of operating system images without giving each application its own instance of Oracle Solaris.

Oracle Solaris Live Upgrade

Traditional operating system upgrades are one of the most time-consuming and error prone tasks facing system administrators. Many mission-critical or datacenter systems cannot afford planned downtime for a day or even an hour to perform software, patch, and operating system upgrades along with necessary testing. Oracle Solaris Live Upgrade promotes greater availability by providing a mechanism to upgrade and manage multiple on-disk instances of the Oracle Solaris — allowing operating system upgrades to take place while the system continues to operate. In fact, Oracle Solaris Live Upgrade can help reduce the planned downtime needed for an operating system upgrade to the time it takes for a simple reboot.

Redundant Networking and Network IP Multipathing

In addition to traditional support for multiple network interfaces connected to different network subnets, Oracle Solaris also provides support for redundant network interfaces that are connected to a single subnet. IP Multipathing provides both failover and IP link aggregation. A number of key features of redundant networking that work to improve the availability and performance of Sun servers are listed below.

- **Failure detection** provides the ability to detect when a network adapter fails and automatically switch (failover) network access to an alternate network adapter.
- **Repair detection** detects the repair of a previously failed network adapter and automatically switch back (fail back) the network access to this interface.
- **Outbound load spreading** distributes outbound network packets across multiple network adapters to achieve higher throughput. Load spreading occurs only when network traffic is flowing to multiple destinations using multiple connections.

Oracle Solaris Security

Oracle Solaris offers an extensive set of built-in security capabilities. The following features of Oracle Solaris work to help organizations secure transaction workloads with less effort and without added cost.

- **Solaris Trusted Extensions** provide labeled security with mandatory access control enforcement for isolating data and applications. The strict data access control offered by Solaris Trusted Extensions drastically minimizes the potential for security breaches and inappropriate use of sensitive enterprise data.
- **Solaris User Rights Management** supports role-based access control over the access rights that both users and applications can exercise by incorporating technology from the time tested Trusted Solaris OS. Solaris User Rights Management helps organizations limit intentional and unintentional security breaches by providing more granular control of superuser capabilities.
- **Solaris Process Rights Management** restricts privileges by pre-assigning rights to user processes, significantly limiting the damage that can be done if the process is somehow compromised by an attack. Solaris Process Rights Management reduces security risks by granting applications only the minimum capabilities needed to perform required duties.

Oracle Solaris Cluster

To address exceptional uptime for mission-critical systems, businesses deploy an increasing number of cluster environments. Designed for high availability as well as manageability, Oracle Solaris Cluster deployments provide the following advantages over single-server solutions.

- **Increased service availability.** Redundant resources in the cluster, coupled with the use of failover mechanisms and scalable services, can help increase application service availability.
- **Greater scalability.** Large SMP servers can be pooled together in a cluster, allowing overall capacity to be expanded either by adding resources to the individual servers (vertical scalability) or by increasing the number of participating servers in the cluster (horizontal scalability).
- **Better resource utilization.** Load balancing and cluster-wide resource management promotes sharing of system resources across multiple application services and allows system resources to be reassigned to other services instead of sitting idle.
- **Improved manageability.** A single management environment for the cluster offers the opportunity to manage components, services, and resources holistically to help improve efficiency and drive down costs.
- **Lower TCO.** The Oracle Solaris Cluster framework can help reduce TCO by simplifying administration. Resources from multiple nodes can be pooled together and administered as a single system. Administrators can access system management tools and Oracle Solaris commands from any system in the cluster. Administrators can add or remove nodes while online and can mix and match servers within the cluster to meet specific needs.

Oracle Solaris Cluster Framework

From a physical perspective, an Oracle Solaris Cluster deployment consists of two or more servers or domains that work together as a single entity to cooperatively provide applications, system resources, and data to users. The configuration of each server generally includes some level of hardware redundancy. In addition, storage is hosted on highly available redundant disk systems that are mirrored, supporting data access in the event of a service interruption on a single disk or storage subsystem. Redundant connections are provided to the disk systems so that data is not isolated in the event of a server, controller, or cable failure. A high-speed, redundant, private interconnect system provides access to resources across the set of servers. Redundant connections to the public network also provide each node with multiple paths for access to outside systems, helping to ensure continued access in the event of a network connection or node failure.

Key to the design of Oracle Solaris Cluster implementations is the fact that no single failure in the hardware, software, interconnect, or network can cause the cluster to fail. Oracle Solaris Cluster solutions prevent loss of service through hardware redundancy, hardware and software failure detection, automatic recovery of services, and failover of applications. Oracle Solaris Cluster software also provides a single management view for all of the services in the cluster. The entire cluster appears as single Sun server, reducing the risk of errors.

From a logical perspective, an Oracle Solaris Cluster deployment consists of a set of domains; where each domain independently executes a copy of Oracle Solaris and communicates through the private interconnect. Utilizing high-speed messaging facilities to build a closely-coupled yet highly-available cluster eliminates the need for domains to share memory. Using fully distributed algorithms, the cluster continuously monitors the health of all members. Failing domains are actively prevented from participating in the cluster to prevent any chance of data corruption.

Oracle Solaris Cluster Software

Available for Oracle Solaris 8, 9, and 10, Oracle Solaris Cluster software extends the capabilities of the operating system to provide enhanced availability of hosted applications. Core Oracle Solaris services, such as devices, file systems, and networks work in a virtually seamless manner across a tightly coupled cluster. While maintaining full Oracle Solaris compatibility for existing applications, Oracle Solaris Cluster software transparently manages shared access to its global file service, global network service, and global devices. Enhanced availability for core services such as file and network service allows existing applications written for Oracle Solaris to benefit from higher availability, more localized failure containment, and faster failover.

Oracle Solaris Cluster supports both scalable and highly-available (failover) services. Scalable services support multiple nodes in a cluster to handle growing numbers of simultaneous users, while failover services can help increase the availability of application services. Oracle Solaris Cluster software includes the following key capabilities.

- **High availability monitoring.** The Oracle Solaris Cluster framework provides extensive monitoring of application processes, disk path integrity, and network availability to detect failures quickly and facilitates migration to another node in the cluster. A distributed set of agents exchange messages over the cluster interconnect, enforcing a consistent membership view, driving synchronized reconfiguration, handling cluster partitioning, and helping maintain full connectivity among all cluster members.
- **Virtualization.** Oracle Solaris Cluster software provides expanded support for Solaris Containers virtualization software. These virtualization capabilities allow scalable and failover applications and associated Oracle Solaris Cluster agents to run unmodified within Oracle Solaris Containers.
- **Global devices, files, and networking.** Oracle Solaris Cluster supports use of globally shared devices, files, and network interfaces as local resources. This capability allows cluster nodes to access and utilize devices that are not physically attached, but instead are attached to another node within the cluster. These facilities create virtually seamless storage and network scalability, as well as improved resource availability and simplified administration.
- **Scalable data management.** With the virtually unlimited scalability of Oracle Solaris Zetabyte File System (ZFS), Oracle Solaris Cluster software offers a file system solution with exceptional availability, data integrity, and flexibility for growth.
- **Application monitoring.** Oracle Solaris Cluster includes dozens of failover and scalable agents for monitoring the continued availability of many Sun and third-party applications.

- **Multisite disaster recovery.** Oracle Solaris Cluster Geographic Edition manages the availability of application services and data across unlimited geographic distances.
- **Protection groups.** Oracle Solaris Cluster software allows an administrator to define groups of resources that must be managed together, facilitating seamless failover of related applications and data.
- **Ease-of-use.** An object oriented command line interface and browser-based configuration wizard ease use and simplify configuration of new data services.

Interconnect Technologies

Clusters rely on interconnects to move information among nodes. Care must be taken to avoid interconnect bottlenecks that can act as obstacles to reaching the potential of otherwise high-performance system designs. The cluster interconnect employed in Oracle Solaris Cluster systems is a physical configuration of independent networks or system interconnections that transfer cluster-private and data service communications between cluster nodes. The Oracle Solaris Cluster software utilizes the interconnect technology for inter-node communications, application level messages, and data and meta-data transfers for global file service and global network service features.

All nodes are connected by the cluster interconnect through at least two physically redundant independent links or paths, eliminating single points of failure. While two links are required for redundancy, Oracle Solaris Cluster software allows up to six to be used to spread traffic, avoiding bottlenecks and improving redundancy and scalability. As a result, the cluster interconnect permits the full utilization of high performance components in cluster systems. A Gigabit Ethernet cluster interconnect provides a cost-effective option that offers reasonable performance during failover and offers effective use of the Oracle Solaris Cluster global devices, global network, and global file system. For Oracle Solaris Cluster implementations that require greater interconnect bandwidth, InfiniBand interconnects are also supported.

Using Agents for Failover or Scalable Services

Agents are software programs that allow applications to take advantage of Oracle Solaris Clusters. Consisting of several methods, these software programs facilitate the starting, stopping, and monitoring of applications on cluster nodes. Combined with cluster framework software and multi-host disks, agents help increase the availability and scalability of applications. Agents help prevent significant application interruption after any single failure within the cluster.

Failover Agents

Oracle Solaris Cluster systems support failover agents for applications and services. When a failure occurs, a running application is either restarted on the same node, or migrated to another node in the cluster without user intervention, depending on application configuration. To accomplish this task, failure services migrate and reconfigure logical hostnames (IP addresses) between cluster nodes. IP addresses are automatically configured down on the original node and up on the replacement node.

Scalable Agents

Oracle Solaris Cluster systems provide scalable agents to support the simultaneous execution of application instances on multiple cluster nodes. Files containing application and network resources (shared addresses) on which the scalable agents depend are used to house the critical information needed to migrate services from one cluster node to another.

All nodes hosting a scalable agent use the same shared address to host the service. Service requests enter the cluster through a single global interface and are distributed to cluster nodes based on a load balancing policy. If the node hosting the global interface fails, the global interface fails over to another node.

Custom Agents

In addition to agents from Oracle and third-party software vendors, new agents can be custom built to allow other applications to take advantage of failover services or scalable services in Oracle Solaris Cluster software. A co-packaged Agent Development Kit provides everything developers need to take full advantage of Oracle Solaris Cluster services. Alternatively, organizations can leverage Sun's expertise and simplify development efforts by contracting Sun Services to build custom agents.

Remote Cluster Services

Local clustering provides a solid level of continuous service availability in the event of application, operating system, or hardware failure in a single datacenter. However, some enterprises require additional protection against disasters such as losing an entire building or campus. In these instances, spreading cluster nodes across multiple geographic locations or providing the capability to failover to a secondary cluster at a remote location makes sense. Oracle Solaris Cluster offers the following options for remote cluster services:

Campus Clustering

Campus clustering essentially extends the distance between cluster components, allowing a cluster system to extend across multiple rooms or buildings in a campus. Campus clustering supports up to 20 kilometers of separation between components, such as nodes and shared storage. In the event of a localized disaster such as a flood, fire, or building power outage, the surviving nodes can support the service for a failed node. This solution offers some site-level tolerance, subject to quorum. However, the short physical distances implied by campus clustering limit the survivability of the cluster and its services for larger disasters, such as earthquakes or grid power outages.

Metro Clustering

For greater availability across an increased distance, cluster nodes can be separated by up to 200 kilometers using Solaris Cluster software and dense wave division multiplexing (DWDM) technology to provide application service continuity in the event of a catastrophic failure. However, a single cluster solution supported across a limited distance can still be affected by a single disaster. For instance, an event similar to the 2003 East Coast power grid failure can impact both sites of a metro cluster running from New York City to Boston. Finally, because the storage systems are shared across the cluster, disk performance is confined by the latency of the connecting links.

Solaris Cluster Geographic Edition

For organizations that must be able to quickly survive disasters that span a wide area, Solaris Cluster Geographic Edition offers a multisite disaster recovery solution that manages the availability of application services and data across unlimited geographic distances. Sun Cluster Geographic Edition is also targeted at organizations that carry out 24 x 7 work around the globe and require a local cluster that replicates the same configuration and data as another cluster located somewhere else in the world. In the event that a primary Solaris Cluster goes down, Sun Cluster Geographic Edition helps an administrator to quickly initialize the business services with replicated data on the secondary Oracle Solaris Cluster.

Customizing Configurations

Organizations can tailor the configuration of domains within SPARC Enterprise M-Series servers to meet specific availability goals. The following examples illustrate the flexibility to use one set of hardware to create solutions with different availability characteristics.

A SPARC Enterprise M8000 server with four CMUs and four IOUs is used for both examples.

- The four SPARC Enterprise M8000 CMUs (CMU0, CMU1, CMU2, and CMU3) each contain the following CPU and memory quantities:
  - Four 2.88 GHz SPARC64 VII processors
  - 32 x 2 GB DIMMs
- The four SPARC Enterprise M8000 IOUs (IOU0, IOU1, IOU2, and IOU3) each contain the following I/O devices:
  - Four 300 GB Disk Drives
  - Two Base I/O Cards (Slots 0, 4)
  - Three 8 Gb FC cards (Slots 1, 3, 6)
  - Three Quad 1 Gb Ethernet cards (Slots 2, 5, 7)

Each example divides the SPARC Enterprise M8000 server into four domains. In addition, each domain supports a single enterprise application and includes four CPUs and 64 GB of memory. Both configurations can also take advantage of IP network multipathing (IPMP) and multiplexed I/O (MPxIO) to create alternate I/O paths that can be utilized in the event that the primary I/O route fails. The major difference between the two configuration examples resides in the fact that the CMUs and IOUs are configured as Uni-XSBs for the first example and Quad-XSBs for the second example.

Example One — Prioritizing Hardware Fault Isolation

In this example, each CMU is placed in Uni-XSB mode and each domain is assigned the full resources of one CMU and one IOU (Table 9). The resulting configuration maximizes fault isolation. Since each CMU and IOU pair only support one domain, a CMU or IOU fault only impacts the application set running in that one domain. Other domains — that utilize separate CMUs and IOUs — continue to run normally. If the failed component is the SC or MAC ASIC, then the domain cannot restart and the

applications in the failed domain remain unavailable. However, Oracle Solaris Cluster software can be utilized to allow another domain or separate server to resume application processing for the failed domain.

TABLE 9. IN EXAMPLE ONE EACH DOMAIN OWNS ALL OF THE RESOURCES OF A SINGLE CMU AND IOU PAIR

DOMAIN	ASSIGNED UNI-XSB	DESCRIPTION
Domain 1	XSB00-0	XSB00-0 contains all the CPUs and memory from CMU0 and all the I/O and disks from IOU0.
Domain 2	XSB01-0	XSB01-0 contains all the CPUs and memory from CMU1 and all the I/O and disks from IOU1.
Domain 3	XSB02-0	XSB02-0 contains all the CPUs and memory from CMU2 and all the I/O and disks from IOU2.
Domain 4	XSB03-0	XSB03-0 contains all the CPUs and memory from CMU3 and all the I/O and disks from IOU3.

Example Two — Simplifying Load Balancing and Resource Allocation Changes

For this example, each CMU is placed in Quad-XSB mode. Each domain is composed of the resources within one Quad-XSB from each CMU-IOU pair (Table 10). Using storage configuration terminology, this is similar to creating a *stripe* across all of the available hardware — each domain relies on a portion of the resources of all four CMU-IOU pairs.

TABLE 10. THE RESOURCES FOR EACH DOMAIN ARE DRAWN FROM MULTIPLE CMU AND IOU BOARD SETS.

DOMAIN	ASSIGNED UNI-XSBS	DESCRIPTION
Domain 1	XSB00-0, XSB01-1, XSB02-2, XSB03-3	Domain 1 contains resources from each CMU and IOU
Domain 2	XSB01-0, XSB02-1, XSB03-2, XSB00-3	Domain 2 contains resources from each CMU and IOU
Domain 3	XSB02-0, XSB03-1, XSB00-2, XSB01-3	Domain 3 contains resources from each CMU and IOU
Domain 4	XSB03-0, XSB00-1, XSB01-2, XSB02-3	Domain 4 contains resources from each CMU and IOU

In contrast to the first example, a hardware failure in any CMU or IOU impacts all domains. However, the advantage of this configuration is the ability to quickly move a single CPU, memory, or I/O device from one domain to another, simplifying load balancing and resource allocation changes. In fact, reconfiguration operations can be completed while applications continue processing. Recovering from a failure of the SC or MAC ASIC requires downtime for hardware replacement or the use of Oracle Solaris Cluster to enable a separate server to take over the workload of all the four domains.

Improving Processes and Operational Efficiency

Putting processes in place that help reduce administrative error and minimize the opportunity to overlook potential system faults is critical to maximizing uptime. System management and monitoring tools for SPARC Enterprise M-Series servers can help organizations automate and optimize IT processes. The extensive built-in management and monitoring capabilities of SPARC Enterprise M-Series servers and tools such as Oracle Enterprise Manager Ops Center and Sun Management Center help organizations streamline procedures, leading to higher levels of availability for IT services. Using the these tools to improve processes can also help increase operational efficiency, lowering overall administrative costs.

Remote Monitoring and Management

Providing hands-on, local administration of systems is no longer realistic for most organizations. With the need to support 24x7 operations, disaster recovery hot sites, and geographically dispersed datacenters, remote management is essential to improving operational efficiency. SPARC Enterprise M-Series servers provide a powerful eXtended System Control Facility (XSCF) to remotely execute and control nearly any administrative task that does not involve physical access to hardware.

The XSCF consists of a dedicated processor that runs independent of the main server and utilizes XSCF Control Package software. The XSCF communicates with the server using the Domain to Service Processor Communication Protocol (DSCP). This protocol runs on a private TCP/IP-based or PPP-based communication link between the service processor and each domain. While input power is supplied to the server, the XSCF constantly monitors the system even if all domains are inactive. The XSCF regularly monitors the environmental sensors, provides advance warning of potential error conditions, and executes proactive system maintenance procedures as necessary. For example, the XSCF can initiate a server shutdown in response to temperature conditions that might induce physical system damage.

The XSCF Control Package running on the service processor helps administrators to remotely control and monitor domains, as well as the platform itself. Using a network or serial connection to the XSCF, operators can effectively administer the server from anywhere on the network. Remote connections to the service processor run separately from the operating system and provide the full control and authority of a system console.

XSCF Control Package

The XSCF Control Package helps users to control and monitor SPARC Enterprise M3000, M4000, M5000, M8000, and M9000 server platforms and individual Dynamic Domains, quickly and effectively. A command line interface (CLI) and Web browser user interface provided by the XSCF Control Package gives administrators and operators access to system controller functionality. Password-protected accounts with specific administration capabilities also provide system security for domain consoles. Communication occurs between the XSCF and individual domains using an encrypted connection based on secure shell (SSH) and secure socket layer (SSL), supporting secure, remote execution of XSCF Control Package commands.

The XSCF Control Package contributes to system reliability, availability, and serviceability by offering the following key functions:

- **Dynamic Reconfiguration** commands facilitate the execution of tasks to logically attach or detach installed system boards from the operating system while the domain continues to run applications without interruption.
- **Domain administration** functions support the creation of logical system boards made up of Uni-XSB and Quad-XSB units.
- **Audit administration** capabilities log interactions between the XSCF and the domains.
- **Power Supply administration** commands allow configuration of the SPARC Enterprise M8000 and M9000 server dual power grid operation as well as the monitoring and controlling of power to the components of SPARC Enterprise M4000, M5000, M8000, and M9000 servers.
- **Hardware status monitoring, reporting, and handling** capabilities provide interpretation of hardware information presented and notification of impending problems such as high temperatures or power supply problems, plus access to the system administration interface.
- **Oracle Solaris Predictive Self Healing** functions coordinate with the XSCF, helping to improve availability through accurate fault diagnosis and predictive fault analysis.
- **Automatic diagnosis and domain recovery** features support the execution and monitoring of diagnostic programs such as the Open Boot Prom (OBP) and POST.
- **Capacity on Demand** operations provide the ability to stage and then later activate additional processing resources.
- **XSCF failover** functions include monitoring of the dual XSCF configuration on SPARC Enterprise M8000 and M9000 servers for failure and execution of an automatic failover as needed.

Flexible Roles for Operators and Administrators

In addition to flexible remote monitoring and management, administrative access control remains an important issue for maximizing system availability. In some environments, even well-intentioned operators can accidently induce downtime. At the same time, inconsistent access control mechanisms can cause the circumvention of established protocols — further jeopardizing availability and security.

The XSCF Control Package simplifies administration of several autonomous domains independently by different system administrators and operators — all cooperating within a single SPARC Enterprise M-Series server. This management software supports multiple user accounts that are organized into groups. Different privileges are assigned to each group. Privileges allow a user to perform a specific set of actions on a specific set of hardware, including physical components, domains, or physical components within a domain. In addition, a user can possess multiple privileges on multiple domains. The specific definitions for each privilege type are listed in Table 11.

TABLE 11. XSCF CONTROL PACKAGE DEFINED PRIVILEGES

PRIVILEGE	CAPABILITIES
NONE	• No privileges
USERADM	• Create, delete, disable, and unable user accounts, user passwords, user privileges • View all platform states
PLATADM	• All service processor configuration other than the useradm and auditadm tasks • Assign and unassign hardware within domains • Domain and Service Processor power operations • All operations on domain hardware • Service processor failover operations on systems with more than one XSCF • View all platform states
PLATOP	• View all platform states
DOMAINADM	• Perform all operations, including hardware changes to domain on which this privilege is held • View all states of hardware assigned to domains on which privilege is held • View all states of the domain on which this privilege is held
DOMAINMGR	• Perform domain power operations • View all states of hardware within domain on which this privilege is held • View all states of the domain on which this privilege is held
DOMAINOP	• View all states of hardware within domain on which this privilege is held • View all states of the domain on which this privilege is held
AUDITADM	• Configure auditing • Delete audit trail
AUDITOP	• View all audit states and the audit trail
FIELDENG	• Perform all operations reserved for field engineers

Platform Management Tools

Sun Management Center and Oracle Enterprise Manager Ops Center software, as well as other third party tools offer advanced management functions that compliment the capabilities of the XSCF Control Package. In order to simplify integration, the XSCF can communicate to system management tools by enabling a Simple Network Management Protocol (SNMP) agent on the service processor. The network interface on the service processor facilitates data transfer to SNMP managers within third-party management applications. SNMP v1, v2, and v3 and concurrent access from multiple SNMP managers is supported.

The service processor SNMP agent can export the following types of information to a SNMP manager:

- System information such as chassis ID, platform type, total number of CPUs, and total memory
- Hardware configuration
- Dynamic reconfiguration information, including the assignment of resources to domains
- Domain status

- Power status
- Environmental status

The service processor SNMP agent can supply system and fault event information using public MIBs. The XSCF supports the configuration of following two MIBs:

- **XSCF extension MIB (SP-MIB)** provides information on the status and configuration of the platform. For fault events, the SP-MIB sends a trap with basic fault information.
- **Fault Management MIB (FM-MIB)** records fault event data. The FM-MIB provides the same detailed information as the FMA MIB in a Solaris domain. This data can help service technicians diagnose failures.

Oracle Enterprise Manager Ops Center

Oracle Enterprise Manager Ops Center software is a highly-scalable datacenter management platform. This software 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 Microsoft Windows, Linux, and Oracle Solaris operating systems. An example of the Oracle Enterprise Manager Ops Center console interface is shown in Figure 14.

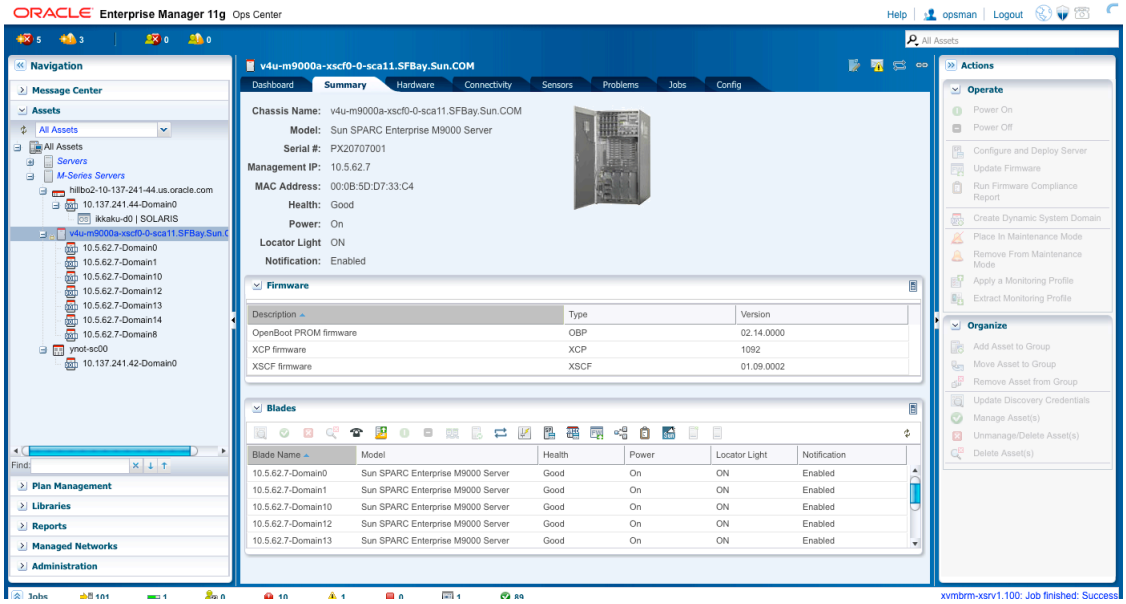

Figure 14. Example view of the Oracle Enterprise Manger Ops Center Console

Oracle Enterprise Manager Ops Center can help enterprises to provision and administer both physical and virtual datacenter assets, and offers the following capabilities:

- **System monitoring.** Oracle Enterprise Manager Ops Center provides monitoring functions while minimizing any incremental network load. Able to detect a system by communicating with its operating system or service processor, Oracle Enterprise Manager Ops Center can provide data about the potential faults and attributes of Sun Fire and SPARC Enterprise servers in a manner that helps operational staff can monitor systems with minimal training. Oracle Enterprise Manager Ops Center also makes it possible to monitor energy use on any capable system, and then aggregate that information relating to groups of servers so that energy use can be associated with a specific application. Power usage is sampled every five minutes, and the hourly average is used to calculate the kWh figure, which is easily mapped to the cost.

  By taking advantage of the capabilities of the service processor, Oracle Enterprise Manager Ops Center allows hardware to be managed centrally—without deploying agents. However, for monitoring at the operating system level, Oracle Enterprise Manager Ops Center uses a software agent, that can be automatically installed from the Oracle Enterprise Manager Ops Center controlling tier.

- **Lifecycle management of both physical and virtual environments.** Oracle Enterprise Manager Ops Center can manage a heterogeneous set of systems (from bare metal through operating system) throughout their lifecycle — from the first time power is applied to the server's retirement. The software lifecycle management capabilities of Oracle Enterprise Manager Ops Center include automation of tasks such as patch downloads, system state backup, patch identification, installation of prerequisite patches, and patch installation. Oracle Enterprise Manager Ops Center can also simplify management of virtual environments through the ability to discover, create, destroy, stop, start, clone, copy, and change the configuration during runtime of an individual VM as well as on logical groupings of systems. Oracle Enterprise Manager Ops Center can help simplify tracking of vendor recommendations or site-specific custom profiles, enforcing policies across the datacenter and providing reports that document, track, and audit compliance measures.

- **Operating system provisioning tasks.** With Oracle Enterprise Manager Ops Center, tasks that once had to be performed on a server-by-server or an operating system-by-operating system basis can now be performed across pools of resources. This not only increases productivity, but it also helps keep environments up to date, including the software and firmware patches that allow datacenters to operate securely. Aspects of system provisioning supported by Oracle Enterprise Manager Ops Center include discovery of datacenter assets using standard mechanisms and protocols, lights-out-management configuration, remote active management, transparent access to multiple operating system provisioning tools within Oracle Solaris, Red Hat Enterprise Linux (RHEL), Oracle Enterprise Linux (OEL), and SUSE Linux Enterprise Server (SLES), and creation of operating system file systems, shares, and services.

Oracle Enterprise Ops Center Architecture

Oracle Enterprise Manager Ops Center is built using a scalable, distributed, and secure architecture. The three-tier architecture uses a centralized controller along with distributed proxies that keep data local and secure. This architecture improves performance by minimizing the amount of data that must be transmitted between the proxy controllers and the enterprise controller features in Oracle

Enterprise Manager Ops Center. In addition, this architecture allows Oracle Enterprise Manager Ops Center to work within complex network and firewall configurations as well as application silos that might be distributed around the world.

When used in conjunction with SPARC Enterprise M3000, M4000, M5000, M8000, and M9000 servers, 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. Oracle Enterprise Manager Ops Center can also be used to discover the embedded service tag technology in the service processor and domains running Oracle Solaris on SPARC Enterprise M3000, M4000, M5000, M8000, and M9000 servers.

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 help insure accountability and auditing. These automation capabilities can be used for knowledge-based change management in conjunction 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.

Sun Management Center

For legacy environments still using Sun Management Center software from Oracle for server management, the SPARC Enterprise M-Series servers provide a Sun MC agent on each XSCF. Sun Management center presents a logical view of Sun systems, where a logical system unit is defined as a single Sun server or as a domain within a server. To support legacy networks and heterogeneous environments, Sun Management Center software tightly integrates with major management frameworks including CA Unicenter TNG, HP Open-View, IBM Tivoli, and BMC Patrol.

For SPARC Enterprise M3000, M4000, M5000, M8000, and M9000 systems, Sun Management Center software retains the model of platform and domain management established by the XCSF and XSCF Control Package. By synchronizing permissions and access control lists, the same capabilities are available to Sun Management Center software users as to users of the command-line interface of the XSCF. Administrators and operators can be given only the access needed to perform tasks.

Sun Management Center software can also be used to manage Dynamic Reconfiguration within SPARC Enterprise M4000, M5000, M8000, and M9000 servers, including domain creation, dynamic board attach, and dynamic detach. (The actual attach and detach operations take place within the affected domain, while Sun Management Center software provides the GUI for initiating these activities.) Sun Management Center software is able to communicate with the XSCF to retrieve status information not only for active domains, but also for boards not currently assigned to any domain — an important requirement for managing the deployment of resources between domains.

Optimizing the Environment with Oracle Services and Training

Reaching 99.99% uptime allows for only an hour of downtime per year. Pushing availability levels a bit further to 99.999% uptime narrows the downtime window to little more than 5 minutes per year. As the room for error decreases, the contributions that people and processes make to availability levels rises in importance. Professional and education services from Oracle can help organizations fine-tune processes as well as train or supplement IT staff. By leveraging Oracle's extensive expertise, best practices, and innovative managed services offerings, enterprises can reduce potential risks and enhance the stability and performance of IT services.

Each Oracle customer engagement for Sun systems includes:

- Working collaboratively with internal teams and partners to share learning
- Addressing business through technology investments
- Accelerating project implementations while reducing risk
- Establishing highly-effective and repeatable datacenter processes based on best practices and automated solutions
- Building in-house expertise through both knowledge transfer and training

Oracle's portfolio of service and support offerings includes technology consulting, educational services and training programs, and service plans for Sun hardware and software solutions. Proactive monitoring and management including both remote management and managed services offerings are also available. Oracle's worldwide service and support organization is complemented by a broad network of service and support partners. Additional details on these offerings are described below and on the following Web site: *http://www.oracle.com/us/support*.

Oracle's Suite of Professional Services

Oracle's professional services accelerate the realization of important business goals through the rapid and seamless deployment of innovative Sun products and technologies. By leveraging Oracle's expertise, new servers can be brought into production quickly with predictable results for reliability and availability. Using this expertise can also help IT staff remain focused on strategic priorities rather than non-core activities such as deploying IT infrastructure components.

The Oracle professional services portfolio for Sun systems consists of services to help enterprises reduce technology adoption risk and optimize investments in Sun hardware and software products. These services not only help relieve IT staff from non-strategic tasks, but also provide a foundation for enhanced system stability and performance. Table 12 offers a list of professional services offered by Oracle and its partners. Detailed information on these offerings is available at *http://www.oracle.com/us/support/systems/advanced-customer-services*.

TABLE 12. PROFESSIONAL SERVICES OFFERINGS

PROFESSIONAL SERVICES	DESCRIPTION
DATACENTER CONSULTING SERVICES	Datacenter Consulting services help organizations plan, implement and manage an

	optimized datacenter environment. Oracle's extensive portfolio of datacenter consulting services cover everything from optimizing existing IT infrastructure to helping plan and build a new datacenter to meet escalating demands. This service can assist with planning and implementation of an end-to-end datacenter infrastructure strategy across compute, networking, and infrastructure components.
SERVICE DEFINITION WORKSHOP	The Sun Service Definition Workshop provides an in-depth assessment of current IT and business goals, and establishes a clear and executable plan to assure that objectives are properly aligned to achieve long term success. Utilizing an interactive workshop format, consultants work with IT leadership to understand the challenges within the current IT environment, articulate the operational vision that best meets the organization's business needs, and develop an action plan that is strongly connected to and driven from actual business requirements for getting there.
OPERATIONS CAPABILITY ASSESSMENT SERVICE	The Operations Capabilities Assessment Service benchmarks an organization's ability to efficiently manage internal IT operations. Well defined criteria and metrics are used to assess the capabilities of people, process and tools. This service can include a complete review of the IT environment or a focused assessment concentrating on specific critical areas.
IT SERVICE MANAGEMENT OPTIMIZATION SERVICE	As the most widely accepted approach to IT service management in the world, the IT Infrastructure Library (ITIL) provides a cohesive framework of best practices, drawn from the public and private sectors internationally. Implementing ITIL processes is key to improving the efficiency and value of IT services. The IT Service Management Optimization Service builds on the established ITIL framework to deliver a rapid, easy-to-use, tailored solution. The accumulated knowledge and experience of Oracle IT service management experts is captured in a customized knowledge base to provide specific, action-oriented processes. These processes can then form the basis of organizational change and training, service management tools customization, and ongoing improvement.
DYNAMIC INFRASTRUCTURE SUITE	The Dynamic Infrastructure Suite of Services helps organizations respond quickly to changes in market demand by facilitating creation of a flexible, efficient, and secure IT infrastructure. As a result, organizations can add, retire, or expand business services up to 80% faster, without new systems overhead.
ENTERPRISE INSTALLATION SERVICES	The Enterprise Installation Service offers fast, basic server set up by experienced engineers. Standardized installation methodologies are applied based on best practices to ensure systems are properly configured and that the latest patches and updates are installed so systems will run at an optimal level right from the start.

SUN IRUNBOOK SERVICE	The Sun iRunbook Service includes the compilation and set up of a customized, Web-based central repository of streamlined technical "how to" information based upon Sun preferred practices and IT processes. The service aims to facilitate more efficient datacenter management, fewer incidences of error, faster problem-resolution, widespread adoption of preferred practices, and increased staff technical knowledge.
APPLICATION MIGRATION SERVICE	The Application Migration Service helps organizations maintain version and platform consistency and provides a tested methodology for a successful migration to Oracle Solaris. The service support efforts to migrate applications from many UNIX® or Linux variants, making the application ready to compile and run on Oracle Solaris 10 or earlier versions.
APPLICATION READINESS SERVICE	The Application Readiness Service helps ensure that a server environment is correctly designed and application-ready prior to deployment. Well-designed server environments help to support faster system implementation, prevent unnecessary downtime, and reduce the risk of system configuration errors.
ORACLE SOLARIS CLUSTER IMPLEMENTATION SERVICE	Oracle Solaris Cluster Implementation Service helps organizations take full advantage of Oracle Solaris Cluster software by designing and testing cluster environments to meet high availability requirements. By using this service, organizations can deploy cluster implementations with confidence, knowing the solution has been tested and verified to meet specific requirements.
SUN PERFORMANCE ANALYSIS & CAPACITY PLANNING SERVICE	Sun Performance Analysis & Capacity Planning Service evaluates an organization's server environment to determine if current assets are sufficient to meet expected growth and business needs, whether additional hardware will be required, or if there are opportunities to consolidate workloads from old or under-used servers.
ENTERPRISE MIGRATION SUITE	Sun Enterprise Migration Suite helps organizations safely migrate to a new IT infrastructure or upgrade to Oracle Solaris 10 while remaining focused on critical business needs. Consultants analyze business objectives, project objectives and possible constraints, including partnerships, resources, road maps, dependencies, and operational requirements, and then demonstrate the benefits, value and risks of the proposed migration.

Management Services

Oracle's Sun Management Services provide dynamic IT infrastructure management that optimizes datacenter operations while helping enterprises retain control over IT program direction and strategy. Built on the principles of flexibility and choice, Sun Management Services aim to reduce IT operational cost through best-in-class managed services. A single point of contact between the customer and service providers helps cut management complexity. Sun Management Services are also anchored in widely-recognized principles for best practices in IT service management, providing structured methods for the continuous evaluation and improvement of IT services.

Many enterprises realize the best results by remaining in control of their IT programs when leveraging the competencies of outside vendors. Business conditions change quickly and the best way to use outside services may change over time. Sun Management Services provide organizations great freedom in modifying sourcing arrangements as business requirements evolve.

The Sun Management Services portfolio includes services ranging from packaged offerings of relatively fixed scope to custom engagements constructed from a wide range of capabilities. Custom engagements are based on service level agreements that define the custom services to be delivered and an objective means of measuring service delivery success. Table 13 describes the packaged offerings from Sun Management Services that can be linked together and custom-tailored to configure a unique solution for each situation. The flexibility offered by Sun Management Services can help organizations improve the efficiency of IT operations and overcome IT skills shortages without giving up control of IT programs and direction.

TABLE 13. SUN MANAGEMENT SERVICES OFFERINGS

MANAGEMENT SERVICE OFFERING	DESCRIPTION
SUN ONSITE OPERATIONS MANAGEMENT	Sun Onsite Operations Management is a customized service providing dedicated resources to help ensure the success of mission-critical environments. Sun technical experts can help reduce operational risk during the implementation of new solutions or simply provide access to key skills sets as needed.
SUN REMOTE OPERATIONS MANAGEMENT	Sun Remote Operations Management offers heterogeneous and remote monitoring and management for all components of an IT infrastructure, including the operating system, third-party, and custom applications, databases, networks, security, and storage. This service creates customized solutions based on the components selected for monitoring and management. In addition, highly-specialized expertise is available to support remote operations management of storage and identity services.
MANAGED HOSTING FROM SUN	Managed hosting capabilities provide network and site security, with protection against fire, earthquakes, flooding, and other natural disasters. In addition, biometric access control options and bullet proofing are available, if needed. Systems are housed in private cabinets, dedicated cages, or private suites. Facilities support 24x7x365 services with redundant power, cooling, and network topologies with no single point of failure and bandwidth expansion capabilities. Service level agreements for power, cooling, and network availability, along with extensive reporting are also available.

Educational Services

People and process contribute heavily to application service availability, making prioritization of IT skills assessment and training important. Oracle's learning management systems and technical learning solutions can help keep IT staff skills up-to-date, accelerating individual performance while increasing productivity.

Many organizations that purchase systems such as SPARC Enterprise M-Series servers engage in large projects that may require additional skills across the organization. It is advisable to begin these types of projects with an assessment of the people assigned to new roles and an evaluation of their current and desired skills sets can. From this starting point, a tailored learning plan can be developed and applied.

The Oracle learning portfolio includes services ranging from up-front analysis and skills assessment to training and certification. Before a project begins, Oracle can help identify skills gaps across the organization by reviewing organizational structure and assessing staff skills in each area. Individual skills can then be built through a variety of training modalities as described below. Finally, Oracle offers certification for the knowledge and skills acquired through formal training and on-the-job

experience as a way to measure success, giving individuals and managers confidence in newly-developed skills.

A Variety of Learning Modes

To help IT managers design, staff, train, and manage a high-performance IT team, Oracle's training solutions include a variety of formats for learning, including:

- **Traditional classroom training.** Oracle traditional classroom training facilities around the world provide access to expertise and courseware in a classroom setting that provides ready access to the technologies explored in the course. Oracle's courseware includes skills training and certifications in leading-edge technologies developed by Oracle.
- **On-site training.** Oracle brings expert instructors on-site to train an organization's IT staff, eliminating the need for employee travel and allows in-tact teams to share in training programs.
- **Customized training.** Oracle provides customized training content and delivery to meet unique project requirements.
- **Online learning.** Individuals or teams of IT staff can access self-paced Web-based content or utilize Live Virtual Class training that combines the traditional strengths of classroom training with the ease of use, and global accessibility of the Internet.
- **CD-ROM courses.** Courses delivered via CD-ROM are an excellent way to provide core competency training for individuals without requiring travel or extended time away.

Certification

The experience, training, and study required to achieve a technical certification increases an employee's skills and knowledge, thus improving self-confidence and productivity. From the individual perspective, certification increases the chance of salary and career advancement. Organizations gain value in a variety of ways. By having Oracle certified professionals on staff, system integrators and software consulting firms can gain credibility in the eyes of their customers, thus improving the opportunity for increased revenue and additional contracts. Certified staff can also give IT managers a greater level of confidence in setting goals for the achievement of high service levels with maximum system performance and availability.

Service and Support

Keeping hardware and software systems up-to-date and running smoothly is critical to application service availability. Oracle offers a range of ongoing management and support services to help IT organizations meet service level agreements and help minimize Total Cost of Ownership (TCO).

Oracle Premier Support for Systems

Oracle Premier Support for Systems provides organizations with complete support for Oracle servers and data storage systems. Organizations can depend on Oracle's expanded technical assistance and comprehensive global support services to help maximize the return on IT investments and minimize risk.

Oracle Premier Support for Systems includes the following services:

- **Complete system support**
  - Fully-integrated support for server hardware, OS, and virtualization software
  - Mission-critical storage system support
- **24/7 technical support**
  - Access to Oracle experts
  - Remote diagnostic tools
  - Lifetime Support for OS and virtualization software
- **24/7 online resources**
  - Oracle knowledgebase
  - Product documentation
  - Oracle support communities
  - Best practices
- **24/7 hardware service coverage**
  - Two-hour Onsite Hardware Response[3]
  - Auto Service Request
  - Experienced field engineers
  - OEM replacement parts
  - Field Change Orders (FCOs)

[3] System must be located within an Oracle two-hour service coverage area to receive two-hour response as a standard service.

- **Proactive support tools**
  - Support alerts
  - Software update tools
  - Security resources
  - Sun Services Tools Bundle
  - Embedded diagnostic tools
- **Operating System, Firmware and Embedded Software updates**
  - Oracle Solaris, Oracle Enterprise Linux, Oracle VM, and Oracle hardware
- **System Lifecycle Policy**
  - Receive support for your Sun hardware systems for 5 years from last shipment date

Oracle Premier Support for Systems offer oragniztions the following benefits:

- **Immediate access to deep expertise.** Receive fast, proven solutions gained from supporting thousands of active customer installations, from applications to disk.
- **Lower Risk.** Unlimited, 24/7 access to systems specialists and valuable online technical resources help organizations proactively avoid business interruptions and enhance security.
- **Greater Performance.** Optimize IT environments with total-system server support featuring fully-integrated hardware and OS coverage.
- **Increased Data Availability.** Enterprise storage system support protects data, helping organizations gain the highest business value.

For more information on Oracle Premier Support for Systems please see:
http://www.oracle.com/us/support/systems/premier/index.html

Oracle Premier Support for Operating Systems

Oracle Premier Support for Operating Systems if for organizations that are running Oracle Solaris, Oracle Enterprise Linux, Oracle VM or any combination of these products and have opted not to purchase complete system coverage. This support program provides award-winning support, access to new software innovation, and proactive support tools for any or all of these three leading software products for one great price. Organizations can protect technology investments and keep business operations running effectively and efficiently with Oracle Premier Support for Operating Systems. For more information on this support program please see:
http://www.oracle.com/us/support/systems/operating-systems/index.html

Oracle Advanced Customer Services for Systems

The availability of IT systems and their efficient operation are key to business survival. Oracle Advanced Customer Services for Systems provides deep know-how in deploying, managing, and

optimizing Sun systems from Oracle—improving performance, increasing availability, and reducing implementation times.

Oracle Advanced Customer Services for Systems delivers support and guidance to improve the implementation and operation of Sun systems. This support program provide a range of mature services addressing specific needs. With onsite as well as remote delivery capabilities, Oracle Advanced Customer Services for Systems gives offers access to the skills and expertise organizations need.

Oracle Advanced Customer Services for Systems offerings include the following:

- **Oracle Expert Service for Systems.** Tap the knowledge of Oracle experts to configure and optimize system environments to meet unique organizational needs. Experts provided by this service can customize a solution, manage administrative tasks, and oversee the logistics and system requirements of growing and changing IT environments.
- **Oracle Operations Management for Systems**. Leverage the deep experience gained from supporting thousands of systems around the world, 24x7. As a part of this service, highly skilled experts in technology and operations, using exclusive Oracle tools, help keep IT infrastructure running smoothly.
- **Oracle Packaged Services for Systems.** Packaged services delivered by Oracle system experts can help organizations speed time to deployment and improve system availability and efficiency.

For more information on Oracle Advanced Customer Services for Systems please see:
http://www.oracle.com/us/support/systems/advanced-customer-services/index.html

Summary

Availability levels for business-critical systems are the result of the interaction of people, process, and product over the life cycle of a project. By taking a comprehensive approach to addressing availability, organizations can improve uptime while also gaining a wide array of business benefits, including greater operational efficiency, lower costs, and better ROI.

Oracle products and services bring the following three major strengths to help enterprises build cost-effective and highly available solutions:

- **Highly reliable server architectures.** Oracle's SPARC Enterprise M4000, M5000, M8000, and M9000 servers, Oracle Solaris, and Oracle Solaris Cluster provide a broad range of reliability, availability, and serviceability features to maximize uptime and minimize recovery time. These capabilities are the result of many years of continuous improvements in hardware, operating system, and software technologies.
- **Automation to solidify system management processes.** Sun management tools can help organizations improve adherence to known best practices and reduce the complexity of supporting and managing highly-available application services. Oracle's SPARC Enterprise M-Series server built-in management capabilities, Sun Management Center software, and Oracle Enterprise Manager Ops Center software help facilitate effective and efficient management of large numbers of servers.
- **Proven expertise and support.** Oracle's professional services, innovative education offerings, support tools, and proactive monitoring and management solutions can help organizations architect, implement, and manage better IT solutions that maximize availability levels and reduce costs.

For More Information

For more information on Oracle's SPARC Enterprise M-Series severs and related software and services from Oracle please see the references listed in Table 14.

TABLE 14. REFERENCES

SPARC Enterprise Servers	http://www.oracle.com/us/products/servers-storage/servers/sparc-enterprise
Oracle Solaris	http://www.oracle.com/us/products/servers-storage/solaris
Oracle Solaris Cluster	http://www.oracle.com/us/products/servers-storage/solaris
Oracle Enterprise Manager Ops Center software	http://www.oracle.com/us/products/enterprise-manager/opscenter
Oracle Support	http://www.oracle.com/us/support/index.html

High Availability for Business-Critical
IT Services Using Oracle's SPARC
Enterprise M-Series Servers
April 2011, Version 1.2

Oracle Corporation
World Headquarters
500 Oracle Parkway
Redwood Shores, CA 94065
U.S.A.

Worldwide Inquiries:
Phone: +1.650.506.7000
Fax: +1.650.506.7200
oracle.com

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

Hardware and Software, Engineered to Work Together