HIGH AVAILABILITY IN THE DATACENTER
SUN™ SPARC® ENTERPRISE™ SERVERS

White Paper
April 2007

Table of Contents

Chapter 1

Higher Availability for Lower Total Cost of Ownership

Enterprises increasingly rely on the availability of IT systems in order to effectively compete in today's fast changing global markets. As a result, improving business-critical services is a major focus for IT organizations. The following key factors heavily influence modern system operations.

- *Downtime in critical systems results in missed market opportunities, revised customer loyalties, and even failed companies*. As businesses embrace network computing and move mission-critical applications online, availability emerges as the number one design criteria for IT architectures. Unavailable systems can compromise product quality and schedules, wasting valuable staff time and causing products to miss market windows.
- *The dynamic nature of the Internet leads to unpredictable user workloads that may compromise essential systems.* IT managers attempting to react to sudden changes in server throughput requirements often find traditional support and upgrade mechanisms are no longer sufficient. Systems must adapt quickly, enabling organizations to readily capitalize on business opportunities.
- *Due to changes in the nature of software architectures, organizations increasingly think in terms of availability of IT services.* The creation of modular software delivered over the network shifts emphasis away from the reliability of specific system components to the availability of end-to-end services. Users want fast, reliable execution of IT services to help accomplish objectives, driving the formation of complex service-level agreements between companies, partners, and customers.

In addition, datacenter managers continually face requirements to do more with less. Year after year, standards for system availability levels rise and budgets continue to tighten. IT managers must operate current deployments and invest in new projects that can provide competitive advantage, all while keeping total cost of ownership (TCO) low and maximizing the return on investment (ROI).

In evaluating TCO, enterprises must not only look at the immediate costs associated with any new or ongoing IT project, but also consider the impact of IT investment decisions on the entire enterprise. A working definition of TCO includes the short-term acquisition costs of hardware and software, as well as the larger, secondary costs that result from selecting a particular technology implementation. TCO definitions vary, but nearly all conclude the short-term, immediate costs of acquiring hardware and software represent a fraction of the broader, long-term costs of operating the asset.

Sun understands availability and TCO are the result of a complex relationship between people, processes, and products (Figure 1-1). Dedicated to improving availability and TCO, Sun provides a comprehensive set of technologies, products, programs, and initiatives to increase service levels while reducing cost and risk.

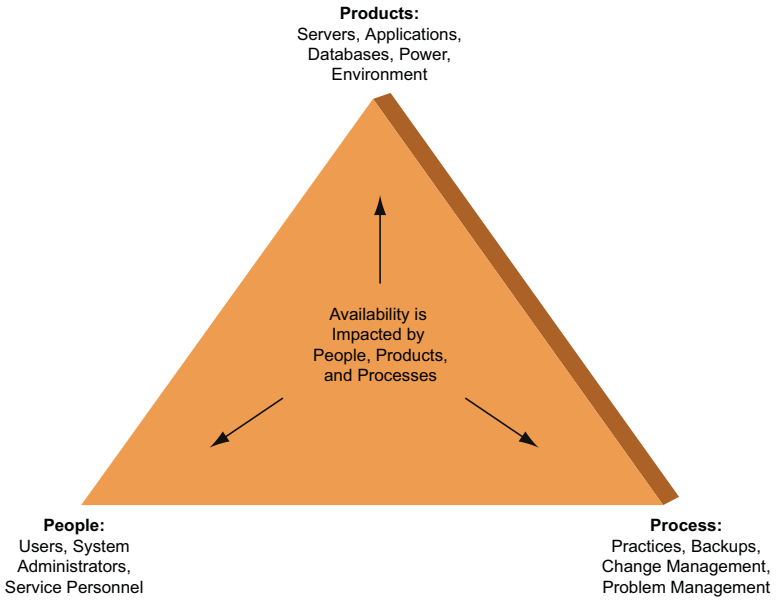

Figure 1-1. To effectively reduce unplanned downtime, organizations must focus on the impact made by people, processes, and products.

By focusing on the following three key elements of IT solutions, Sun delivers technologies and services that help businesses achieve high availability while lowering TCO.

- *Server architectures* — The reliability features of the Sun™ SPARC® Enterprise™ server line bolster the availability levels of hosted applications. Design elements which contribute to uptime include redundant and hot-swappable components, end-to-end data protection, advanced CPU and memory error detection and recovery, Dynamic System Domains, and state-of-the-art fault isolation.
- *Training and processes* — The actions of system administrators contribute to the availability levels of systems. Unfortunately, human error is unavoidable. As complexity increases, people often misunderstand the importance of critical system elements and the ramifications of administrative processes on system uptime. Sun works to reduce both the occurrence and impact of human error by providing greater technical expertise and training along with advanced management tools that help administrators follow defined procedures.
- *Automation* — Technologies that automate management procedures not only help prevent errors, but also provide efficiencies that enable system administrators to manage a greater number of servers.

Sun's solution portfolio — Sun SPARC Enterprise servers, key software technologies, cluster solutions, training programs, and service offerings — helps enterprises achieve high levels of availability without breaking budgets.

Sun™ SPARC® Enterprise™ Servers — A Reliable Foundation

Reducing downtime — both planned and unplanned — is critical for IT services. Business success often depends upon the continuous availability of applications and systems. As a result, system designs must include mechanisms that foster fault resilience, quick repair, and even rapid expansion, without impacting the availability of key services. As systems become more reliable, unplanned downtime decreases as a proportion of overall downtime. Today, planned downtime represents approximately 80 percent of all service interruptions in a typical organization. Decreasing planned downtime through best practices and system features that enable online upgrades and repairs often provides the greatest improvement in software service availability. A simplified administrative environment contributes to management efficiency and automates tasks, helping prevent process errors.

Sun SPARC Enterprise servers are highly reliable, easy to manage, vertically-scalable systems with all of the benefits of traditional mainframes and none of the associated cost, complexity, or vendor lock-in. In fact, Sun SPARC Enterprise servers deliver a mainframe-class system architecture at open systems prices, and run the robust, flexible Solaris™ 10 Operating System (OS). Furthermore, the range of compute power offered by the Sun SPARC Enterprise server family provides the levels of vertical scalability required for server consolidation and many other deployment classes, enabling enterprises to match the right system to the job at hand. Sun SPARC Enterprise M4000 and Sun SPARC Enterprise M5000 servers fulfill mid-range system requirements, while Sun SPARC Enterprise M8000 and Sun SPARC Enterprise M9000 servers deliver the massive processing power needed for high-end computing (Figure 1-2).

Figure 1-2. Sun SPARC Enterprise servers are a scalable family of reliable and high-performance symmetric multiprocessing (SMP) servers that deliver low TCO.

The Sun SPARC Enterprise server family scales from 1 to 64 processors and can incorporate up to 2 TB of memory. These servers are powered by the new SPARC64 VI processor which combines chip multi-processing (CMP) and vertical multi-threading technology (VMT) to enable four threads to run in parallel. This innovative CPU design helps Sun SPARC Enterprise servers offer significantly better performance than competing systems. At the same time, these servers offer full binary compatibility and complete investment protection for owners of previous generations of Sun systems.

Specifically designed to support complex, networked computing solutions and stringent high availability requirements, the systems in the Sun SPARC Enterprise server family 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. The advanced architecture of these reliable servers

enables high levels of application availability and rapid recovery from many types of hardware faults.

Sun SPARC Enterprise servers provide built-in features to support enterprise virtualization and consolidation strategies, easing operations and improving flexibility. Dynamic System Domains and Dynamic Reconfiguration (DR) technology provided with Sun SPARC Enterprise servers simplify resource sharing and create efficiencies that help reduce TCO. In addition, Sun SPARC Enterprise servers introduce eXtended System Board technology, enabling administrators to assign compute and memory resources to Dynamic System Domains with a finer level of granularity than previous generations of Sun servers. Dynamic System Domains and Dynamic Reconfiguration technology increase availability by supporting flexible and agile resource management, allowing rapid response to changing business conditions by quickly shifting resources to systems in need of additional compute power.

Sun SPARC Enterprise servers are ideal for server consolidation, mainframe rehosting, scientific and high performance computing applications, and many other software services. These servers provide the availability features needed to support commercial computing needs along with the raw performance demanded by the high performance computing community.

Sun℠ Upgrade Advantage Program

Continued reliance on old systems can diminish an organization's ability to meet service level agreements, yet some organizations spend more to maintain older systems than the cost of a system upgrade. The Sun℠ Upgrade Advantage Program offers an easy way for organizations to adopt new technology, such as Sun SPARC Enterprise servers, without breaking budgets. When combined with other Sun services, such as Sun Microsystems Finance lease packages and Sun Enterprise Services extended warranties, the Sun Upgrade Advantage Program delivers a complete solution for bringing computing infrastructures up-to-date at minimum cost.

Every aspect of the Sun Upgrade Advantage Program reduces the cost, risk, and complexity of transitioning to new technology. The following features of the Sun Upgrade Advantage Program help enterprises migrate to the latest technology from Sun.

- *Aggressive trade-in allowances* — Whether trading in Sun or non-Sun products, Sun offers excellent value for previous generation systems and components.
- *Up front savings* — Sun delivers savings up front in the form of a percentage taken off list price.
- *90 day return policy* — Sun reduces risk by giving organizations the time and flexibility to test and migrate to a new platform.

- *Free shipping and safe disposal* — Sun covers the cost to ship systems from a customer's dock to the Sun warehouse, and disposes of trade-in assets in an environmentally safe way.
- *Consolidation options* — Sun makes deploying applications and services onto fewer servers more cost-effective than ever before by providing financial incentives to consolidate onto the latest Sun products.

Consulting Services and Training — Improving Processes

Delivering highly available application services requires new ways of thinking about the configuration, deployment, and management of systems. Organizations must design administrative processes which minimize operational impact.

Sun works directly with its customers and partners to assess overall requirements and institute methodologies and strategies that achieve the level of availability organizations demand. The servers in the Sun SPARC Enterprise family can be purchased with Sun Services offerings that include all of the core services and methodologies recommended to provide optimum levels of availability. Sun blends consulting, training, and services to deliver end-to-end support at every phase of the lifecycle. Sun offers many services that enable businesses to leverage the expertise of Sun technical consultants, learning specialists, and support professionals to architect, deploy, and manage a successful implementation.

Beyond professional services, education plays a key role in optimizing application performance and availability. Sun offers a full suite of education curriculum and training services to help organizations develop the skills needed to build and maintain secure, scalable, and highly available business solutions. With one of the world's largest UNIX® system training organizations, Sun trains over 85,000 people every year on the effective use of Sun products.

Automation — Reducing Errors and Lowering TCO

From the Solaris Operating System (Solaris OS) to a host of powerful management applications, Sun supplies tools that provide critical automation for administrative tasks and simplify the management of complex environments. The following automation tools help reduce the degree of human interaction and improve procedural accuracy, resulting in higher application availability and more efficient operations.

- The Sun SPARC Enterprise server eXtended System Control Facility provides built-in automated management and monitoring.
- Sun N1™ Service Provisioning System simplifies application life cycle management by rapidly provisioning new business services and enabling quick replication of a complete software stack on servers throughout the IT infrastructure.

- Sun N1 Grid Engine 6 software optimizes utilization by intelligently mapping compute-intensive workloads onto networked resources based on predefined policies.
- Sun™ Management Center software provides a flexible graphical user interface (GUI) for monitoring, administration, and management for all of Sun systems, enabling true remote system management and lights out operation. Sun Management Center software also integrates easily with third-party enterprise management tools.
- Solaris Cluster software facilitates scalable and highly available applications and enables administrators to manage and administer complete end-to-end services in clusters of up to eight nodes.
- Sun Update Connection is designed to maintain an accurate and up-to-date IT infrastructure while providing maximum flexibility for accomplishing system updates. Building on previous patch management tools from Sun, the Sun Update Connection service integrates all aspects of patch management into a seamless architecture that simplifies the process of staying up-to-date and secure with the latest software updates from Sun.

Chapter 2

Sun SPARC Enterprise Servers — Architected for High Availability

Balancing uptime requirements against hardware acquisition, deployment, and maintenance costs contributes to keeping TCO in line for server deployments. Systems must maximize availability while not introducing prohibitive acquisition costs, complexity, or proprietary architectures.

System uptime depends not only on seamless recovery from failures and rapid service times, but also on the ability of administrators and operators to anticipate and respond to changing situations. Recognition of these factors provides the basis for Sun's approach to availability. As such, the Sun SPARC Enterprise server design features listed below work together as a part of a comprehensive and integrated architectural approach designed to reduce downtime and improve the availability of key systems at lower total costs.

- *Built-in hardware redundancy and hot-swappable components* in key subsystems lay the foundation for system resiliency and increased serviceability.
- *Dynamic System Domains* enable a single system to be physically divided into multiple electrically isolated partitions, each running independent instances of the Solaris OS. Hardware or software failures in one Dynamic System Domain do not affect applications running in other domains, unless the failed resource is shared across both domains.
- *eXtended system boards* allow more granular, sub-system board assignment of individual CPU, memory, and I/O components to Dynamic System Domains, increasing flexibility and optimizing resource utilization.
- *Dynamic Reconfiguration* enables multiple simultaneous changes to hardware configurations without interrupting critical systems, easing maintenance.
- *Reduced component count and complexity* within the server architecture contributes to better overall stability and reliability of the platform.
- *Advanced CPU integration and guaranteed data path integrity* provide for autonomous error recovery by the SPARC64 VI processor, reducing the time needed to initiate corrective action and subsequently increasing uptime.
- *Memory patrol, chipkill, and memory mirroring* all work to enable system operations to continue uninterrupted in the face of memory subsystem error conditions.
- *End-to-end data protection* detects and corrects errors throughout the system, ensuring complete data integrity.

- *State-of-the-art fault isolation* identifies faults at the chip level rather than the component level, enabling execution of offline operations on just the errant Application Specific Integrated Circuit (ASIC), improving system resilience.
- *Remote monitoring and management* capabilities enable lights-out management and proactive administration, providing considerable flexibility.

Built-in Hardware Redundancy and Hot-swappable Components

Today's IT organizations are challenged by the pace of non-stop business operations. In a networked global economy the enterprise revenue stream remains active around the clock, forcing planned downtime windows to shrink and in some cases disappear entirely. To meet these demands, the Sun SPARC Enterprise family of servers employ redundant and hot-swappable hardware to help mitigate the disruptions caused by individual component failures or changes to system configurations (Table 2-1). In fact, these systems are able to recover from hardware failures — often with no impact to users or system functionality.

Table 2-1. The redundant and hot-swappable components in the Sun SPARC Enterprise server family increase system availability and ease maintenance.

	Redundant Components	**Hot-Swappable Components**
Sun SPARC Enterprise M4000 servers and Sun SPARC Enterprise M5000 servers	• 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 • PCI Cards
Sun SPARC Enterprise M8000 servers and Sun SPARC Enterprise M9000 servers	• Disk Drives • Power Supply Units • Fan Units • Dual Power Feed (optional) • CPU Memory Board Units • I/O Units • External I/O Expansion Unit (optional) • eXtended System Control Facility Units • Crossbar Units (M9000 only)	• Disk Drives • Power Supply Units • Fan Units • CPU Memory Board Units • I/O Units • External I/O Expansion Unit • PCI Cards • DVD-ROM drive • Tape drive • eXtended System Control Facility Units

If a fault occurs, corrective action is taken by the system. Depending upon the component and type of error, the system may continue to operate in a degraded mode or may reboot — with the failure automatically diagnosed and the relevant component automatically configured out of the system. Indeed, Sun SPARC Enterprise servers recover quickly from many component failures, including serious faults such as the failure of a CPU or a critical ASIC. In fact, no single hardware component failure prohibits the Sun SPARC Enterprise M8000 and Sun SPARC Enterprise M9000 servers from booting.

Furthermore, hot-swappable hardware within the Sun SPARC Enterprise servers speed service and allow for the simplified replacement or addition of components — without a need to stop the system. Devices designated as hot-swappable are easily configured for use or disabled and removed from running systems. Also, the configuration of multiple network interfaces and redundant storage devices, and dual-pathed storage networks that offer RAID capabilities provide additional opportunities for redundancy in the Sun SPARC Enterprise servers.

Resilient Interconnect Architecture

All members of the Sun SPARC Enterprise server family include a level of redundancy within the system bus. The interconnects utilized across the Sun SPARC Enterprise server line vary in implementation but share many of the same characteristics. The system interconnects are at the heart of each system, providing combined address, data, and control transport. All buses are uni-directional paths with multiplexed address and data, plus control and ECC in each direction. The motherboard in the Sun SPARC Enterprise M4000 and Sun SPARC Enterprise M5000 servers, and the CPU/Memory board unit in the Sun SPARC Enterprise M8000 and Sun SPARC Enterprise M9000 servers contain multiple system controllers (SC), enhancing reliability and performance (Figures 2-1, 2-2, and 2-3).

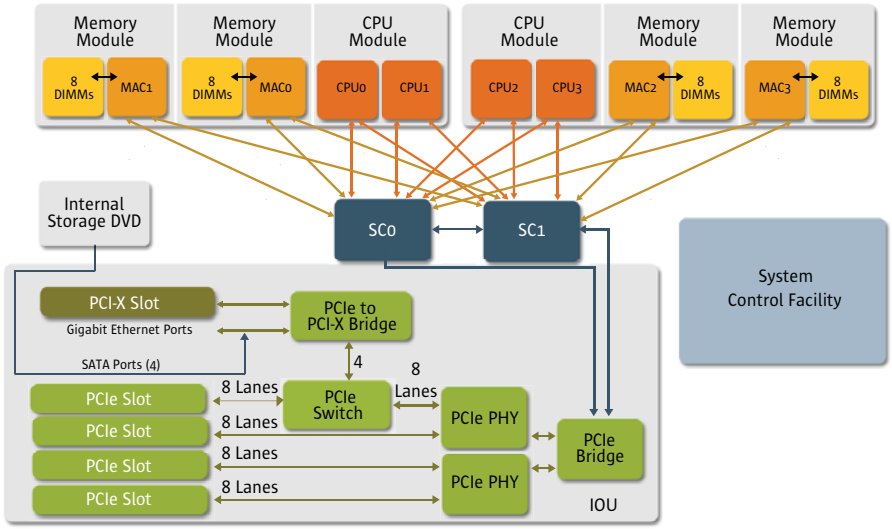

Figure 2-1. Sun SPARC Enterprise M4000 system interconnect.

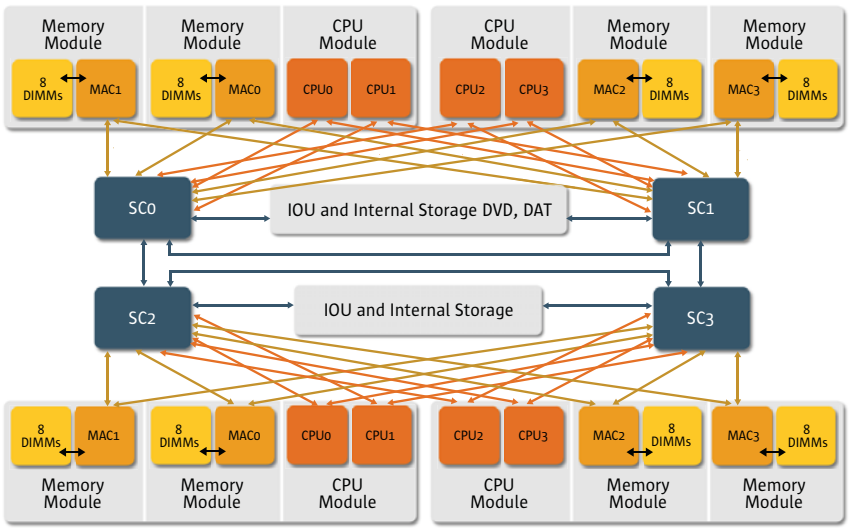

Figure 2-2. Sun SPARC Enterprise M5000 system interconnect.

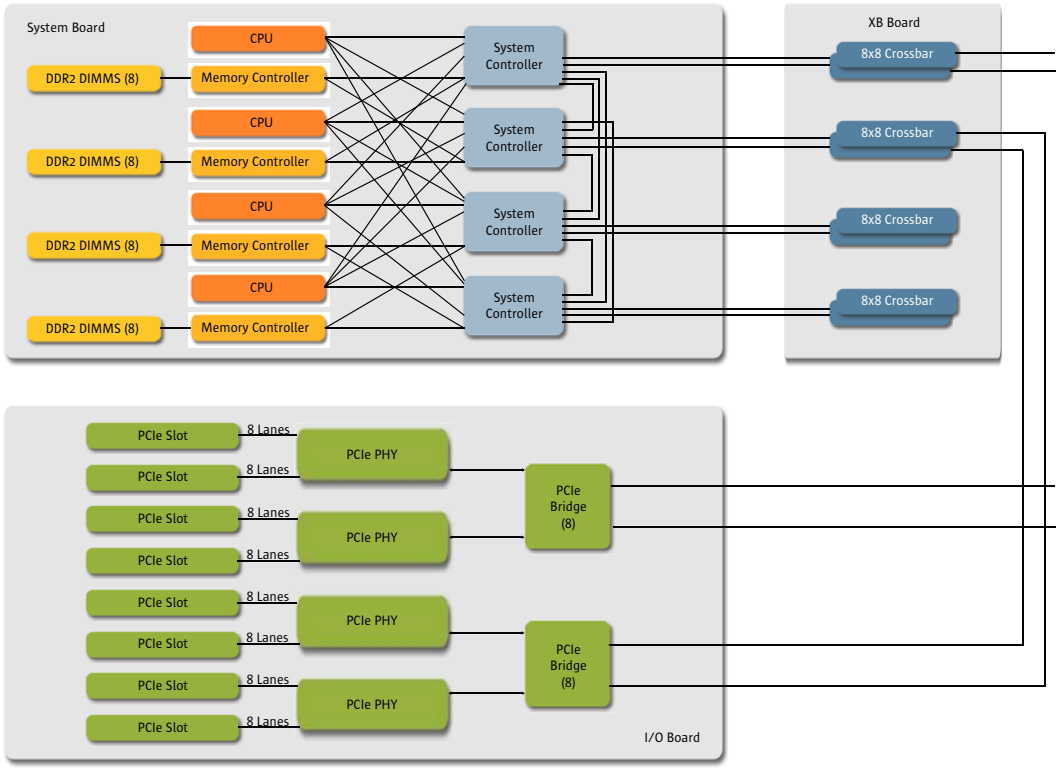

Figure 2-3. Sun SPARC Enterprise M8000 and Sun SPARC Enterprise M9000 system interconnect.

The multiple paths within the interconnect of Sun SPARC Enterprise servers increase system availability. In the rare event of a hardware failure within the interconnect, the system uses the surviving bus route on restart. This isolates the faulty crossbar and enables operations to resume. In the Sun SPARC Enterprise M8000 and Sun SPARC Enterprise M9000 servers the system controllers all link back to degradable crossbar switches, enabling these systems to continue to operate even after the loss of one crossbar switch.

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

Reliable System Clock Operation

The high-end servers in the Sun SPARC Enterprise server family provide redundancy within the clock control system. The Sun SPARC Enterprise M8000 and Sun SPARC Enterprise M9000 servers use a common clock chip in which most of the internal parts are redundant. In the Sun SPARC Enterprise M8000 server, the clock circuit is directly mounted on the backplane. In the Sun SPARC Enterprise M9000, the system clock is provided by the clock control unit. In both servers, 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 system can automatically restart in the event one route fails. The surviving route is used to continue operation.

Configurable CPU, Memory, and System I/O Redundancy

Automatic system recovery (ASR) enables Sun SPARC Enterprise servers to reboot immediately following a failure, automatically configure around a failed component, and notify the specified system administrator of the event. This approach prevents faulty or marginal hardware from causing repeated failures or keeping an entire system down.

The CPU, memory, and I/O expansion capability of the Sun SPARC Enterprise servers ensures that configurations can include redundant processors, memory DIMMs, and I/O devices. Even after the loss of one or more of these components, the Sun SPARC Enterprise servers can fully resume business processing. The following list highlights the processing and I/O expansion capabilities of the Sun SPARC Enterprise server line which enable organizations to create levels of power and system redundancy to match enterprise needs. Specific expansion capabilities of individual servers are found in Table 2-2 and Table 2-3.

- CPU and memory — The scalable Sun SPARC Enterprise server line enables enterprises to design redundant processing configurations with support for up to 64 processors and as much as 2 TB of memory. The Sun SPARC Enterprise servers also feature dynamically reconfigurable CPU and memory resources which enable the active addition or removal of compute power without stopping system operation.
- Built-in I/O — Sun SPARC Enterprise server internal PCI expansion slots enable connectivity to external devices. Combining these expansion slots with optional disk management software, system configuration can easily include multiple connections to peripheral devices, enabling redundant controllers and channels. Software tools can also maintain the multiple paths and enable 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 Sun SPARC Enterprise M8000 or Sun SPARC Enterprise M9000 system during active operation.
- External I/O Expansion Unit — Each system in the Sun SPARC Enterprise server family supports the connection of one or more optional, rackmount External I/O Expansion Unit. These units are hot-swappable and can mount up to twelve PCI-X or twelve PCI-Express cards. In addition, the External I/O Expansion Unit includes redundant hot-swappable power supplies and fan units.
- PCI cards — PCI cards which 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 Sun SPARC Enterprise 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 Sun SPARC Enterprise midrange servers, an add-on base I/O card enables access to internal devices on high-end Sun 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 within every model in the Sun SPARC Enterprise server line. Also, hot-swappable DVD-ROM and tape drive units in the Sun SPARC Enterprise M8000 and Sun SPARC Enterprise M9000 increase the flexibility of these systems.

Table 2-2. Expansion capabilities of mid-range Sun SPARC Enterprise servers.

	Sun SPARC Enterprise M4000 Server	**Sun SPARC Enterprise M5000 Server**
Processors	• SPARC64 VI • 2.15 GHz • 5 MB L2 cache • Up to four processors	• SPARC64 VI • 2.15 GHz • 5 MB L2 cache • Up to eight processors
Memory	• 128 GB maximum • 32 DIMM slots	• 256 GB maximum • 64 DIMM slots
External I/O Chassis	• Up to two units	• Up to four units
Internal I/O Slots	• Four PCI-Express • One PCI-X	• Eight PCI-Express • Two PCI-X
Hard Disk Drives 2.5-inch SAS	• Up to two drives	• Up to four drives

Table 2-3. Expansion capabilities of high-end Sun SPARC Enterprise servers.

	Sun SPARC Enterprise M8000 Server	**Sun SPARC Enterprise M9000-32 Server**	**Sun SPARC Enterprise M9000-64 Server**
Processors	• SPARC64 VI • 2.28 GHz with 5 MB L2 cache • 2.4 GHz with 6 MB L2 cache • Up to 16 processors	• SPARC64 VI • 2.28 GHz with 5 MB L2 cache • 2.4 GHz with 6 MB L2 cache • Up to 32 processors	• SPARC64 VI • 2.28 GHz with 5 MB L2 cache • 2.4 GHz with 6 MB L2 cache • Up to 64 processors
Maximum Memory	• 512 GB maximum • 128 DIMM slots	• 1 TB maximum • 256 DIMM slots	• 2 TB 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 2.5 inch SAS	• Up to 16 drives	• Up to 32 drives	• Up to 64 drives

Redundant Service Processor

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

Advancing one step further, the Sun SPARC Enterprise M8000 and Sun SPARC Enterprise 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 Sun SPARC Enterprise M8000 and Sun SPARC Enterprise M9000 servers is

also hot-swappable, enabling maintenance of the XSCF without affecting server operation.

Non-Stop Power and Cooling

Sun SPARC Enterprise servers use separate modules for power and cooling. Redundancy in each of these subsystems combined with temperature sensors placed throughout the system — measuring temperatures on processors and key ASICS as well as the ambient temperature at several locations — keep servers operating even under power or fan fault conditions.

Fan Unit

Sun SPARC Enterprise server midrange and high-end servers use fully redundant, hot-swappable fans as the primary cooling system (Table 2-4 and Table 2-5). 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. Sun SPARC Enterprise servers operate normally under these conditions, enabling ample time to service the failed unit. Once more, replacement of fans units can occur without interrupting application operation.

Power Supply

The use of redundant power supplies and power cords adds to the fault resilience of Sun SPARC Enterprise servers (Table 2-4 and Table 2-5). Power is supplied to Sun SPARC Enterprise servers by redundant hot-swap power supplies, enabling 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, Sun SPARC Enterprise M8000 and Sun SPARC Enterprise 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 Sun SPARC Enterprise M8000 and Sun SPARC Enterprise 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.

Table 2-4. Redundant power and cooling found in Sun SPARC Enterprise mid-range servers increases uptime.

	Sun SPARC Enterprise M4000 Server	**Sun SPARC Enterprise M5000 Server**
Fan Units	• Four fan units • Two 172 mm fans • Two 60 mm fans • One of each type is redundant	• Four fan units • Four 172 mm fans • Two fan groups, each containing two fan units • One redundant fan per fan group
Power Supplies	• 2350 watts • Two units • 1+1 redundant • Single-phase	• 4590 watts • Four units • 2 + 2 redundant • Single-phase
Power Cords	• Two power cables • 1+1 redundant power cables	• Four power cables • 2+2 redundant power cables

Table 2-5. Redundant power and cooling found in Sun SPARC Enterprise high-end servers increases uptime.

	Sun SPARC Enterprise M8000 Server	**Sun SPARC Enterprise M9000-32 Server**	**Sun 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 • 9 units • N+1 redundant	• 21,300 watts • 15 units • N+1 redundant	• 42,600 watts • 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)

Dynamic System Domains

Server consolidation is one of today's hottest IT trends. The economic advantages realized by consolidation in the areas of administration and system maintenance make this strategy valuable to enterprises. However, organizations require tools that increase the security and effectiveness of hosting multiple applications on a single server. Dynamic System Domains enable the division of a single system into separate, virtual servers, each running an independent copy of the Solaris OS and possessing use of designated I/O devices. When system components are exclusively dedicated to individual Dynamic System Domains, hardware or software faults in one domain remain isolated and unable to impact the operation of other domains, providing

increased application availability. Each domain within a single server platform can run a different version of the Solaris OS, making this technology extremely useful for pre-production testing of new or modified applications or system software. Domains, used with great success in the mainframe world for a number of years, prove equally useful in the world of open systems.

eXtended System Boards

Dynamic System Domains provide a very effective tool for consolidation and enable ideal separation of resources. In order to achieve this high level of isolation, previous generations of Sun servers designated entire system boards as the smallest unit assignable to a domain. However, some organizations do not require complete hardware isolation and can benefit from the ability to create a higher number of domains with compute power that more precisely matches current workloads. To meet these needs, the Sun SPARC Enterprise server line introduces eXtended System Boards (XSB).

There are two types of eXtended System Boards. An XSB that consists of an entire system board or logical system board is called a Uni-XSB. Alternatively, a system board or motherboard that is logically divided into four parts is called a Quad-XSB. Logical division lines for the creation of Quad-XSBs in each type of Sun SPARC Enterprise server are shown in Figure 2-4, Figure 2-5, and Figure 2-6.

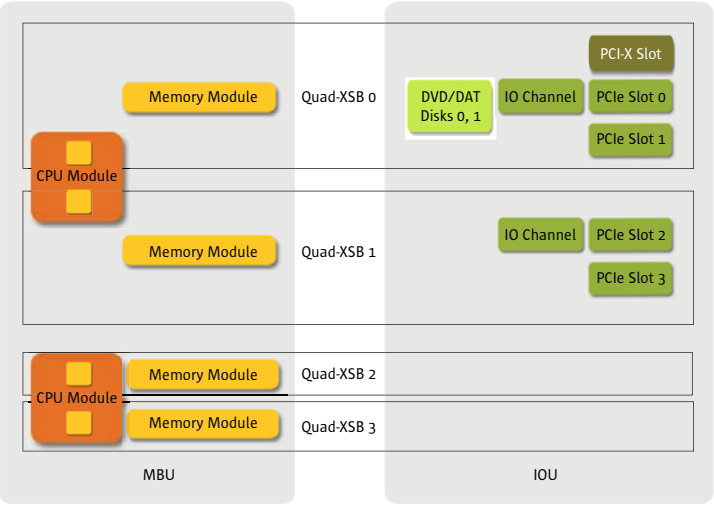

Figure 2-4. A Sun SPARC Enterprise M4000 server can be configured as a Uni-XSB or divided into Quad-XSBs.

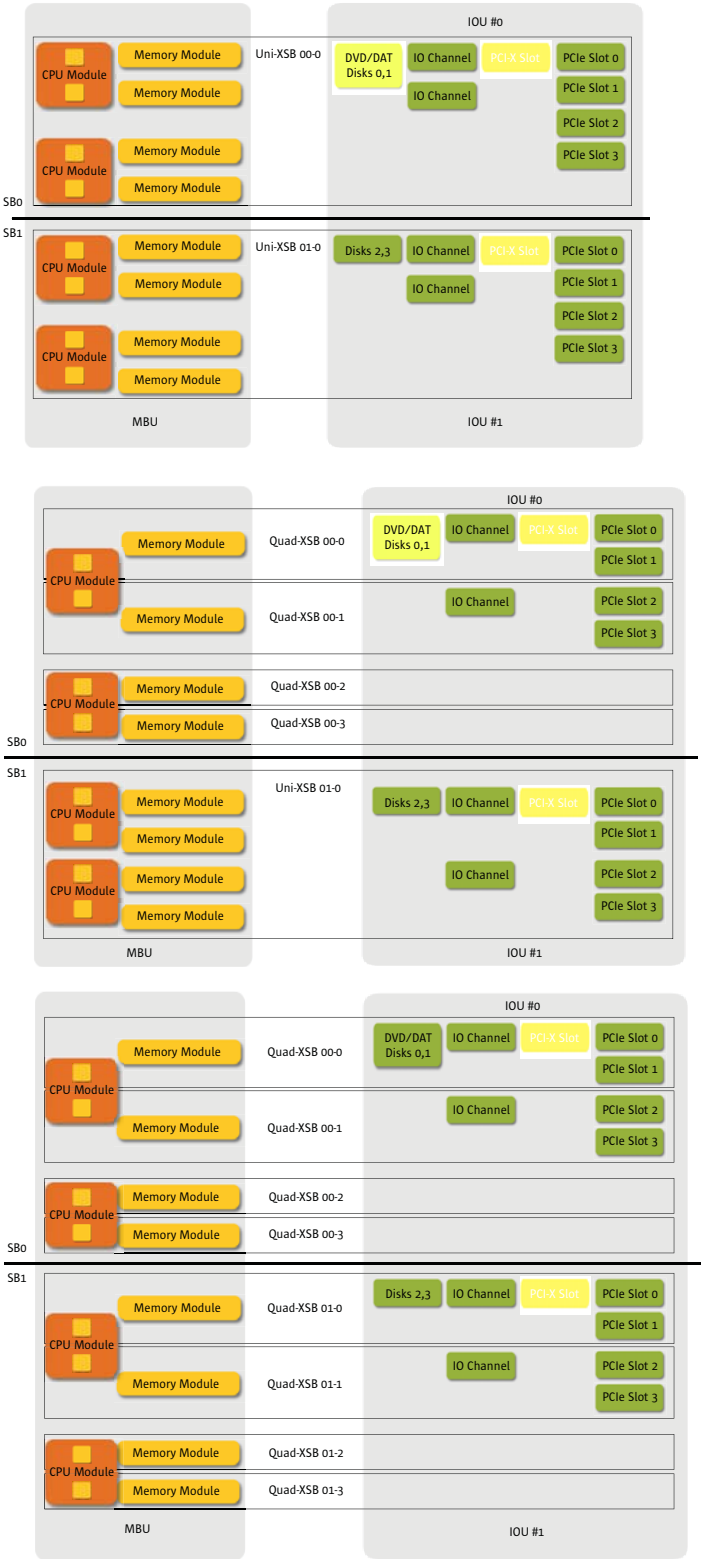

Figure 2-5. A Sun SPARC Enterprise M5000 server can be configured as two Uni-XSBs, all Quad-XSBs, or a combination of a Uni-XSB and Quad-XSB.

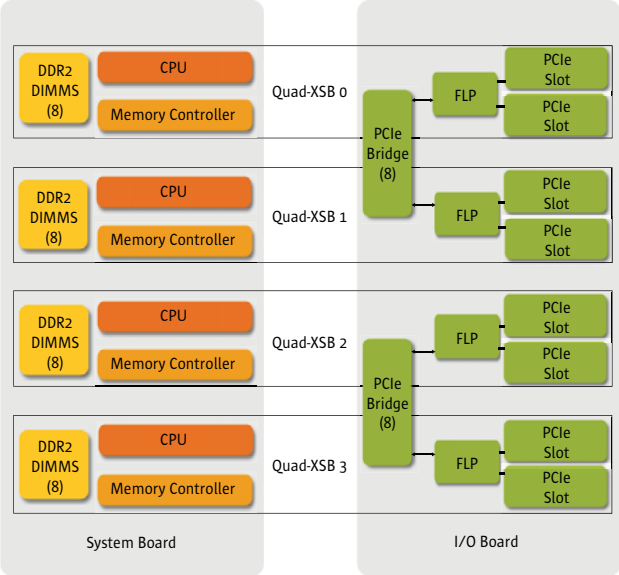

Figure 2-6. Each Sun SPARC Enterprise M8000 and Sun SPARC Enterprise M9000 system board can be configured as a Uni-XSB or divided into Quad-XSBs.

A Dynamic System Domain can consist of any combination of Uni-XSBs and Quad-XSBs, enabling sophisticated asset allocation. 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 2-6).

Table 2-6. Characteristics of Uni-XSBs and Quad-XSBs.

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

Dynamic System Domains At Work

Dynamic System Domains enable organizations to custom-tailor the compute capacity of a Sun SPARC Enterprise server to meet specific enterprise needs. For instance, a Sun SPARC Enterprise M9000 server can be configured as a single domain with up to 64 SPARC64 VI processors in order to host an exceptionally compute-intensive application. Alternatively, an organization with multiple databases which require isolation from one another might divide a single Sun SPARC Enterprise M9000 server into as many as 24 domains. More typical configurations involve three or four domains in a single Sun SPARC Enterprise server. The number of domains configurable within each type of Sun SPARC Enterprise server can be found in Table 2-7.

Table 2-7. Dynamic System Domains increase reliability and lower costs by providing fault isolation and consolidating system administration efforts.

Sun SPARC Enterprise Server Model	Maximum Number of Domains
M4000	2
M5000	4
M8000	16
M9000	24

Using Dynamic System Domains and eXtended System Boards in conjunction with Solaris Containers enables further virtualization within each copy of the Solaris OS. Together, these two methods of server virtualization represent an important cost advantage. Businesses can use fewer physical systems while still providing dedicated resources along with isolated and secure data and programs to customers. This approach leads to lower TCO by simplifying and unifying administration while at the same time providing needed privacy and support for service-level agreements.

Dynamic Reconfiguration — Online Reconfiguration, Maintenance, and Upgrades

The fluid nature of business operations challenges IT organizations to accommodate constant changes to application workloads. As a result, enterprises often need to adjust the amount of compute power assigned to each system domain. Dynamic Reconfiguration technology provides added value to Dynamic System 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. Removing and adding components such as CPUs, memory, and I/O subsystems from a running system helps reduce system downtime. In addition, Dynamic Reconfiguration makes maintenance and upgrades easier 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 enables independent domain administrators to perform DR operations simultaneously without concern for the status of DR requests or executions in other domains.

Concurrent Maintenance

Administrators are increasingly challenged to carve out planned downtime windows to perform necessary service on essential systems. Concurrent maintenance refers to the ability to perform hardware configurations without impacting application availability. Since Sun SPARC Enterprise servers are equipped with redundant components, failures no longer necessarily result in server downtime. However, the replacement of failed or downgraded components must occur at some point. Dynamic Reconfiguration enables

maintenance operations to complete while the system continues to operate. Even the assimilation of new components into running operating environments need not interrupt critical services.

Automatic Dynamic Reconfiguration

Automatic Dynamic Reconfiguration (ADR) enables 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. 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. 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.

Capacity on Demand

The Capacity on Demand (COD) option enables Sun SPARC Enterprise servers to 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. These resources remain offline at no charge until the resources are actually needed. When additional processing power is required, COD system boards are immediately configured into the system and a right-to-use license must be purchased for each extra CPU.

Sun SPARC Enterprise server configurations can include a mix of both standard and COD system boards. At least one active CPU is required for each domain in the system. Dynamic Reconfiguration operations configure COD system boards into and out of domains in the same method as standard system boards move between domains.

Whenever a COD board is configured into a domain, the administrator must make sure enough RTU licenses are available to the target domain to enable each active CPU on the COD board. If the number of RTU licenses assigned to the target domain falls short, the system displays a status message for each CPU that cannot join the domain.

Online Upgrades

Many essential compute systems need the capability to respond quickly to unpredictable loads while continuing to deliver critical services. Online upgrades enable administrators to rapidly increase the capacity and capabilities of Sun SPARC Enterprise systems — sometimes instantly — without interruption. Using this feature, configuration of additional processing power for individual domains can occur without impacting software operation. For example, online upgrades enable the configuration of additional processors, memory, I/O, and mass storage devices without requiring the system to reboot.

Many enterprise environments also require the installation of software upgrades without disruption to operations. Toward this end, features of the Solaris OS enable operating system software upgrades and patches to execute without taking the system offline. In addition, by combining the Solaris Predictive Self Healing features of the Solaris 10 OS with the Sun Update Connection, online updates complete with less effort.

Advanced Reliability Features

In addition to redundant components and features to ease maintenance, Sun SPARC Enterprise servers feature other important technologies that correct failures early and keep marginal components from causing repeated downtime. Architectural advances which inherently increase reliability are augmented by error detection and recovery capabilities within the server hardware subsystems. In the end, these features work together to raise application availability.

Reduced Component Count and Complexity

The base architecture of Sun SPARC Enterprise servers provides the foundation for achieving outstanding system reliability. The server design boasts 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.

Advanced CPU Integration and Data Path Integrity

The Sun SPARC Enterprise server line features SPARC64 VI processors. The design of this processor module increases system reliability by delivering improved fault avoidance and error correction capabilities. In fact, much of the area on the SPARC64 VI processor is dedicated to error detection and correction of data within the CPU. All RAM units within the CPU 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 CPU can autonomously offline bad cores without intervention from the service processor.

Other reliability features of the SPARC64 VI CPU include support for error marking, instruction retry, and preventive maintenance. 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. In addition, when a hardware error is detected, all instructions that are currently in execution are cancelled, and retried automatically to prevent

transient errors. Error data generated by the SPARC64 VI processor is also sent to the service processor to support preventive maintenance.

Memory Patrol, Chipkill, and Mirroring

The memory subsystem of the Sun SPARC Enterprise servers includes multiple technologies to increase stability. 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 chipkill —The memory chipkill function of these servers enables single-bit error correction, enabling processing to continue despite events such as burst read errors that are sometimes caused by memory device failures.
- Memory mirroring — When memory mirroring mode is enabled on Sun SPARC Enterprise 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 Sun SPARC Enterprise M4000 and Sun SPARC Enterprise M5000 servers, memory is mirrored within the same memory module, using the common memory address controller (MAC) ASIC (Figure 2-7 and Figure 2-8).

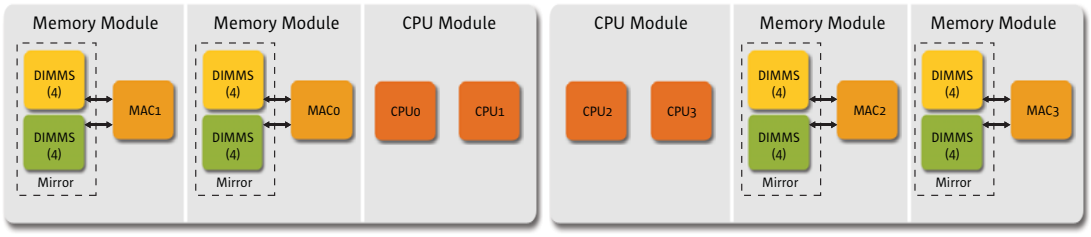

Figure 2-7. Sun SPARC Enterprise M4000 server memory mirroring architecture.

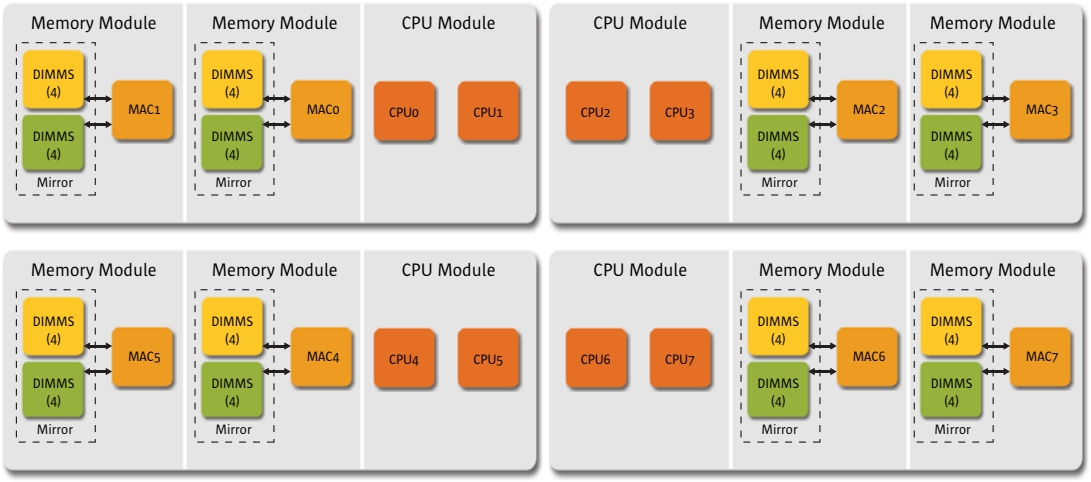

Figure 2-8. Sun SPARC Enterprise M5000 server memory mirroring architecture.

On Sun SPARC Enterprise M8000 and Sun SPARC Enterprise M9000 servers, memory is mirrored across adjacent MAC ASICs to increase reliability (Figure 2-9). However, the configuration of Quad-XSBs in Sun SPARC Enterprise high-end server system boards prevents the use of memory mirroring since each Quad-XSB only contains one MAC.

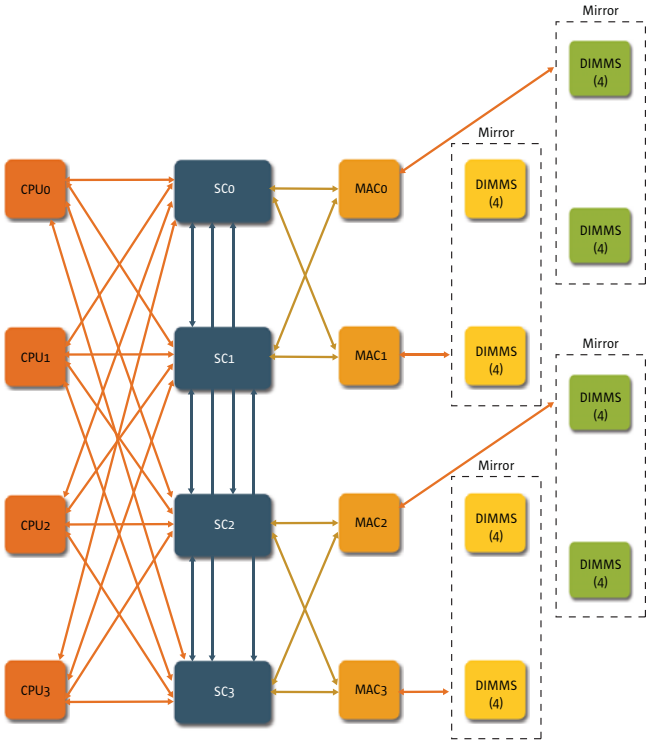

Figure 2-9. Sun SPARC Enterprise M8000 and Sun SPARC Enterprise M9000 server memory mirroring architecture.

Additional Error Detection, Diagnosis, and Recovery Features

Modern IT organizations face requirements for extreme levels of uptime on a regular basis. Datacenter platforms must provide a solid foundation fortified by built-in advanced reliability features in order for IT organizations to achieve stated goals. The Sun SPARC Enterprise family of servers are among the most dependable platforms in the industry. These 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 — Enterprise computer systems such as the Sun SPARC Enterprise servers routinely store and move significant amounts of data. The servers in this product family provide end-to-end error detection and correction on all data paths within the system, including the SPARC64 VI processor, main memory, and all 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 Sun SPARC Enterprise midrange and high-end 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 all pertinent environmental and error conditions.
- Host watchdog — The host watchdog feature of the Sun SPARC Enterprise server line 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 — The Sun SPARC Enterprise servers support Dynamic CPU resource deallocation which includes 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 component and detect signs of an impending fault. Recovery mechanisms are triggered to prevent system and application failure.
- Error identification methods — Error logging, multi-stage alerts, electronic FRU identification information, and system fault LED indicators all contribute to rapid problem resolution.

Remote Monitoring and Management

Providing hands-on, local system administration for server systems is no longer realistic for most organizations. Around the clock system operation, disaster recovery hot sites, and geographically dispersed organizations lead to requirements for remote management of systems. One of the many benefits of Sun servers is the support for *lights out* datacenters, enabling expensive support staff to work in any location with network access. The Sun SPARC Enterprise system design, coupled with the powerful XSCF, XSCF Control Package (XCP), and Sun Management Center software, as well as key functionality in the Solaris OS allow administrators to remotely execute and control nearly any task that does not involve physical access to hardware.

The XSCF provides the heart of remote monitoring and management capabilities in all midrange and high-end Sun SPARC Enterprise servers. This unit consists of a dedicated processor that is independent of the server system and runs the XSCF Control Package. 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 which might induce physical system damage.

The XSCF Control Package running on the service processor enables an administrator 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.

Sun Management Center software (Version 3.0 or later) provides platform agents that run on established domains. In addition, the XSCF Control Package contains a Sun Management Center agent for the service processor. Together, these agents provide data to the remote system for both monitoring and management. By using Sun Management Center software, administrators can take advantage of a consistent user interface and automate key tasks. As a result, organizations can reduce TCO and the likelihood of human error causing unplanned downtime.

Flexible Roles for Operators and Administrators

In addition to flexible remote monitoring and management, administrative access control remains an important issue for providing available system services. In some environments, even well-intentioned operators can accidently induce downtime. At the

same time, inconsistent access control mechanisms can impair administrators and cause circumvention of established protocols — further jeopardizing availability and security.

The XSCF Control Package enables administration of several autonomous domains independently by different system administrators and operators — all cooperating within a single Sun SPARC Enterprise platform. This management software supports multiple user accounts which 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 (Table 2-8). In addition, a user can possess multiple privileges on multiple domains.

Table 2-8. 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

With the remote monitoring and management capabilities of the XSCF and tools like Sun Management Center software, on-site tasks are reduced to swapping field replaceable components and conducting physical system maintenance.

Robust Diagnostics Support

In addition to interacting with the System Controller directly or through Sun Management Center software, the following diagnostic tools are also available for execution on individual domains of Sun SPARC Enterprise systems.

- 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 Sun SPARC Enterprise servers. The primary goal of the SunVTS software is to create an environment in which Sun systems are thoroughly tested to enable proper operation or to find elusive problems.

Architected for Reliability and Serviceability

In summary, the architecture of the Sun SPARC Enterprise servers surpasses all previous generations of servers in reliability and serviceability. These qualities are key drivers for improved availability and TCO. Features and technologies that speed diagnosis and repair help increase system availability and reduce TCO by minimizing downtime required for maintenance.

Sun continues to enhance the reliability and serviceability features of its computing products. Sun is relentlessly committed to creating technology which helps organizations deploy solutions that realize better performance, availability, and TCO. The tight integration of Sun SPARC Enterprise servers with the Solaris OS allows Sun to provide unique features such as Dynamic Reconfiguration and Capacity on Demand. These engineering innovations provide evidence of Sun’s continued leadership position in high-end UNIX servers.

Chapter 3

Software Support for Highly Available Applications

Highly available systems carry great importance in meeting uptime requirements. However, any discussion of availability must also focus on the delivery of applications and services.

- IT organizations need solutions that increase the availability of existing or commonly used applications and services without requiring undue expense and effort. Cost-conscious enterprises need off-the-shelf availability solutions for popular network services, including file services, network services, databases, Web services, and more.
- Application developers need stable and sophisticated software platforms and tools that enable the development of the application as the primary focus. Turning an application into a highly available application service should not require re-engineering applications or re-inventing complex high-availability mechanisms.
- Administrators need solutions that simplify management as IT infrastructures continue to grow both in size and complexity. Today, administrators struggle to manage the constantly changing requirements of ever-increasing numbers of systems and services.

Sun's software strategy leverages the availability features found in Sun SPARC Enterprise systems and at the same time provides real-world solutions that facilitate highly available application and service deployment. The Sun software solutions described in this chapter reduce the complexity and cost of deploying highly available solutions on Sun systems. In addition to the software solutions discussed in this section, Sun also offers cluster solutions based on Solaris Cluster software and a range of interconnect technologies as described in chapter four.

XSCF Control Package

Requests to complete maintenance tasks, configuration changes, and other system operations bombard system administrators. IT organizations need platforms that simplify and expedite completion of these types of activities. The XCP enables users to control and monitor Sun SPARC Enterprise server platforms and individual Dynamic System Domains, quickly and effectively.

The XSCF Control Package provides a command line interface (CLI) and Web browser user interface that gives administrators and operators access to all 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), enabling secure, remote execution of commands provided with XSCF Control Package.

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

- Dynamic System Domain reconfiguration — Execution of Dynamic Reconfiguration tasks to logically attach or detach installed system boards from the operating system while the domain continues to run applications without interruption.
- Domain configuration services — Domain administration which consist of creating logical system boards made up of Uni-XSB and Quad-XSB units.
- Audit administration — Logging of interactions between the XSCF and the domains.
- Hardware control capabilities — Configuration of the Sun SPARC Enterprise M8000 and Sun SPARC Enterprise M9000 server dual power grid operation. Monitor and control of power to the components of all mid-range and high-end Sun SPARC Enterprise servers.
- Hardware status monitoring, reporting, and handling — 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.
- Solaris Predictive Self Healing — Integration with the Fault Management Architecture of the Solaris OS to improve availability through accurate fault diagnosis and predictive fault analysis.
- Automatic diagnosis and domain recovery — Execution and monitoring of diagnostic programs such as the Open Boot Prom (OBP) and POST.
- Capacity on Demand — Execution of Sun Capacity on Demand operations which provide the ability to stage and then later activate additional processing resources.
- XSCF failover — Monitoring of the dual XSCF configuration on Sun SPARC Enterprise M8000 and Sun SPARC Enterprise M9000 servers for failure and performing an automatic failover if needed.

Sun N1™ Software

Intelligently managing a rapidly changing IT infrastructure requires the ability to provision or reprovision servers efficiently. Sun N1 software automates complex software installation and configuration across heterogeneous network infrastructures, helping to speed service deployment. This software also enables sharing of system resources for greater flexibility in allocating or reallocating resources as business requirements shift.

Sun N1 software products include the following packages.

- Sun N1 Service Provisioning System — Sun N1 Service Provisioning System simplifies application life cycle management by rapidly provisioning new business services and enabling quick replication of a complete software stack on servers throughout the IT infrastructure. This software offers organizations a standard method of deploying business services and tracking changes throughout the deployment process and includes an audit trail to help enterprises meet regulatory compliance. N1 Service Provisioning System provides pre-built models for deploying industry leading Java™ 2 Platform Edition (J2EE™ platform) application servers, Web servers, and databases and enables administrators to create custom XML models to deploy, install, upgrade, start and stop other applications. Every action taken by administrators across applications and managed servers is recorded, simplifying rollback to previous states in the event of a service incident.
- Sun Management Center — Sun Management Center software provides in-depth monitoring and advanced management capabilities for Sun enterprise servers to enhance system performance, reliability, security, and utilization. Further details about Sun Management Center software are included below.
- Sun N1 Grid Engine — Sun N1 Grid Engine software helps organizations optimize resource utilization by intelligently mapping compute-intensive workloads onto networked resources based on predefined policies. This software provides the ability to incorporate thousands of desktops and servers into compute grids enabling Sun SPARC Enterprise servers to augment the power of additional resources on the network.
- Sun N1 System Manager — Sun N1 System Manager software can help simplify infrastructure life cycle management for Sun servers in the context of the enterprise network of IT systems. The features of Sun N1 System Manager software enables administrators to provision, monitor, patch, and manage hundreds of servers from a single console. A hybrid user interface that integrates a GUI and a CLI in one console executes management tasks. In addition, administrators can create logical groups of systems and perform actions across a group of systems as easily as performing actions on a single system.

Sun™ Management Center Software

Too often overlooked within the larger picture of IT implementations is the impact of manageability on overall service availability and TCO. Whether onsite or in a remote location, system administrators need visibility and control of resources to quickly redistribute compute capacity, manage service levels, and prevent downtime.

Sun Management Center software improves management efficiency by offering an aggregate view of the entire network of Sun components from the heart of the datacenter to remote locations at the edge of the network. A logical view of Sun

systems enables a logical system unit to be defined as a single Sun server or as a domain within a server. Integration with Solaris Cluster systems also provides visibility and control over cluster resources.

Sun Management Center software provides IT professionals with a single interface to proactively manage and monitor remote Sun systems, storage components, the Solaris OS, and applications. Remote access services enable system administrators, onsite or offsite, to gain protected access through administrative networks. From this interface technicians can monitor system health, perform remote bring up, and restart or take down individual machines. In addition, Sun Management Center software offers the ability to drill down for fine-grained control of Sun systems including management at the board level in hardware and at the kernel level in the Solaris OS. This level of control can help administrators optimize resource utilization and improve availability.

To support legacy networks and heterogeneous environments, Sun Management Center software tightly integrates with all major management frameworks including *CA Unicenter TNG, HP Open-View, IBM Tivoli, and BMC Patrol.*

Flexible Three-tiered Architecture

Sun Management Center software is based on a three-tiered design with an agent-based framework that provides a single point of management for the enterprise (Figure 3-1). This approach simplifies systems management by delivering a higher level of scalability and availability while reducing cost and complexity. Sun Management Center software provides two separate agents for Sun SPARC Enterprise servers. A *domain agent* provides monitoring and management for individual domains. A separate *platform agent* runs on the XSCF.

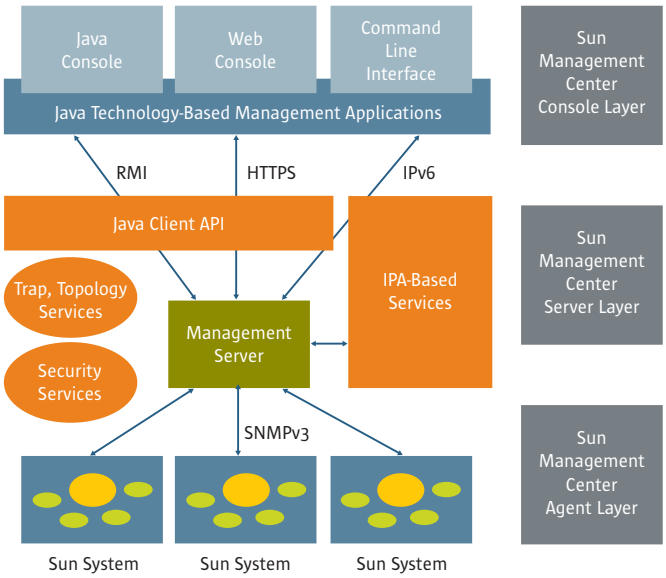

Figure 3-1. Sun Management Center three tier architecture.

A major strength of Sun Management Center software is that these autonomous agents continue to operate even when they lose contact with the central software server. Should a disaster wipe out the central server, remote agents continue to collect data and take action if an event is triggered. When the server for Sun Management Center software is restarted, perhaps at a backup site, the agents are reconnected and nothing is lost. The central server collects data on everything that occurred during the downtime, allowing historical perspectives to be maintained.

Sun SPARC Enterprise Server Functionality Support

For Sun SPARC Enterprise systems, Sun Management Center software retains the model of platform and domain management established by the XCSF and XSCF Control Package. Uniform permissions and access control lists mean that 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 Sun SPARC Enterprise servers, including domain creation, dynamic board attach, and dynamic detach. (The actual attach and detach operations take place within the affected domain, but 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.

Solaris™ 10 Operating System

Sun's 20+ years of engineering investment in the Solaris Operating System and in Sun server technology contribute to the exceptional reliability of Sun systems. The Solaris OS contains a small, compact kernel that limits the exposure to errors that can bring the operating system down. In addition, the Solaris OS maintains a clear distinction between the kernel, shared libraries, and applications in order to limits the impact of application failures. Many Sun customers can point to systems running the Solaris OS that execute continuously for months or even years without a restart, testimony to the robust nature of the Solaris OS.

The strengths of the Solaris Operating System lie in its enterprise-class reliability, scalability, and performance. The following list highlights key features of the Solaris OS.

- Performance scales to match application workload requirements, bringing mission-critical reliability, availability, and serviceability to systems of all sizes.

- Dynamic Reconfiguration, IP Multipathing (IPMP), and the ability to install most patches and other incremental software updates without taking the system offline help organizations increase uptime and ease serviceability.
- Solaris Predictive Self Healing software proactively monitors and manages system components for maximum availability.
- Solaris Containers technology enables virtualization of resources within a single image of the Solaris OS and provide secure isolation of independent software workloads.
- Scalability and performance with a 64-bit kernel enables access to more system resources and a Dynamic Tracing framework which enables developers to squeeze more performance out of applications.
- Ease-of-use features, including Web-based installation, text and voice notes, and a graphical process manager, can help boost availability by reducing the risk of operator error and by reducing service times.
- Software investment protection with complete binary compatibility enables all of today's 32-bit applications to continue to run on the Solaris OS without modification.
- Extended security features through authentication, data integrity, data privacy, and single sign-on capabilities to prevent tampering, snooping, and eavesdropping from compromising data or associated transactions. Support for IPSec prevents identity spoofing and helps to build virtual private networks.

The Solaris OS delivers the power, flexibility, availability, and compatibility to support enterprise-wide computing. Key computing elements — operating system, networking, and user environment — combine within the Solaris OS to provide a stable, high-quality foundation. By taking advantage of the following Solaris OS features, organizations can develop, deploy, and manage of a wide range of computing solutions more reliably.

Solaris Predictive Self Healing Software

When things go wrong on traditional systems, administrators 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.

Long standing features of the Solaris OS provide for system self-healing. For example, the kernel memory scrubber constantly scans physical memory, correcting any single-bit errors in order to reduce the chance of those problems turning into un-correctable double-bit errors. The Solaris 10 OS takes a big leap forward in self-healing with the introduction of Solaris Fault Manager and Solaris Service Manager.

Solaris Fault Manager, a common system that works across all platforms running the Solaris OS, reduces complexity 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 Solaris Fault Manager diagnosis engine produces a fault diagnosis once discernible patterns are observed from a stream of incoming errors. Following diagnosis, the 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.

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 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. Solaris Service Manager is integrated with Solaris Fault Manager. When a low-level fault is found to impact a higher-level component of a running service, Solaris Fault Manager can direct Solaris Service Manager to take appropriate action.

Solaris Containers Technology

Availability is often a function of application performance, and fine-grained resource management is important to building highly available systems. Partitioning resources to particular tasks can further refine a Dynamic System Domain. A Solaris Container provides a virtualized image of the Solaris Operating System, including a unique root file system, a shared set of system executables and libraries, and whatever resources the root administrator assigns to the created container. Solaris Containers boot and shut down just like any instance of the Solaris OS, and reboot in only seconds if the need arises.

Solaris Containers can 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 the Solaris OS. This technology provides a secure, virtualized environment for application execution while also giving administrators almost unlimited flexibility to assign and isolate resources to particular containers.

Redundant networking and network IP Multipathing

In addition to traditional support for multiple network interfaces connected to different network subnets, the Solaris OS 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, the ability to detect when a network adapter fails and automatically switch (failover) network access to an alternate network adapter.
- Repair detection, the ability to detect the repair of a previously failed network adapter and automatically switch back (fail back) the network access to this interface.
- Outbound load spreading, outbound network packets spread across multiple network adapters to achieve higher throughput. Load spreading occurs only when network traffic is flowing to multiple destinations using multiple connections.

Solaris Live Upgrade Software

Traditional operating system upgrades are one of the most time-consuming and error prone tasks facing system administrators. Most 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. The Solaris OS promotes greater availability by providing a mechanism to upgrade and manage multiple on-disk instances of the Solaris OS — allowing OS upgrades to take place while the system continues to operate. The time needed for an operating system upgrade is reduced to that of a simple reboot.

Sun Update Connection

Maintaining up-to-date software versions and patches on Sun SPARC Enterprise servers can help avoid downtime by avoiding known problems addressable with patches. Sun simplifies the process of keeping systems current through a new online facility called Sun Update Connection which is described further in Chapter five.

Chapter 4

Solaris Cluster — Pooling Resources for High Availability and Lower TCO

To address expanding requirements for mission-critical systems, businesses deploy an increasing number of high-availability cluster environments. Using Solaris Cluster software, organizations can mix Sun SPARC Enterprise servers with other servers from Sun's broad server line to assemble clusters of up to sixteen nodes. Cluster solutions from Sun include a combination of Solaris Cluster software, Sun servers, Sun StorageTek™ storage systems, and high-speed interconnect technology.

Solaris Cluster Systems

From a physical perspective, a Solaris Cluster system 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 (Figure 4-1). Each server generally provides some level of redundancy. In addition, storage is hosted by highly available redundant disk systems which are mirrored, enabling 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 ensure continued access in the event of a network connection or node failure.

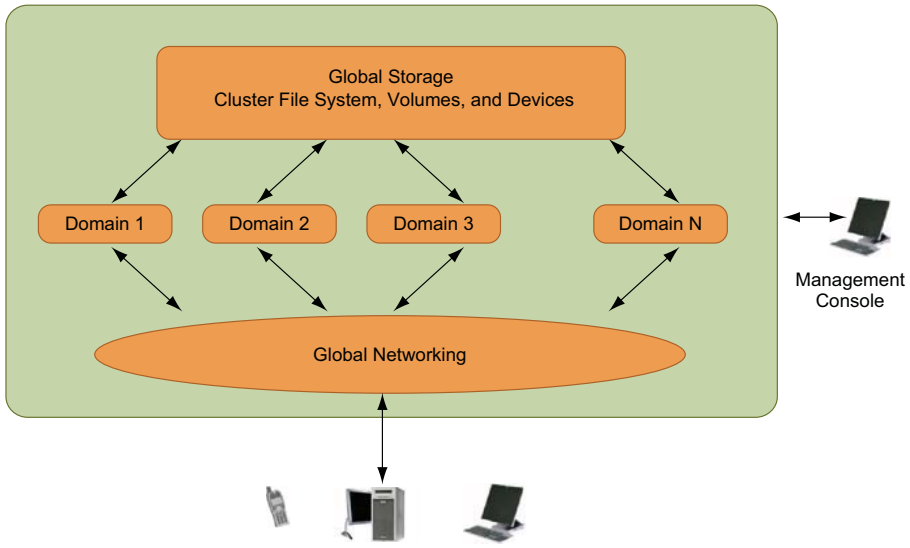

Figure 4-1. Solaris Cluster system overview.

Each server component or node in the cluster environment contains processors, memory DIMMs, an operating system, and application software. A node consists of a domain within a server or a complete server. Key to the design of Solaris Cluster systems is the fact that no single failure in the hardware, software, interconnect, or network can cause the cluster to fail. Solaris Cluster systems prevent loss of service through hardware redundancy, hardware and software failure detection, automatic recovery of services, and failover of applications. Solaris Cluster systems also provide 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.

Designed for high availability as well as manageability, Solaris Cluster systems 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.
- Increased scalability — Large SMP servers can be pooled together in a cluster, enabling 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 enables 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.

From a logical perspective, the Solaris Cluster system consists of a set of domains, where each domain independently executes a copy of the Solaris OS and communicates through the private interconnect (Figure 4-2). 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.

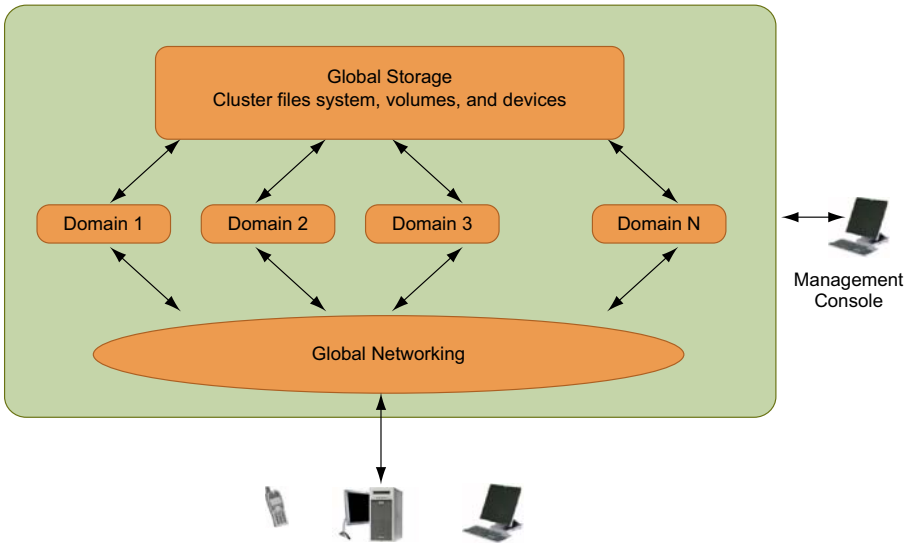

Figure 4-2. Logical view of Solaris Cluster system components.

Solaris Cluster Software

Solaris Cluster software enables enterprises to deliver highly available application services. Available for the Solaris 8, 9, and 10 Operating Systems, Solaris Cluster software extends the Solaris OS to provide enhanced availability of hosted applications. Core Solaris services, such as devices, file systems, and networks work in a virtually seamless manner across a tightly coupled cluster. While maintaining full Solaris OS compatibility for existing applications, 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 the Solaris OS to benefit from higher availability, more localized failure containment, and faster failover.

The 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 Solaris OS 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.

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

- High availability framework — Software which detects node failures quickly and creates a new server for framework resources on another node in the cluster. The framework includes a Cluster Membership Monitor, a distributed set of agents which

exchange messages over the cluster interconnect to enforce a consistent membership view, drive synchronized reconfiguration, handle cluster partitioning, and help maintain full connectivity among all cluster members.

- Virtualization support — Solaris Cluster software provides expanded support for Solaris Containers virtualization software, enabling the consolidation of servers even in high-availability environments. These virtualization capabilities allow scalable and failover applications and associated Solaris Cluster agents to run unmodified within Solaris Containers.
- Global devices, files and networking — All global devices, files, and network interfaces can be seen as local resources. This enables 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.
- High availability Solaris Zettabyte File System (ZFS) — With the virtually unlimited scalability of Solaris ZFS, Solaris Cluster software offers a file system solution with exceptional availability, data integrity, and flexibility for growth.
- Component monitoring — Extensive monitoring of application processes, disk path integrity, and network availability.
- Failover and scalable agents — Software programs that enable Sun or third-party applications to take full advantage of Solaris Cluster software features.
- Cluster-aware applications — Third-party applications with direct knowledge of Solaris Cluster systems, such as Oracle 9*i* Real Application Clusters (RAC) software.
- Multi-site disaster recovery — Solaris Cluster Geographic Edition manages the availability of application services and data across unlimited geographic distances.
- Protection groups — Solaris Cluster software enables an administrator to define groups of resources that must be managed together, enabling seamless failover of related applications and data.
- High-speed interconnect — Multiple types of interconnect technologies are supported by Solaris Cluster software to establish a private communication channel between cluster nodes.
- User interface — An object oriented command line interface and browser-based configuration wizard ease use and simplify configuration of new data services.

Free Download for Developers

Solaris Cluster software is available to developers at no cost through the Solaris Enterprise System, offering a risk-free method to evaluate, develop, and deploy this sophisticated infrastructure software. Licenses and support can be purchased when the software is deployed in a production environment.

Global Devices

Solaris Cluster software is intimately tied to the Solaris OS. In fact, much of the Solaris Cluster software functionality occurs within the Solaris OS kernel. The inherent benefits of this architecture include improved scalability, reliability, manageability, and excellent networking capabilities. There are also distinct advantages with respect to storage devices.

Solaris Cluster software can automatically detect all storage devices on system boot up, including disk drives, tape drives, and CD-ROM devices. Devices are assigned global names and are integrated into and made part of the high availability environment by default. Because Solaris Cluster software builds on standard functions of the Solaris OS, cluster nodes can increase capacity by adding devices without shutting down the server.

Global devices offer users the ability to see all global storage devices as local devices. Commands such as remote shell (`rsh`) are not required to access devices not physically attached to a local system. Note that devices can be excluded from a cluster and the root file system is not included. Global devices improve ease of use for users and enable virtually seamless storage scalability, resulting in improved availability and easier administration.

Global File Service

Solaris Cluster software abstracts the data location from the data services by employing a global file service that operates at the virtual node (vnode) layer before the traditional UNIX File System (UFS). The global file service is seen by the same mount point throughout the domains. Data no longer needs to be attached to the server that hosts the services.

By decoupling the file service from the application, software can run on a server which is not directly connected to the storage device that houses the application's data. A file system can be mounted by any domain in a cluster and every domain in a cluster sees the global file system with the same mount point, making the file system easy to administer and manage. Also, if a disk volume fails, the requests can be redirected to another volume without restarting the application. The application is typically unaware of any faults and service interruption is reduced.

Global Network Services

The ability of Solaris Cluster software to abstract a data service from a network interface means an IP service can float anywhere in a cluster high-availability environment. In Solaris Cluster software, there is a many-to-many relationship between IP services and IP addresses. Multiple IP addresses can be used for a single service, or multiple services can use a single IP address. IP interfaces can be started and stopped

and dynamically migrated from one machine to another without bringing down the service.

Redundant network interfaces enable a cluster to survive the failure of any single network interface component without causing a service interruption. In addition, Solaris Cluster software builds on the existing IP Multipathing capabilities in the Solaris OS to provide network failover features within a single Solaris domain. The global networking feature in Solaris Cluster software facilitates IP traffic distribution within a cluster, laying the foundation for scalable services that are virtually seamless to the application and resulting in improved service availability for applications that run in a cluster environment.

Interconnect Technologies

Clusters rely on interconnects to move information among nodes. Care must be taken to avoid interconnect bottlenecks which can act as obstacles to reaching the otherwise high performance potential of system designs. The cluster interconnect employed in 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 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, 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. Solaris Cluster Software supports a variety of cluster interconnect technologies, enabling high performance private cluster communications for virtually any cluster size or class. Cluster interconnects can be implemented using industry standard Fast Ethernet, Gigabit Ethernet, or Infiniband technology. Additional highlights about each interconnect offering are in the following subsections.

Ethernet

Ethernet technology whether Fast Ethernet, or Gigabit Ethernet, is a good choice for a cluster interconnect when the primary purpose of the interconnect is to support application failover. Ethernet is a cost-effective option that offers reasonable performance during failover and enables effective use of the Solaris Cluster global devices, global network, and global file system. Unfortunately, Ethernet is not a good choice when requirements include high performance data transfer between cluster nodes.

Infiniband

InfiniBand is a high-speed switch fabric architecture that offers advanced features for I/O interconnects including a mechanism to share I/O components in a cluster environment. The new architecture is designed to create a more efficient way to connect storage, communications networks, and server clusters together and deliver an I/O infrastructure that produces the efficiency, reliability and scalability that datacenters demand. With link speeds from 2.5 GB/second to 30 GB/second, InfiniBand is designed to efficiently integrate with existing Ethernet and Fibre Channel infrastructures. InfiniBand's high-bandwidth, low-latency fabric can offer significant improvement to application performance and is a good choice for cluster applications such as Oracle 9i RAC software.

Using Agents for Failover or Scalable Services

Agents are software programs that enable applications to take advantage of 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 multihost disks, agents enable applications to become highly available and scalable. Agents help prevent significant application interruption after any single failure within the cluster.

Failover Agents

Solaris Cluster systems support failover for application agents 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

Solaris Cluster systems support scalable agents, enabling application instances to run on multiple cluster nodes simultaneously. 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.

Oracle® RAC Software

Solaris Cluster Advanced Edition for Oracle® RAC software is the industry's most complete high availability offering for Oracle databases. With built-in support for Sun's Cluster File System, Sun StorageTek QFS software, and support for multiple redundant

interconnects, quorum server, and I/O fencing the advanced edition provides improved scalability, eliminates single points of failure and helps to maintain data integrity for Oracle RAC software deployments.

Available Solaris Cluster Agents

A wide range of qualified failover and scalable agents are available from Sun and third party vendors, with more expected in the near future. Specific agents are available for the following list of applications supported with Solaris Cluster software and the Solaris 10 OS running on SPARC systems.

- Apache Tomcat
- Apache Web/Proxy Server
- BEA WebLogic
- DHCP
- DNS
- IBM WebSphere MQ
- IBM WebSphere MQ Broker
- Kerberos
- MySQL
- NFS server
- Oracle 9i, 10g (High Availability and RAC)
- Oracle 9iAS Application Server
- Oracle Ebusiness Suite
- PostgresSQL
- Samba
- SAP
- SAP Database
- SAP Enqueue Server
- SAP liveCache
- Siebel
- Sybase Adaptive Server Enterprise (ASE)
- Solaris Containers
- Sun Java System Application EE/PE/SE Server
- Sun Java System Calendar Server
- Sun Java System Directory Server
- Sun Java System Message Queue
- Sun Java System Messaging Server
- Sun Java System Web Server
- Sun N1 Grid Engine
- Sun N1 Service Provisioning System

For the most up-to-date information on available Solaris Cluster agents, contact a local Sun sales representative.

Custom Agents

In addition to the agents listed above, new agents can be custom built to enable any application to take advantage of failover services or scalable services in Solaris Cluster software. A co-packaged agent development kit provides everything developers need to take full advantage of 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. Solaris Cluster offers several options for remote cluster services and includes support for industry-leading data replication software from Sun, Hitachi, and EMC.

Campus Clustering

Campus clustering essentially extends the distance between cluster components, enabling a cluster system to extend across multiple rooms or buildings in a campus. Campus clustering enables components, such as nodes and shared storage, to be located up to 20 kilometers apart. 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 distance limits 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, this is a single cluster solution supported across a limited distance that 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.

Sun Cluster Geographic Edition

To provide true business continuance for applications that must continue to run despite the scope of a disaster, unlimited distance clustering is required. For organizations that must be able to quickly survive disasters that span a wide area, Sun Cluster Geographic Edition offers enterprises a multi-site disaster recovery solution that manages the availability of application services and data across unlimited geographic distances. Sun Cluster Geographic Edition is also targeted at enterprises that carry out 24 x 7 work around the globe and require a local cluster which 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 enables an administrator to quickly start up the business services with replicated data on the secondary Solaris Cluster.

Sun Cluster Geographic Edition also is available for independent purchase or as part of Sun Java Availability Suite which can be purchased on a subscription basis for the price of $50 USD per employee per year[1]. Sun Java Availability Suite includes Solaris Cluster software, the new Sun Cluster Geographic Edition, Solaris Cluster agents for over 50 industry-leading applications, and Sun's award-winning developer tools that are contained within Sun Java Studio Enterprise, Sun Java Studio Creator, and Sun Studio software.

Protection Groups

Protection groups enable a set of clusters to map specific groups of resources for recovery in the event of a loss of service on a partner node. A protection group contains application resource groups and properties for managing data replication for those application resource groups. The configuration for a protection group is identical on partner clusters. Protection groups link an application in a resource group with the application data targeted for replication. This linkage and replication enable the application to fail over seamlessly from one cluster to another cluster.

A cluster can be defined in multiple protection groups and posses a different role in each. For example, the primary cluster of one protection group can also be the secondary cluster of another protection group.

1.Advertised pricing is for customers based in the United States only. For non-U.S. based customers, please refer to a local sales office for local area pricing.

Chapter 5

Services and Training for Managing Availability and TCO

Sun understands that building an IT infrastructure capable of delivering high service levels requires a comprehensive approach — one that addresses people, processes, and technologies. By leveraging Sun's extensive expertise, best practices, and innovative managed services offerings, enterprises can reduce potential risks of downtime and enhance the stability and performance of IT services.

Each Sun Services customer engagement includes the following primary goals.

- Work collaboratively with internal teams and partners to share learning.
- Ensure business needs are addressed by technology investments.
- Accelerate project implementations while reducing risk.
- Establish highly effective and repeatable datacenter processes based on best practices and automated solutions.
- Build in-house expertise through both knowledge transfer and training.

Sun's portfolio of service and support offerings includes technology consulting, educational services and training programs, service plans for Sun hardware and software solutions, and proactive monitoring and management including both remote management and managed services offerings.

Many of these services can be delivered or accessed in the near future through Sun Connection, an always-available, secure services connection that links customers, partners, developers, and Sun in a dynamic and collaborative network-based community. Sun's worldwide service and support organization is also complemented by a broad network of service and support partners. Additional details on Sun Services offerings described below can be found on Sun's Web site at *sun.com/service*.

Sun's Suite of Consulting Services

In a world in which technology choices can make or break an organization's ability to compete, Sun's suite of consulting services rapidly closes the gap between new technology adoption and business results. Sun's consulting services accelerate the realization of important business goals through the rapid and seamless deployment of Sun's innovative products and technologies. By leveraging Sun'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 Sun consulting services portfolio 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 5-1 describes some of the services offered by Sun Services and its partners.

Table 5-1. Sun's suite of consulting services.

Sun Service Offering	Description
Sun Operations Capabilities Assessment (OCA) Service	The Sun Operations Capabilities Assessment Service benchmarks the ability of a customer to efficiently manage their IT operations. Well defined criteria and metrics are used to assess the capabilities of people, process and tools. It can include a complete review of the IT environment or a focused assessment concentrating on specific critical areas.
Sun Enterprise Installation Services (EIS)	Sun Enterprise Installation Services are designed to speed time to deployment and provide the foundation for enhanced stability and performance. EIS is based on Sun's proven methodology including site audit, installation planning, system installation specification, statement of installation, installation and customization of Sun hardware and software, and installation verification testing and system turnover.
Solaris 10 OS Upgrade Service	The Solaris 10 upgrade service facilitates the automated, rapid upgrade of existing SPARC and/or X86-based platforms running a previous version of the Solaris OS (7, 8 or 9) to the newly released Solaris 10 OS. By creating a dual boot environment that supports both the old version of the OS and Solaris 10, this service preserves all the features and functionality of the current deployment (same hardware, same software stack) by replacing just the operating system and allows the administrator to effortlessly switch between versions of the OS.
Sun i-Runbook Service	The Sun i-Runbook Service creates a knowledge power-base that enables the entire IT team to quickly find the vital information needed to manage the datacenter. A centralized source of data provides the latest Sun preferred practices and is always available via a simple browser interface to enable IT staff to better understand and manage Sun and third-party hardware and software.
Solaris Cluster Implementation Service (CIS)	The Solaris Cluster Implementation Service is designed to provide a flexible framework for configuring Solaris Cluster Software. Comprised of six modules, CIS provides the design, configuration, and testing to ensure that a new Solaris Cluster system is functioning correctly with a customer's applications to help create a robust, operational environment.

Sun Service Offering	Description
Sun Application Readiness Services (ARS)	The Application Readiness Service for Sun SPARC Enterprise servers provides the services required to configure a new Sun SPARC Enterprise server up to the point of application installation. (This service does not load the application.) This service focuses on the availability, security, and storage requirements of the application environment and enables faster time-to-production by helping to ensure that the system is ready to support mission-critical applications.
Sun Reference Implementation Service for Data Centers (RISDC)	Sun Reference Implementation Service for Data Centers provides services to safely and reliably build, test, and implement datacenter configurations. These Reference Implementations are proven configurations comprised of a flexible combination of Sun servers, Sun storage, Sun Java Enterprise System, the Solaris OS and other software and may also include Local Area Network (LAN) and Storage Area Network (SAN) infrastructure.
Sun Performance Analysis and Capacity Planning Services	The Performance Analysis Service includes collection of a variety of data points in customer systems (e.g., CPU utilization, memory utilization, and swap rates) as well as analysis of the data. Sun makes recommendations on ways to improve server performance using the customer's existing hardware configurations and recommends additional hardware if necessary. The Capacity Planning Service includes measurement of the growth potential of the customer's current server environment and identifies equipment that may be needed to meet future business requirements. (The Performance Analysis Service can be delivered independently of the Capacity Planning Service, but the Capacity Planning Service can only be delivered along with the Performance Analysis Service.)

Sun Service Offering	Description
Sun Enterprise Migration Suite	Sun Enterprise Migration Suite of services can help organizations safely migrate to a new IT infrastructure while staying focused on critical business issues. Sun consultants carefully analyze business objectives, business processes, and operational requirements, then architect and implement a solution based on the organization's unique environment. The goal is to help achieve a successful migration that protects legacy application investments, mitigates future migration efforts, improves IT efficiency and service delivery, and allows IT organizations to more easily adapt to evolving corporate standards, strategies, and regulatory pressures.
Sun Application Migration Service	The Application Migration Service helps customers maintain version and platform consistency and provides a tested methodology for a successful migration to the Solaris 10 OS. The Application Migration Service helps customers migrate their application from many UNIX or LINUX variants to the Sun platform, making the application ready to compile and run on Solaris 10 or earlier versions.
Infrastructure Delivery Option (IDO)	Infrastructure Delivery Option (IDO) is a partner-only program offered by Sun to enable qualified Sun partners to sell and deliver the offerings from the IDO service portfolio. After a clearly defined training path, ending in a structured certification, partner engineers gain access to Sun's service methodologies. Sun developed these methodologies based on Sun's experience from years of delivering installation and infrastructure services. The methodologies are documented and crafted in a selection of support documents, checklists and media (CD/DVD). The result is a structured path for delivering the highest service quality, regardless of the delivery resource.

Sun Management Services

Sun Management Services provide dynamic IT infrastructure management that optimizes datacenter operations while enabling enterprises to 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 while cutting management complexity through collaborative single point of contact relationship between the customer and the service providers. 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.

Sun's multisourcing approach enables organizations to select service providers that specialize in different areas and combine individual optimized services into an

integrated solution that is managed by a single lead vendor (Sun, a partner, or an in-house IT team). Sun works with partner IT service providers and in-house services teams, creating an environment in which each organization can contribute individual core competencies. This approach also enables enterprises to source individual IT programs or source the entire integrated IT environment.

Sun believes enterprises should remain 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. Organizations can also retain control over assets and human resources since Sun provides IT infrastructure services only.

The Sun Management Services portfolio includes services ranging from packaged offerings of relatively fixed scope to fully 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 5-2 describes some of the packaged offerings from Sun Management Services that can be linked together and custom-tailored to configure a unique solution for each situation.

Table 5-2. Sun Management Services offerings.

Sun Management Service Offering	Description
Sun Service Definition Workshop	This interactive workshop is designed to review the management requirements of an organization's complete multivendor environment to develop priorities and build a strategy for optimizing IT service management and aligning IT with business goals.
Sun Interim Operations Management	Sun Interim Operations Management is a short-term service designed to help customers reduce operational risk and accelerate return on investment when implementing new IT solutions. Sun IT operations experts temporarily run the new datacenter environment to bridge any transitory skill gaps, transfer knowledge to internal staff, and bring best practices to IT processes.
Sun Managed Operations Services — Managed Storage Services	Sun experts use the Sun Managed Operations Services platform to perform day-to-day tasks for remote management, monitoring, and reporting on the operation of Sun StorageTek products and SANs.
Sun Managed Operations Services — Managed Security Services	Sun security experts provide 24 x 7 surveillance and analysis including event monitoring and management, penetration testing, and security audits to help prevent security breaches and ensure early detection.

Sun Management Service Offering	Description
Sun Managed Desktop Services	Sun experts manage heterogeneous desktop environments and provide full support for end users including a customer Web portal, service desk, inventory management, image management, and onsite support and administration.
Sun Preventive Services	Sun Preventive Services provide an advanced methodology to improve availability by proactively identifying risk factors in people, processes, and technologies. Customers are given financial incentives for staying healthy and sustaining improvements over time.
Sun Site Support	The Site Support Program is managed by an assigned Sun Service Account Manager and carried out by site-resident Sun support engineers. By placing Sun engineers and a managed parts inventory on-site, Sun can respond quickly to problems that arise while delivering a range of services that can proactively avoid trouble and optimize system performance.
Sun Managed Site Infrastructure Services	Sun Managed Site Infrastructure Services include pre-architected solutions for specific components of datacenter operations to help reduce cost and complexity. The pre-defined solutions which can be customized to meet unique IT operational requirements include reactive hardware and Solaris support, service management, on-site systems administration, service desk, and heterogeneous IT operations management.
Sun Operations Management Solution	Sun Operations Management Solution is an integrated set of workshop, assessment, transformational and operational offerings designed to help organizations optimize their IT infrastructure to support operational best practices and create a flexible IT environment that can respond quickly to changing market conditions. Using a phased approach, Sun helps customers identify areas that need improvement, remediate gaps, and provide ongoing service management support.

With the flexibility offered by Sun Management Services, Sun can help organizations improve the efficiency of IT operations and overcome IT skills shortages without giving up control of IT programs and direction.

Sun Educational Services

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

Many enterprises that purchase servers such as the Sun SPARC Enterprise servers engage in large projects that may require additional skills across the organization. Sun recommends organizations begin such a project with an assessment of the people assigned to new roles and an evaluation of their current and desired skills sets. From this starting point, a tailored learning plan can be developed and applied from the outset of the project.

The Sun learning portfolio includes services ranging from up-front analysis and skills assessment to training and certification. Before a project begins, Sun 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. And finally, Sun 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, Sun's training solutions include a variety of formats for learning.

- Traditional classroom training — Traditional classroom training at Sun and Sun partner facilities around the world provides access to Sun expertise and courseware in a classroom setting that provides ready access to the technologies explored in the course. Sun's courseware includes skills training and certifications in leading-edge technologies developed by Sun and its partners.
- On-site training — Sun brings expert instructors on-site to train an organization’s IT staff, eliminating the need for employee travel and enable in-tact teams to share in training programs.
- Customized training — Sun provides customized training content and delivery to meet unique project requirements.
- Online learning (Sun Learning Connection) — Online learning options from Sun provide a cost effective, versatile, and efficient means to improve and accelerate human performance. Individuals or teams of IT staff can access self-paced Web-based content or utilize Live Virtual Class training which combines the traditional strengths of classroom training with the ease of use, and anywhere accessibility of the Internet.
- CD-ROM courses — Courses delivered via CD-ROM are also an excellent way to get the core competency training that individuals need without requiring travel or extended time away from the office.

The Value of Certification

The experience, training, and study required to achieve a Sun 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 Sun 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. Sun offers a range of ongoing management and support services to help IT organizations meet service level agreements and help minimize TCO.

SunSpectrum℠ Service Plans

SunSpectrum℠ Service Plans help ensure systems run smoothly and meet service availability demands. Through a combination of technical support, hardware service coverage, and Solaris updates, SunSpectrum service plans can resolve technical issues quickly and effectively. From mission-critical services to basic self-maintenance support, Sun offers four different support plans that address a range of requirements.

- SunSpectrum Platinum℠ Service Plan — For critical systems, Sun recommends SunSpectrum Platinum which includes complete 24 x 7 coverage, our highest priority response, and additional specialized services to help proactively maintain high system availability.
- SunSpectrum Gold℠ Service Plan — Going beyond basic support, SunSpectrum Gold provides extended coverage hours and interoperability assistance to help organizations successfully manage key business systems.
- SunSpectrum Silver℠ Service Plan — Sun's most popular Service Plan, SunSpectrum Silver provides basic service and support with business-hours coverage. This plan includes all of the essentials for one low price.
- SunSpectrum Bronze℠ Service Plan — Designed to support companies that use in-house personnel to perform hardware service, SunSpectrum Bronze support provides technical support, self-help tools, and parts replacement coverage.

All levels of SunSpectrum support include access to the Solaris OS enhancement releases, patches, and maintenance releases at no additional charge. These updates are made available through Sun Update Connection, an online update service that proactively manages software updates from Sun to help drive greater availability of IT services and reduce the risk of security breaches.

Sun Update Connection

Sun Update Connection simplifies the process of staying up-to-date and secure with the latest software updates from Sun. This service builds on previous patch management tools from Sun and provides an easy to use GUI as well as a CLI. All aspects of patch management are integrated into a seamless architecture to deliver the following features.

- Notifications to let administrators know when new updates become available for systems
- Automated procedures that greatly simplify the task of keeping systems current
- Fast intelligent software performs dependency checks so that updates are automatically deployed along with all dependent updates that are prerequisites
- Optional local caching of updates to help minimize network traffic and enhance security for the update process
- A Web hosted service that provides a centralized view of connected systems and enables consistency in applying updates across multiple systems

Sun Update Connection is designed to maintain an accurate and up-to-date IT infrastructure while providing maximum flexibility for accomplishing system updates. Sun Update Connection is part of a broader suite of service offerings called Sun Connection which is described below.

Sun Remote Services

The Sun Remote Services (Sun RS) program is a set of tools, technologies, and expertise designed to support Sun's remote services strategy. The strategy aims to improve mission-critical support and increase the overall availability of the environments of SunSpectrum Gold or SunSpectrum Platinum support customers. Sun RS methodologies allow Sun experts to detect and manage Sun system events faster and more proactively, resolving potential problems before impact to Sun environments and business operations occurs. With Sun RS, this rapid response can be accomplished through remote, real-time monitoring of systems and improved, ongoing communication between Sun and the enterprise’s IT staff. As an added benefit, Sun RS provides enterprises with increased access to monitoring data, alerts, and reports, and to Sun technical support.

Sun Remote Services Net Connect (Sun RS Net Connect) provides a collection of services designed to help organizations proactively manage Sun systems, including self-monitoring, event notification and training, and comprehensive reporting.

RAS System Analysis is a new automated server assessment service that leverages Sun's best practices and extensive RAS knowledge to provide detailed reports of specific, actionable issues affecting reliability, availability and serviceability.

Sun RAS Profile

Sun RAS (Reliability, Availability, Serviceability) Profile is designed to regularly identify those areas where, based on industry standards and best practices, businesses can optimize operation of in-house Sun platforms.

Available as an addition to SunSpectrum Platinum and SunSpectrum Gold contracts, Sun RAS Profile is an ongoing, comprehensive server assessment, leveraging Sun's unique knowledge bank of product configuration data, and is conducted by senior Sun engineers. Armed with this data, Sun experts drill deeply into the configuration issues, as well as the potential operational and environmental issues that can affect the reliability, availability, and serviceability of Sun systems. Sun RAS profile includes the following deliverables.

- Server configuration analysis
- Comparison to a database of Sun and industry best practices
- A management level report highlighting areas of risk
- A technical report of risk analysis and recommendations for fixes

Sun Connection

Sun Connection, Sun's new vision for always available service and support, delivers intelligent IT services over the network in an on-demand access model that includes multiple modes of connectivity — always connected, occasionally connected, and disconnected. The Sun Connection program includes a trusted service connection and shared access to a community of users, enabling organizations to explore new opportunities for creative solutions to business needs.

Continually streaming services provide a rich connection between customers, partners, developers, and Sun. The user interface is an always-available network connection instead of a human being on an 800 telephone number. Technology is used intelligently to radically reduce complexity and provide a highly scalable service delivery mechanism that can dramatically reduce the cost of service delivery and support. This new economic model enables smaller IT environments and even individual systems to receive service and support benefits, once only available to large datacenter operations.

Services are delivered according to intelligent rules maintained within Sun Connection, enabling the right services to be offered at just the right time. In addition to timely delivery, the services are highly predictive, proactive, and prescriptive in nature, enabling IT staff to stay on top of service level performance indicators and identify and remediate risk factors before IT service problems are experienced by end users. Sun Connection can be thought of as a central nervous system that connects customers and partners with intellectual property, sophisticated analytical and diagnostic capabilities, and streaming services designed to identify risk factors, predict and prevent service downtime, and proactively manage high quality IT services.

Sun Connection consolidates services into one easy-to-access portal that organizations can personalize to meet enterprise needs and achieve the following benefits.

- Cost-saving efficiency in notification and delivery of software updates
- Instant access to sophisticated services via remote service delivery
- Faster problem diagnosis and resolution through remote servicing and an always-available support connection
- A proactive approach that can help IT staff focus on strategic priorities instead of reacting to problems
- Prescriptive training delivered at the precise moment of need to help build IT skills
- Better visibility into IT risk factors
- Improved self-service support using a direct link to self-serve knowledge and information repositories
- Services delivered within the context of the current situation

Chapter 6

Summary

Downtime for mission-critical services impacts businesses immediately with the potential for missed opportunities, customer dissatisfaction, and increased operational costs. IT departments continue to focus on highly available services to meet service level agreements, but need ways to reduce costs at the same time.

Sun understands that delivering highly available services is not enough. In order to meet cost objectives, businesses must also efficiently manage IT services. TCO is the result of the interaction of people, process, and product over the life cycle of a project and ultimately drives many investment decisions. Sun is dedicated to providing effective solutions that can help address both availability and TCO across the entire spectrum of an IT implementation.

Sun brings the following three major strengths to help enterprises build cost-effective and highly available solutions.

- Highly reliable server architectures — Sun SPARC Enterprise servers, the mid-range and high-end of Sun's extensive server line, 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 Sun's server technologies.
- Automation to simplify system management — Sun provides management solutions that can help businesses reduce the cost of building and managing highly available application services. Sun's system management offerings enable detailed control of Sun hardware components and the ability to efficiently manage an entire network of Sun systems.
- Proven expertise and support — Sun’s technology consulting, innovative education offerings, support tools, and proactive monitoring and management solutions can help organizations architect, implement, and manage IT investments to help improve availability and reduce TCO.

With Sun SPARC Enterprise servers, Solaris Cluster software, XSCF Control Package firmware, Sun Management Center software, and an extensive line of enterprise-class services and training, Sun provides an end-to-end solution for helping organizations realize the most demanding availability requirements while minimizing TCO.

For More Information

For more information on how Sun can help build availability into an IT infrastructure, visit the Web sites in Table 6-1 or contact a local Sun sales representative.

Table 6-1. Web Links for Additional Information.

Web Site URL	Description
sun.com/servers/	Sun servers
sun.com/servers/sparcenterprise	Sun SPARC Enterprise Servers
sun.com/javaenterprisesystem/	Sun Java Enterprise System
sun.com/software/n1gridsystem/	Sun N1 software
sun.com/solaris/	Solaris Enterprise System
sun.com/cluster/	Solaris Cluster software
sun.com/sunmanagementcenter/	Sun Management Center software
sun.com/service/	Sun Services solutions

Sun Microsystems, Inc. 4150 Network Circle, Santa Clara, CA 95054 USA **Phone** 1-650-960-1300 or 1-800-555-9SUN (9786) **Web** sun.com

Printed in USA 04/07