Availability Features in the Sun™ Enterprise™ X500 Server Family

Technical White Paper

Printed in the United States of America.
901 San Antonio Road, Palo Alto, California 94303 U.S.A.

Contents

Introduction 1

In today's highly competitive marketplace, Information Technology managers struggle with the challenge of balancing expenses against the demands for improved service levels. Computing resources are increasingly distributed within the enterprise, adding to the complexity and cost of system management. Users, on the other hand, want systems and applications to be continuously accessible and available, and they need high levels of compute performance as well.

Recognized for outstanding performance and scalability, Sun™ Enterprise™ servers are used today in mission-critical computing applications — in retail, banking and insurance, manufacturing, healthcare, telecommunications, and a variety of other markets. Sun Enterprise servers also form the cornerstone of many corporate infrastructures, providing database, application, file, and web services throughout many local and wide area networks. In all of these environments, systems availability is critical.

Without necessary data and applications, workers cannot be productive — orders cannot be processed, customers cannot be serviced, and products cannot be designed, manufactured, and delivered. Unplanned outages are extremely costly, both in lost productivity and, all too frequently, direct revenue.

Sun Enterprise X500 Servers — Unprecedented Availability

Sun Microsystems, Inc., a world leader in enterprise system computing, is responding to customers' growing needs for higher system availability. In addition to excellent performance and scalability, Sun is now offering unprecedented availability with the Sun Enterprise X500 family of servers — the Sun Enterprise 3500, 4500, 5500, and 6500 systems.

New availability features such as Dynamic Reconfiguration and Alternate Pathing were previously exclusive to the Sun Enterprise 10000, a mainframe-class system designed for the data center. Sun is now incorporating these availability features in the full Sun Enterprise server product line, addressing availability requirements throughout the enterprise, from individual departments through the corporate data center. The availability offered by the Sun Enterprise X500 product family is unmatched by the current generation of competitive server products.

Sun Products and Services for High Availability

Availability means something different to almost every customer. For this reason, Sun provides a full range of product and service offerings to meet customers' varying availability requirements. Sun offers these single node and clustered server solutions:

- *Sun Enterprise X500 Family*

  The Sun Enterprise X500 servers incorporate a number of sophisticated reliability features, as well as refinements to enhance system serviceability. The end result is a line of scalable, high-performance, highly available computing platforms. Chapter 3 discusses the availability features of the Sun Enterprise X500 servers and describes how these features are implemented.

- *Sun Enterprise 10000*

  The Sun Enterprise 10000 offers the most scalability of all the Sun Enterprise servers, with up to 64 processors, and the highest availability of the Sun Enterprise single-node systems. In addition to leading-edge availability features such as Dynamic Reconfiguration and Alternate Pathing, the Sun Enterprise 10000 includes Dynamic System Domains, which allow a single machine to be divided into logical partitions, with each partition totally electrically isolated from the others. Much like logical partitioning available on mainframes, with Dynamic System Domains, environments can be safely isolated from one another, and resources quickly reallocated to meet changing requirements for performance or throughput.

- *Sun Enterprise Clusters*

  To achieve even higher levels of availability, Sun also offers Sun Enterprise Clusters. Sun Enterprise Clusters offer increased availability by supporting application failover between nodes. Up to four Sun Enterprise server nodes are combined in a cluster configuration with high-speed interconnects, high-speed networking, and mirrored storage. Specialized software provides comprehensive database (Oracle, Informix, Sybase, or IBM DB2), SAP R/3, NFS™, or web service failover between nodes. Both standard and parallel database applications are supported. Sun Enterprise Cluster software includes a robust fault management API, allowing other applications to be enhanced for failover support, either by customers themselves or experienced consultants from Sun Enterprise Services.

The following products and services can be added to the server configurations above to provide a complete solution:

- *Sun™ StorEdge™ Products*

  The Sun StorEdge™ product line is a complete family of scalable, highly available storage solutions. Sun StorEdge products feature fully redundant, hot swappable power supplies, fans, and disk drives. Data protection is available through mirroring with either controller or host-based RAID. When installed with Sun StorEdge storage solutions, the Sun Enterprise servers provide robust, highly available computing solutions.

- *Sun Mission-Critical Services*

  Sun provides leading-edge, mission-critical services. Customers can choose from multiple levels of global services to provide maximum uptime. At the highest level of support, Sun offers 7x24 coverage, proactive account management, remote system monitoring and assistance, and a customized guarantee of system availability.

This paper focuses on the innovative availability features of the Sun Enterprise X500 servers. For more information on other Sun products, see the related white papers listed under *References*.

Availability Overview 2

Introduction

Enterprises continuously seek increased levels of availability in the systems they deploy. To meet these demands for higher availability, platforms must have special features that enhance system uptime. In addition, systems must be highly reliable and easily serviced, so that when problems do occur, they can be quickly diagnosed and repaired.

This chapter broadly discusses the topic of availability, while Chapter 3 explores specific availability features and technologies used in the Sun™ Enterprise™ X500 servers.

Availability

Availability is the time a particular resource, such as a system, application, or data, is accessible and usable. Obtaining higher levels of availability begins with core system design and extends to the overall software application architecture.

Availability is typically measured as *a percentage of total uptime over the course of a year.* System uptime requirements are almost always stated this way. For example, a 99.99 percent availability requirement translates to 52.8 minutes of downtime per year, where a 99.9 percent availability requirement means about 8 hours of downtime per year.

What's often unclear is whether or not the downtime factor includes planned maintenance or just unscheduled outages, and if the downtime represents a single outage or a total of outages through the year (e.g. , 8 x 1 hour outages or 1 x 8 hour outage). Of course, such questions can only be answered on a case-by-case basis, with the customer assessing their business requirements and the cost of downtime for particular applications.

Reliability and Serviceability

To a great extent, *reliability and serviceability drive availability*. The more failure reduction built into a system, the more available it is likely to be. Features and technologies that reduce failure recovery time, or that speed up diagnosis and repair, increase a system's availability. For this reason, the acronym *RAS*, which stands for *reliability, availability, and serviceability*, is often used in describing availability features of a system.

Reliability

Reliability is the starting point for building increasingly available systems, since a measure of a system reliability is how long it has been up, and/or how long it typically stays up between failures. The nature of the failure is not important — any failure affects the system's overall availability.

Mean Time Between Failure, or MTBF, is often considered an important metric with respect to measuring system reliability. However, there is currently no industry adopted standard for measuring MTBF, which makes the MTBF number for a given system or component of questionable use for comparison against other vendors.

There are two primary means of achieving greater reliability:

- Building high MTBF components into the system, and adding them in redundant (N+1) configurations, and
- Including technology that maintains data integrity, reducing the probability that bad data will flow undetected through the system

Data integrity is key in measuring a system's reliability. For example, a system with an uptime history of 99.999 percent (5 minutes per year downtime) would be useless if it produced wrong answers or provided no mechanism for integrity checking as the data moved through the various system components. In a highly available system, there must be safeguards to ensure that I/O operations are not lost, for example, if a failure does occur.

Serviceability

Serviceability defines the time it takes to isolate and repair a fault, or, more succinctly, the time it takes to restore a system to service following a failure. Mean Time To Repair, or MTTR, is considered an important metric when discussing the serviceability of a system or some component of the system. MTTR, however, is a unit of time and does not factor in the cost of service.

Consider the cost of a 7x24 (seven days a week, 24 hours a day) service contract versus 5x8 coverage. In some environments, a system that recovers automatically from most failures can obviate the need for 24 hour service coverage, and can reduce support staff requirements. The net result is a highly available system with a lower cost of ownership.

Over the last few years, Sun has delivered increasingly modular systems with better serviceability. Continuous innovations — ranging from reducing the number of jumpers and module slot dependencies to tighter integration and greater environmental tolerances — have shortened the time required for the replacement of a failed module. These enhancements, coupled with improved diagnostic capabilities, have significantly reduced the service cycle on systems. This has had the net effect of increasing overall system availability.

With the addition of Dynamic Reconfiguration and Alternate Pathing, Sun is providing a quantum leap in system serviceability. With these features in a properly configured system, a faulty component can often be replaced and brought on line, *without disrupting system operation, while users continue to access their data and applications.*

Defining Availability Requirements

The terminology used to describe availability can be confusing, and tends to be inconsistent within the industry. Sun Enterprise Clusters are said to provide *highly available* application services — when a node in the cluster fails, applications are failed over to another node. The term *fault tolerant* is typically used to describe features where automatic failover occurs without any intervention or disruption of service (e.g., the Alternate Pathing feature enables disk storage controllers to be fault tolerant). However, in discussing overall system availability, these terms can have different meanings to different people. For this reason, *system uptime* is the best method of specifying overall system availability goals.

System uptime is usually given as the percentage of uptime that the system provides over the course of a year. The following two tables list terminology used when discussing availability and corresponding uptime metrics. Table 2-1 is the work of Jim Gray, an acknowledged expert in computers and transaction oriented systems. Table 2-2 comes from a very large telecommunications company.

System Type	Unavailability Minutes/Year	Availability (Percent)	Availability Class
Unmanaged	50,000	90	1
Managed	5,000	99	2
Well-Managed	500	99.9	3
Fault Tolerant	50	99.99	4
Highly Available	5	99.999	5
Very H.A.	0.5	99.9999	6
Ultra H.A.	0.05	99.99999	7

Table 2-1 Availability terms and metrics (Source: Jim Gray)

Class	Downtime/Year	Percentage
Conventional	3.5 days	99.000
Highly Available	8.5 hours	99.9
Fault Resilient	1 hour	99.99
Fault Tolerant	5 minutes	99.999

Table 2-2 Availability terms and metrics (Source: major telecommunications company)

Looking at both tables, very different uptime percentages are used for the same terms. For example, Jim Gray defines *fault tolerant* as providing 99.99 percent uptime, while the telecommunications company suggests it refers to 99.999 percent uptime. Hence, when discussing availability, it is important that system uptime goals are clearly stated *as a percentage* to ensure that a specific availability target can be met.

Achieving Availability Requirements

Meeting clearly defined availability requirements is principally a matter of *configuring for availability*. For example, base systems can be configured for availability by adding extra power supplies, I/O controllers, mirroring disks, etc. Redundant system components minimize single points of failure in the hardware configuration. To take advantage of software enhancements such as Alternate Pathing and Dynamic Reconfiguration, configuring redundancy into the system becomes extremely critical. (Chapter 4 discusses configuring for availability in more detail.)

Even higher levels of availability can be achieved through further redundancy — for example, adding an entire system as a hot spare backup — at an additional cost. As an example, Sun Enterprise Clusters combine up to four Sun Enterprise servers, mass storage, high-speed interconnects, and specialized software. These systems offer fully automated fault detection and recovery from a single point of failure for highly available web, file, database, and application services. Sun Enterprise Cluster systems provide the highest database availability for transaction processing or decision support applications, and support both standard and parallel applications on the same node. (For more information on Sun Enterprise Clusters, refer to the white papers listed under *References*.)

All of these approaches increase availability, but equally expand system cost and/or complexity. As with any other critical business issue, availability decisions must be made based on the cost of downtime versus the investment required to adequately safeguard against it.

Availability Features in the Sun™ Enterprise™ X500 Product Line 3

Introduction

The Sun™ Enterprise™ X500 servers offer many innovative availability features. These include:

- *Fast system recovery*, through Automatic System Recovery and Power On Self Test (POST)
- *Enhanced system reliability*, through a simple system design, sophisticated environmental monitoring, and a new software enhancement called CPU Power Control
- *Improved data integrity* with ECC circuitry and parity checks
- *"No outage" servicing*, enabled by enhancements such as Alternate Pathing, Dynamic Reconfiguration, and hot-pluggable and hot-swappable components
- *Robust diagnostics support* through a comprehensive diagnostic tool (SunVTS™)
- *Intelligent systems monitoring and management* with built-in hardware monitors, Sun Enterprise SyMON™, and remote console capabilities

These features are shared across the entire Sun Enterprise X500 product family — the Sun Enterprise 3500, 4500, 5500, and 6500. Many features existed in the previous Sun Enterprise X000 server designs (the Sun Enterprise 3000, 4000, 5000, and 6000). Some, such as the sophisticated environmental sensing and shared power design, are tightly integrated into the hardware of these systems.

With this high level of technology and integration, availability was clearly a prominent design goal and not an afterthought for the Sun Enterprise server family.

All Sun Enterprise servers run Solaris™, a robust and mature operating environment that supports thousands of customers in mission-critical operations today. In recent Solaris operating system releases, Sun is offering several software enhancements to promote even higher levels of availability: CPU Power Control, Alternate Pathing, and Dynamic Reconfiguration. In addition, Sun is simplifying the task of monitoring system availability in the enterprise with a new comprehensive system management framework — Enterprise SyMON.

This chapter describes all of the features that contribute to system availability in the Sun Enterprise X500 systems. It explains how they are implemented, and where appropriate, how they will evolve over time to provide even higher levels of availability.

Fast System Recovery

All Sun Enterprise X500 systems incorporate the ability to detect failed hardware components and *boot around the failed components*. Therefore, a failure of a CPU/Memory Board, processor, memory bank, or faulty I/O Board does not keep the entire system down.

Automatic System Recovery (ASR)

Specific types of hardware failures, such as a processor failure, may bring down a system. Automatic System Recovery (ASR) enables Sun Enterprise X500 systems to reboot immediately following a failure, automatically configuring around a failed component. This approach prevents faulty hardware from keeping the entire system down or causing the system to fail repeatedly.

ASR tests various hardware components when the system is first powered on, or when an external reset is generated. The primary component of ASR is Power-On Self Test (POST). The POST code is resident in Flash PROMs that are built into each core system board, and which share PROM space with Open Boot PROM (OBP) code. Both CPU/Memory Boards and I/O Boards include a Flash PROM.

Using Flash PROM in the Sun Enterprise X500 systems adds to the serviceability of these machines. With Flash PROM, upgrades are simply a matter of loading in new PROM code, which can be done via a `tftp` process, or through a flash update utility on CD-ROM. This facilitates a fast and easy way to incorporate Open Boot PROM and POST code changes for new functionality or bug fixes.

Power On Self Test (POST)

Power On Self Test, or POST, is a suite of hardware integrity tests implemented in system firmware. POST is executed at different points in time, such as when the system first powers up. The POST code is rule-based, and possesses different levels (tests) of execution. POST testing improves system reliability by preventing the system from booting with faulty hardware as part of the active configuration, which could compromise uptime or data integrity.

POST was developed as a stand-alone set of programs entered through a section of the Open Boot PROM code called the *reset dispatcher,* which, as the name implies, is the section of system firmware executed whenever a reset condition occurs (e.g., power-on, fatal error, reboot, etc.). The reset dispatcher enters POST under the following conditions:

- Following a fatal system error
- The system's key switch is in the DIAG position
- An Open Boot PROM (OBP) environment variable is set
- Other resets (such as a system reboot)

To support Dynamic Reconfiguration, the POST code is also executed to test a newly inserted system board prior to its activation (see *Dynamic Reconfiguration* later in this section). This testing reduces, but does not entirely eliminate, the possibility of a failed component on a newly inserted system board bringing the system down.

POST performs testing at various levels, depending on environment variables and how the testing is initiated. To restore the system to operation quickly, a minimal amount of testing occurs after a soft reset and system reboot. More extensive testing, however, is done after a fatal system error to identify the problem and potentially to configure around a failed component.

Typically, POST initializes and tests all core system board ASICs, memory SIMMs, UltraSPARC™ processor modules, and the Clock Board. POST code does not, however, test SBus™ or PCI cards and associated I/O devices (this testing is done through card-resident FCode).

In addition to testing system components, POST firmware also probes, initializes, and configures system memory. The level of SIMM testing is determined by the current *diag-level* value, which can vary the number of patterns run through the memory SIMMs. POST also includes a *diag-level* that causes POST menus to be displayed for manual testing.

POST code executes out of the Flash PROM and, for some tests, the external cache on each CPU/Memory Board. The Boot Bus, which provides access to the Flash PROM that contains the POST code, also contains 32 KB of SRAM and the board temperature register (see *Sophisticated Environmental Monitoring*, below). This allows POST to run and identify problems even on systems that are experiencing severe memory or bus problems.

POST code is written such that every operation is checked for success or failure. POST also includes built-in timers to prevent broken hardware from causing it to hang. This allows POST to identify (at a reasonably fine level of granularity) which component failed at a given point in time during the testing.

POST testing is hierarchical in nature, starting with basic initialization and functionality testing and moving forward based on the success or failure of previously tested components. As it proceeds, POST disables failed components that could interfere with future testing.

As a part of its base of rules, POST code checks the installed components to confirm that they are compatible with the clock speed of the system. The enhanced Gigaplane™ bus in the Sun Enterprise X500 family can be clocked at speeds ranging from 84 to 100 MHz. (The actual clock speed is determined by the CPU modules, the CPU/Memory Boards, and the Gigaplane bus itself.) POST will automatically sense if a slower component is installed and configure the system to run at the slower rate.

JTAG

JTAG (Joint Test Action Group) is an IEEE standard for testing electronic circuits. JTAG is used by POST in the Sun Enterprise X500 systems for initialization and diagnostics, and during system operation to verify the integrity of hot swap boards.

JTAG testing provides the ability to test for essential bit *set* and *clear* functionality in the various board and ASIC registers in a system. JTAG testing shifts bit patterns through the ASICs and reads back the patterns to confirm they are correct. Connectivity between ASICs and boards is also tested by moving the data pattern out from one ASIC into another and testing for correctness in the target location. Multiple ASICs are tested serially through a JTAG concept called *scan rings*. Each scan ring represents some number of ASICs chained together sharing a common test signal. Each core system board can contain a maximum of 16 scan rings.

JTAG control and command is done via JTAG registers in the Boot Controller ASIC on each core system board. This allows for JTAG scan tests to be run in parallel, so larger systems with more boards do not require longer POST run times.

Fatal System Errors

POST also functions as the end point for handling fatal system errors, whether they occur during POST testing or while system software (e.g., the Solaris operating environment) is up and running. Following a fatal system error, POST will reconfigure a system around the failed component so the system can be restored to service. A fatal system error is an error that causes the system to be in an illegal hardware state, such as:

- UPA address parity error
- Gigaplane address or control parity error
- DTag parity error
- External cache tag parity error
- UPA Master Port time-out
- Internal error

When a fatal system error occurs, the hardware state is set to indicate the error (i.e., status bits in hardware registers) to the POST firmware. The POST code will initiate JTAG ring scans on the system boards, saving the scan data in local SRAM memory. A message is then printed indicating the nature of the error, and whether the error occurred while Solaris was running.

What POST does at this point depends on the status of key environment variables. Correctly configured, POST can run the self tests and in many cases identify the failed component and configure around it. POST has the ability to configure around failed processor modules, banks of memory, CPU/Memory Boards, I/O UPA Ports (SBus, PCI, and Creator), I/O Boards, TOD/NVRAM chips, and DC power supplies. POST can even be configured to use an alternate boot disk to work around the failure of a disk controller or an I/O Board (obviously the contents of the boot disk must be mirrored in such a case).

Enhanced System Reliability

Simple Design

The first step in building reliability into a system is simplicity of design. Large, complex systems with numerous components and interconnects are inherently less reliable since they have more potential areas of failure. Additionally, with large systems, diagnosing a problem is a more complex effort, increasing service times and reducing overall system availability.

Sun Enterprise X500 systems implement a *minimum* number of active components and interconnects, using highly reliable VLSI designs. Fewer components means a higher MTBF. MTBF is further increased by a simple, elegant centerplane bus, which reduces the amount of bus circuitry required, and by the passive nature of the centerplane/backplane design. The simplicity of the Sun Enterprise X000 and X500 architectures clearly demonstrates Sun's commitment to significantly reduce the component count over previous generation systems. At the same time, Sun Enterprise X000 and X500 systems provide significant gains in performance.

Of particular importance, Sun Enterprise X500 servers are built with a *jumperless configuration* and *no slot dependencies*. Adding, removing, and replacing core system components is an easy task that can typically be done in minutes using only a screwdriver. The net result of this “keep it simple” design is greater uptime.

Sophisticated Environmental Monitoring

Sun Enterprise X500 systems use sophisticated system environmental monitoring to protect against failures due to temperature extremes, lack of air flow through the system, or random fluctuations in AC power. (Insufficient air flow can occur if the appropriate installation clearances are not observed. These clearances are critical — in particular the side clearances — to prevent the outgoing warm air from one system entering the air flow intake of another.)

Temperature Monitoring

Sun Enterprise X500 systems are tested for operation in temperatures ranging from 5 to 35 degrees Centigrade (41 to 95 degrees Fahrenheit). Each CPU/Memory module has a thermistor installed below each of the processor boards. The analog output from each thermistor is fed to an analog-to-digital converter and the resulting value is placed in a system register for software access.

The same implementation is used on the I/O Boards and the Clock Boards, so accurate temperature readings are maintained for all core system boards. The sampled temperature is used to drive the speed of the cooling fans enclosed in the 300 watt Power/Cooling Modules (PCMs). Thus it is possible to have different PCMs running fans at different speeds to provide sufficient cooling. (Chapter 4 describes the redundant power and cooling functionality in more detail.)

Note that a memory-only CPU/Memory Board will not provide any temperature data, as no thermistors are installed for monitoring memory SIMM temperatures. This does not adversely effect system reliability in any way, since memory SIMMs do not generate a significant amount of heat.

A polling mechanism implemented in the Solaris operating system reads the temperature registers every 2 seconds. The component temperature data is then available for CPU Power Control (described below) or system monitors such as Enterprise SyMON (described later in this section).

Power Supply Status

All of the PCMs present status bits to indicate their status. For example, a *present* bit and a `dc_ok` bit indicate if the PCM is installed and working. A change in the state of the `dc_ok` bit — which may occur during random or periodic power fluctuations — generates a hardware interrupt. When this occurs, the Solaris operating system is notified and a message can be printed.

A `fan_fail` bit is also used in the PCMs, but requires a board to be installed in one of the two adjacent slots for the bit to be visible to the software. (In a redundant power configuration, one PCM is installed without any boards in the two adjacent slots.)

Any change in the status of the `fan_fail` bit generates a system interrupt. However, the failure of a fan will not necessarily cause a Sun Enterprise X500 system to shut down. A new software enhancement — CPU Power Control — may prevent a shutdown from occurring as a result of an over-temperature condition.

CPU Power Control

The Solaris 2.6 operating environment adds a new safeguard — CPU Power Control — for all Sun Enterprise X500 servers. CPU Power Control helps to ensure that the temperature of any CPU/Memory Board never goes above a safe operating level.

As Chapter 4 describes, a Power/Cooling Module (PCM) provides cooling for every two adjacent board slots. An over-temperature condition most often results from a fan failure in a PCM, especially when the remaining fans are not able to provide sufficient cooling or the system air flow is blocked. When an over-temperature condition occurs, CPU Power Control takes the processors off-line on the associated CPU/Memory Board, and powers them down.

The Solaris operating environment checks the system temperature registers every 2 seconds. If the core temperature of a processor on a CPU/Memory Board exceeds 72 degrees Centigrade (162 degrees Fahrenheit), the Solaris operating system emits an over-temperature warning to the console, and then automatically powers down the processors on the board. Since processors are the chief source of heat on a CPU/Memory Board, powering them down lowers the temperature to a safe operating range. While the processors are powered off, the memory banks on the board are still fully accessible, and users may continue to access their applications using other processors in the system. While the system continues to operate (although in a degraded mode), the associated Power/Cooling Module (PCM) can be replaced, assuming the system is configured with redundant power. When the problem is fixed, the system administrator can bring the processors on-line again with the `psradm` command.

If the problem is not resolved and the temperature hits 83 degrees Centigrade (181 degrees Fahrenheit), the system repeats the warning. If a processor's core temperature stays at this level for 20 seconds or longer, the system will power itself down entirely. This prevents any physical damage to the system.

CPU Power Control cannot take off-line and power down the last processor in the system. For example, if there is only one board in the system with two processors, CPU Power Control will power down the first processor but not the second. Similarly, CPU Power Control cannot take off-line any processors that have threads bound specifically to that processor. In these cases, if the temperature continues to climb, the system will initiate a shutdown if necessary.

Improved Data Integrity with ECC

ECC Circuitry and Parity Checks

Sun Enterprise X500 systems are designed throughout to help ensure that data is not lost or corrupted. To this end, the Gigaplane backplane/centerplane and the UPA interconnect are fully protected by ECC circuitry, as is main memory. Even within the tightly controlled environments within chips such as the UltraSPARC processor, cache SRAMs, and the various bus interface ASICs, data paths are fully protected by parity checks. Address lines on the Gigaplane and the SBus peripheral bus are also protected by parity.

An uncorrectable data error may result in a fatal system error. Should an uncorrectable ECC error occur, the processor will panic and bring the system down. POST will examine the error state on reboot and intensify the POST level of testing on the component that generated the ECC error.

"No Outage" Servicing

With new features like Alternate Pathing and Dynamic Reconfiguration, the Sun Enterprise X500 servers provide improved serviceability. In conjunction with the hot-pluggable design of the Sun Enterprise X500 hardware, these features allow many operations to be safely performed while a system is running, minimizing downtime.

Alternate Pathing to Network and Disk

Alternate Pathing (AP) promotes system availability by maintaining access to critical network and storage resources even after a failure. Just as mirroring provides a redundant copy of the data in the event of a disk failure, Alternate Pathing provides redundancy in the event of a *controller* failure. With Alternate Pathing, if a disk controller fails, disk operations are automatically redirected to a predefined alternate path (Figure 3-1). *Users continue to access disk storage without interruption.*

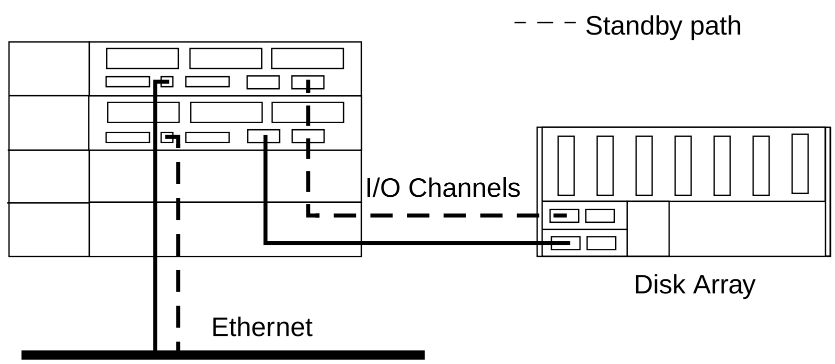

Figure 3-1 Alternate Pathing provides redundancy in the event of a storage or network controller failure

In the case of network controllers, the path must be manually switched. If the primary network interface fails, the administrator switches network traffic to the standby alternate path.

Alternate Pathing is implemented via a database that maps logical device paths (*meta-disks* or *meta-networks*) to physical device paths. The database is replicated on multiple raw disk partitions so that it is still accessible if a copy gets corrupted. The Alternate Pathing daemon (`ap_daemon`) acts as an intermediary between commands issued by the system manager to alter the database and the AP librarian (`ap`), which actually maintains the AP database.

All I/O requests to a meta-device are sent to the appropriate *meta-driver*, which passes them to the corresponding physical device driver. The meta-drivers select which physical path to use and re-route traffic when a path becomes unavailable. The meta-drivers make these decisions based on information in the AP database.

Alternate Pathing is the foundation for Dynamic Reconfiguration of I/O Boards (see *Dynamic Reconfiguration* below). To take advantage of Dynamic Reconfiguration, if a system I/O Board contains a network or storage controller, the system manager must first disable the controller prior to logically detaching the I/O Board. To maintain uninterrupted access to the network or data attached to that controller, the system manager can first switch traffic from the active controller to the standby alternate path. Until the I/O Board is reinstalled and reactivated, all traffic is redirected to the alternate path.

The Alternate Pathing 2.1 software included with the Solaris 2.6 operating environment works with Sun network adapters, dual-ported SPARCstorage™ Array devices, and Sun StorEdge™ A5000 devices. AP is not required with the StorEdge A3000 (the RSM Array 2000) because that array includes built-in controller failover functionality. AP 2.1 is compatible with the following storage management software:

- Sun™ StorEdge™ Enterprise Volume Manager™ 2.4 and later (Version 2.5 is suggested, but Dynamic MultiPathing, which provides similar functionality, should not be used)
- Veritas VxVM versions 2.3, 2.4, and 2.5 (Version 2.5 is suggested, but Dynamic MultiPathing, which provides similar functionality, should not be used)
- Solstice™ DiskSuite™ 3.0 (Version 4.0 is not supported)

For more information, refer to the *Sun Enterprise Servers Alternate Pathing User's Guide*.

Dynamic Reconfiguration

Dynamic Reconfiguration (DR) software enables changes to a system's hardware resources *while the system is up and running, without a system reboot*. In discussing Dynamic Reconfiguration, it is helpful to understand the following terminology:

- *Dynamic Attach* — The process of logically attaching new components to the Solaris operating environment, and making those components available for use (e.g., I/O Boards, processors, memory, etc.)
- *Dynamic Detach* — The process of logically detaching components that were available and on-line to the operating system, making them ready for removal
- *Hot Plug* — The ability to physically add or remove a component on a powered-up system without disrupting system operation
- *Hot Swap* — The ability to add or remove a component from a powered-up system without any notification to the operating system — *the operating system automatically recognizes the component*. This term is typically used to describe disk drives and power supplies.

Today the system boards in the Sun Enterprise X000 and X500 servers are all hot pluggable — a new board can be inserted or an inactive board may be removed at any time. However, a board cannot be activated or deactivated while the system is running — a system reboot is required for the operating system to recognize the changed hardware state. With the addition of Dynamic Reconfiguration functionality, the hot pluggable system boards in the Sun Enterprise X000 and X500 servers can be inserted AND configured for use (dynamically attached), or a board can be disabled (dynamically detached) and then removed from a running system. Dynamic Reconfiguration eliminates the need for system reboots after hardware configuration changes, and allows users to continue running their applications without interruption.

Sun is implementing DR functionality in phases. Dynamic Reconfiguration is first available for I/O Boards with the Solaris 2.6 operating environment. In a subsequent release, Sun will enhance DR to include support for CPU/Memory Boards.

Adding a System Board

The process of Dynamic Reconfiguration differs depending on whether a board is being added or removed from the system. When a system board is added, there are three primary steps: insertion, connection, and configuration.

First the system administrator confirms the status of an empty and enabled slot using `cfgadm`, the DR administration interface. Then the board is physically inserted by an authorized Sun service provider. Next, the system administrator uses `cfgadm` to connect the board (i.e., it is powered up and electrically connected). At this point the temperature sensors on the board are operational. Finally, the administrator configures the board for use by the operating system, again using `cfgadm.` During the configuration process, the board is checked via the same POST testing that occurs on power-up, the OBP device structure is built, and device drivers are dynamically loaded.

During a dynamic attach operation, the POST testing tries to identify any faulty component prior to activating the board in the system. However, it is possible that POST may not detect a failed component in absolutely every instance, which could result in a system failure when the board is activated. This risk, although slight, should be kept in mind when planning DR attach operations.

When adding or removing system boards, it is critical that *trained Sun service personnel should perform the insertion or removal*. In the process of adding a board, inserting the board too slowly or bending a pin could result in a fatal system error. Therefore only qualified Sun service personnel should perform these actions.

Removing a System Board

The process is generally the reverse for removing a system board; the board must be unconfigured and disconnected before it may be physically removed.

When a system board is to be removed, the system administrator must first quiesce all activity to the board. (It is important to note that Dynamic Reconfiguration does not automatically quiesce activity on the board — the administrator must do this prior to dynamically detaching the board.) In the case of a system I/O Board, all traffic to devices on the board must stop. If the system is configured for Alternate Pathing, the administrator can switch network and storage traffic to predefined alternate paths — the alternate paths would consist of different Ethernet and disk controllers on other I/O Boards.

Using `cfgadm`, the administrator then unconfigures the board, which takes the device drivers off-line and deletes the OBP device structure. At this point the board may be disconnected using `cfgadm`. Disconnecting the board removes power from the board and puts it in a disabled state. An LED indicates that the board is in a powered down state and that it may be unplugged at any time by a qualified Sun service technician.

Dynamic Reconfiguration — I/O Boards

Initially Dynamic Reconfiguration will be supported for system I/O Boards with three SBus slots. To determine whether a site is qualified to use DR, an important first step is to analyze the system's device drivers. To support DR, all device drivers for the I/O Board and its resident SBus cards must implement these four operations:

- `DDI_SUSPEND`
- `DDI_RESUME`
- `DDI_DETACH`
- `DDI_ATTACH`

In addition, all device drivers *within the system* must have at least SUSPEND/RESUME functionality to use DR.

Sun has modified device drivers for many Sun network devices, I/O controllers, and other Sun SBus cards, as well as for the three-slot SBus system I/O Board itself. Sun has enhanced the following device drivers to support all four DR operations — therefore any of these devices may be resident on an SBus I/O Board that is dynamically reconfigured:

- SunSwift™ 100BaseT Fast/Wide SCSI SBus Adapter
- SBus Single-Ended Fast/Wide Intelligent SCSI-2 Host Adapter (SWIS/S)
- SBus Differential Fast/Wide Intelligent SCSI-2 Host Adapter (DWIS/S)
- SBus UltraSCSI Differential Host Adapter (UDWIS/S)
- SBus Differential SCSI-2 Host Adapter (DSBE/S)
- 25 MB/s Fiber Channel Controller
- 100 MB/s Fiber Channel Controller
- SBUS SunFastEthernet™ 10/100 SBus Adapter
- SBUS Sun™ Quad FastEthernet™ (QFE)
- TurboGX™/TurboGXplus™ (TGX, TGX+)

The following additional devices support only SUSPEND/RESUME. They may be resident in the system, but DR cannot be performed on the particular I/O board on which they reside:

- SBus Fast SCSI Host Adapter (FSBE/S)
- PCI SunSwift 100BaseT Fast/Wide SCSI
- PCI SunFastEthernet
- PCI Sun Quad FastEthernet
- Creator (FFB)
- Creator3D (FFB2)

If a third party SBus device is used, it is important to check with the vendor to determine whether the driver has been modified to support Dynamic Reconfiguration operations.

DR cannot be performed on the central I/O board (usually the I/O board in slot 1) that provides the electrical path to the Clock Board. In addition, DR operations should be performed on only one board at a time. See the *Dynamic Reconfiguration User's Guide* for more information.

Dynamic Reconfiguration — CPU/Memory Boards

In a subsequent Solaris operating system release, Sun will enhance DR to provide support for CPU/Memory Boards. Implementing DR for CPU/Memory Boards is more complex than for I/O Boards, primarily because of the difficulty of relocating memory. Because DR requires that all activity to the board first be quiesced, all memory accesses to the board must be stopped. To do this while the system is running requires that all memory pages be temporarily locked and then mapped to other locations on another board.

DR for CPU/Memory Boards will support generic memory addition/deletion, and will allow *non-interleaved, non-permanent memory* to be relocated. (At boot, memory interleaving may be disabled by setting a system variable, but it may be undesirable to do so because of its performance benefits. Permanent memory is non-pageable memory — some low-level kernel and OBP memory allocations are non-pageable and are thus considered permanent.) The implementation of DR for CPU/Memory Boards will facilitate adding boards to the server while it is running, and will make new processors and memory available to the operating system without a system reboot. CPU/Memory Boards that are dynamically attached can be subsequently detached, even if memory interleaving is enabled in the rest of the system.

Hot Pluggable and Hot Swappable Hardware

Electrically, all core systems components in the Sun Enterprise X500 platforms are *hot pluggable*. This means that core components can be safely inserted and removed from a powered-up system. Dynamic Reconfiguration, as discussed above, will enable I/O and CPU/Memory Boards to be activated and deactivated by the operating system without a system reboot, while users continue to access their data and applications. Assuming a system has been configured for redundant power and cooling, Power/Cooling Modules are already *hot swappable* (see *Redundant Power and Cooling,* Chapter 4*)*.

The insertion of hot swap boards causes the board's pre-charge blades (connector pins), which are slightly longer than the board's other blades, to make contact first with the backplane or centerplane. This results in the initiation of 5v and 3.3v to the board before the signal pins meet. This helps prevent a newly inserted board from dragging down the signal voltage below an acceptable level, and causing a bus error which would otherwise cause a fatal system error.

When the hardware senses that a board has been inserted, the system suspends all activity for a short amount of time (typically less than 200 milliseconds). This short term freeze prevents transactions from taking place during the pin insertion process to avoid the possibility of a signal line skewing and causing data corruption (or other fatal system errors). For most applications, the temporary suspension due to hot swapping should not be noticeable. Only real-time applications with critical response time requirements may be affected by this delay.

The removal of a board requires that the board first be disabled and set to a powered down state. This is a very important point and must be clearly understood: *powered-up, on-line boards cannot be simply removed from a running system*, even with Dynamic Reconfiguration. Boards must first be disabled by the operating system and transitioned to the powered down state. Because of the risks associated with inserting and removing boards, these operations should only be done by qualified Sun service personnel.

When the administrator initiates a dynamic detach and removal of a board, the system switches the board's power off through the setting of several bits in the board's control and status registers. This then turns off the on-board DC-to-DC converters and shuts down the Address Controllers (ACs) and Data Controllers (DCs) so that the board can be safely removed.

In summary, hot swap technology provides these key advantages:

- Adding additional hardware dynamically enables the system to handle an increasingly larger workload, and improves the performance of a running system.
- Removing/replacing faulty hardware with minimal system service disruption means less downtime. This requires the administrator to quiesce activity to the board and to delete the board from operating system use — significantly less time consuming than a full shutdown and cold reboot.

Robust Diagnostics Support

SunVTS™ Diagnostic Tool

To supply robust on-line diagnostic capabilities for Sun Enterprise X500 systems, Sun provides the Sun Validation Test Suite (SunVTS™). The primary goal of the SunVTS software is to create an environment in which Sun systems can be thoroughly tested to provide proper operation and to find elusive problems or avoid imminent failures.

SunVTS offers an easy-to-use graphical user interface (GUI) for initiating and logging test results from various subsystems (e.g., processors, memory, I/O, etc.) or the system as a whole. Some of the key features of SunVTS are:

- *UNIX® level diagnostics*, where system tests execute real UNIX code under the Solaris operating environment
- *Automatic system probing*, enabling the system configuration to be displayed via a user interface
- *Both a GUI-based interface and a character-based interface*. The SunVTS kernel is cleanly separable from the user interface such that multiple user interfaces can communicate with the same SunVTS kernel. The character-based interface permits shell scripts to control SunVTS
- *An Application Programming Interface*, providing a well-defined interface into the SunVTS kernel from other processes, as well as from the user interfaces. SunVTS execution can be initiated in a `cron`-like fashion, with no direct user interface at all
- *Advanced configuration and execution control*, allowing tests to be grouped together based on user requirements, with fine grained execution control for status and logging information

SunVTS Architecture

SunVTS is designed in a modular fashion, providing a clearly defined API into the SunVTS kernel. Functions of the SunVTS kernel include hardware system configuration probing and saving, message logging, maintaining test status, and test scheduling. The SunVTS kernel is implemented as a daemon. The SunVTS daemon spawns threads dynamically to start tests which run as separate processes, and it communicates with the spawned test processes through standard IPC facilities.

The SunVTS user interface provides the means to start up the SunVTS kernel and provides an interface through the SunVTS API. The user interface (either graphical or character-based) can be run from a remote system and multiple user interfaces can communicate with the same SunVTS kernel concurrently, thus enabling the monitoring of test execution from multiple locations.

SunVTS ships with an initial suite of tests designed to test the system's processors, memory and I/O interfaces (e.g., disks, networks, tapes, etc.). These tests can be run independently or bundled together in any combination to stress test the system at different levels. A published API provides the necessary interfaces for the development of new tests, such as for third-party I/O interfaces. A documented set of guidelines, including message formats, as well as utilities to integrate new tests into the SunVTS Shared Object Library, minimizes the effort involved in adding new tests and provides a means by which test message output can have a common look and format.

Advanced Systems Monitoring and Management

Performance management is crucial to overall system availability. System performance can make a system appear to be "down" to end users — substandard response times become indistinguishable from a system that is unavailable. As a result, hardware monitors and powerful system management tools become an important component in system availability.

Hardware Monitoring

The Sun Enterprise X500 platform is designed with several hardware registers and counters that are used for monitoring system activity. These registers are located in strategic locations throughout the system and are accessible with software tools such as Enterprise SyMON to extract performance-related information.

The internal registers monitor the use and efficiency of the following hardware subsystems:

- *UltraSPARC processor*
  - Cache hit rates
  - Translation Lookaside Buffer (TLB) hit rates — the TLB is effectively a cache for virtual-to-physical address translations
- *UPA & Gigaplane* counter information[1] on:
  - Memory bank reads/writes
  - System-wide packet counts (data, address, control, etc.)
  - Per-port packet counts (read, write, etc.)
  - Gigaplane utilization
- *SBus and PCI interface* statistics[2] on:
  - DVMA read/write transfers
  - Bytes transferred
  - TLB misses
  - Interrupt count
  - PIO read/writes
  - Cycles used for DVMA & PIO

Using system management tools like Enterprise SyMON, system administrators and performance analysts can access detailed technical information and locate system bottlenecks. They can achieve a better understanding of application behavior and the load an application imposes on the system.

1. This data is extracted from the Address Controller (AC) ASIC. An AC exists on every core system board regardless of type, and interfaces the Gigaplane to two UPA ports

2. This data is extracted from the SYSIO ASIC, which is the I/O bridge between the UPA and the SBus, or from the PSYCHO ASIC, which is the I/O bridge between the UPA and PCI bus

Hardware Watchdog

The hardware watchdog mechanism is a time-out facility designed to prevent data corruption and avoid system “hard hang” conditions. It does so by watching selected system events and not letting them take an unacceptable time to complete.

Like the hardware counters, the hardware watchdog mechanisms are located strategically around the system. To prevent data corruption, the system watches for fatal data-related errors and forces a system reset when they occur. For example, a parity error on the Gigaplane address lines will force a reset condition on the system to prevent corrupt data from flowing through the system or getting written to disk.

Other hardware watchdogs are used. A TOD (Time-of-Day) watchdog resets the system if the operating system software fails to respond at regular intervals. This prevents the system from hanging as a result of the operating system sitting in a tight loop. Additional time-out protection exists within the Address Controllers (ACs) that interface the UPA to the Gigaplane. A programmable timer is used to detect UPA Master Port Time-outs (which are fatal errors) and force a system reset. The timer will expire when requests for a port are queued up and no requests have been serviced in the time-out period. The synchronous design of the Sun Enterprise X500 systems increases the predictability of service request times for Gigaplane transfers. An excessive delay on servicing a request is indicative that something is seriously wrong, and that operator intervention is required.

Enterprise SyMON 2.0

System managers face an increasingly difficult task in managing highly distributed computing resources and applications. Monitoring distributed systems, networks, and applications is key in providing predictable service levels and maintaining systems and application availability.

To address this critical need for enterprise-level manageability, Sun has developed Enterprise SyMON 2.0. SyMON 2.0 is an open, extensible, and standards-based server management solution that facilitates enterprise-wide management of Sun Enterprise servers, their subsystems, and components. SyMON 2.0 leverages some of the user-level system management capabilities of the previous SyMON system monitor, but has been totally re-architected to use Java™ technology and the SNMP protocol.

An Extensible Framework for Enterprise System Management

SyMON's new object-oriented, three-tiered architecture provides a comprehensive system management framework, enabling *a single point of management for all Sun elements* — Sun Enterprise servers, storage, and the Solaris operating environment. The framework is fully extensible, allowing SyMON to be integrated with other network and system management platforms. To monitor business-critical ERP or database applications, new SyMON agents can be easily implemented.

Figure 3-2 shows an architectural overview of Enterprise SyMON, revealing three major layers in its architecture:

- The console layer
- The server layer, and
- The agent layer

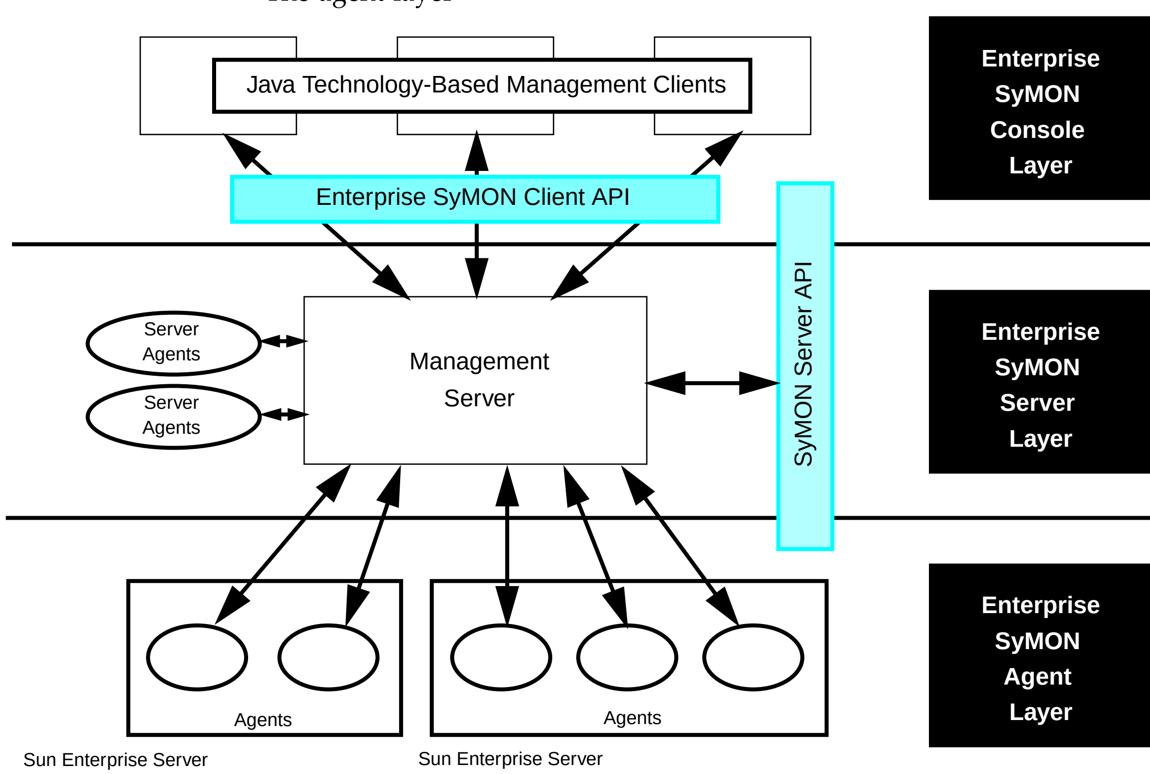

Figure 3-2 Enterprise SyMON 2.0 architecture

At the SyMON console layer, the clients are Java technology-based, providing platform-independent, network-accessible, browser-based access to system management information. The clients display the managed objects and allow system administrators to manipulate object attributes and properties. As shown, multiple client views can exist for different users, enabling distributed as well as centralized system management.

At the core of the SyMON server layer is the SyMON server itself. Written in the Java language and multi-threaded, the server intelligently processes and routes requests between console clients and agents. One of the server's advantages is its ability to recognize redundancy — it recognizes when identical information requests are made independently by different clients and merges these requests. In multi-user environments, this can yield a significant savings in systems resources.

The SyMON server also dispatches events and alarms to the subscribing consoles. The server relates incoming alarms to client jobs in a process known as *Trap Correlation*. By identifying jobs that are affected by system events, the server can immediately re-issue these requests and update SyMON clients immediately after a system event. This has the important effect of making the polling-based mechanisms of SNMP appear to be subscription-based to the SyMON clients.

The SyMON agent layer consists of the SyMON agents and probe components deployed on the managed nodes. SyMON agents manage objects by intelligently and autonomously collecting and monitoring data. The agents use rule-based technology to determine the status of the managed objects. SyMON agents store the data and status of managed objects in a Management Information Base (MIB) — this information can then be accessed by the SyMON server via standard SNMP protocol. The agent may automatically generate alarms or perform actions based on the detected conditions, thereby enabling predictive failure capabilities and proactive system management.

SyMON Features and Benefits

Enterprise SyMON 2.0 provides the following features and benefits:

- *A topology viewer* shows graphical representations of Sun servers throughout the enterprise. A "discovery" process automatically builds the topology view, which can be divided into several administrative domains for distributed management

- *A physical system viewer with per-device information* provides graphical "pictures" of system components, with relevant information such as network interface types, disk types, processor module speeds, etc.
- *A logical system viewer* shows the system components and their relationships in a symbolic, hierarchical representation. Because they provide quick, easy-to-assimilate graphical representations of the network and specific server configurations, the topology viewer, physical system viewer, and logical system viewer are especially useful to administrators who may not be familiar with Sun servers throughout the enterprise
- *A health monitor* displays the general health of the enterprise, the administrative domain, and each server and its components. The health monitor allows an administrator to identify problems at a glance, and enables predictive failure. For example, an administrator might receive early warning when the temperature of a CPU/Memory Board hits a certain level, before CPU Power Control takes effect and disables the processors
- *Event/alarm management* includes the ability for events to trigger emails, to dial pagers, etc. Enterprise SyMON can be easily configured to proactively respond to problems
- *A per-process display* extracts specific information on process resource usage and behavior. An administrator can easily monitor the load an application imposes on the system for the purposes of performance analysis, capacity planning, etc.
- *A graph viewer* displays CPU, memory, disk, and network performance metrics. This viewer allows an administrator to chart trends or correlate application behavior, enabling more effective performance analysis
- *Historical data logging* provides agent, intermediate, or server-level logging. Log file scanning enables efficient free-text searches of any log file, helping administrators identify recurring problems
- *On-line diagnostics* (via an interface with SunVTS) enables troubleshooting directly from SyMON
- *A configurable, customizable, and extensible framework* allows new agents to be easily developed and facilitates the integration of SyMON with other system management tools. This framework allows administrators to monitor additional applications, or to use familiar system management tools

Enterprise SyMON 2.0 greatly simplifies the administration and management of Sun Enterprise systems. Administrative domains allow system management functions to be safely distributed and delegated throughout the enterprise. By monitoring system resources and setting thresholds which trigger events, administrators can automatically be notified when a serious condition exists. In many cases administrators can react to such conditions and take action before users even realize anything is wrong.

Enterprise SyMON's hardware failure prediction and capacity tools increase system reliability by alerting administrators of a potential system component failure. By taking advantage of the new Alternate Pathing and Dynamic Reconfiguration features, administrators can dynamically detach faulty components and remove them while the system continues to run. With Enterprise SyMON, administrators will be able to initiate DR operations and easily keep track of configuration changes.

The overall effect of better systems monitoring and management with Enterprise SyMON is that systems will be more *available*.

Remote Console Support

The Sun Enterprise X500 systems include extra hardware on the Clock Board to support remote console operations. The remote console interface is made up of a hardware state machine that is initialized at boot time and does not require any level of operating system support when the system is up and running. The interface scans the incoming serial data stream, looking for special character sequences that can be used to interrupt, reset, or power-on-reset the system from a remote location.

The commands themselves are made up of a series of uncommon keystrokes and must be spaced in a time critical fashion (more than 0.5 seconds but less than 2.0 seconds between keystrokes). The actual commands recognized by the state machine are:

- *Software Reset* — Resets console software (does not cycle power)
- *XIR Reset* — Preserves caches and memory. Useful for a "hung" system where it is necessary to "unjam" the system to perform diagnosis and it is desirable to preserve as much state as possible
- *Power Cycle Reset* — Initiates a power cycle (powers down a running system)

Sun’s remote console functionality allows for true “lights out” management of a Sun Enterprise X500 system, giving systems administrators more flexibility in where they are physically located. Remote console support, along with the browser-based Enterprise SyMON client, will allow administrators to do their jobs regardless of location.

Configuring Sun™ Enterprise™ X500 Servers for Maximum Availability 4

A system's configuration is critical in providing higher availability. Most importantly, the system must be configured with redundant components so that there is no single point of failure. The layout of components is also important, especially to take advantage of new availability enhancements like Alternate Pathing and Dynamic Reconfiguration. This chapter introduces some common configuration issues in configuring for high availability.

Configuring for Redundancy

Optional Component Redundancy

The Sun™ Enterprise™ X500 systems feature optional component redundancy. All systems may be configured with multiple system boards, processors, memory banks, network and I/O controllers, and redundant power and cooling. The first step in configuring for availability is to configure the system for no single point of failure (SPOF).

At the extreme, entire systems may be redundant, as in the Sun Enterprise Cluster configurations. Redundant components, however, can often provide significantly enhanced system availability more economically. Configuring extra system I/O Boards, disk and network controllers, CPU/Memory Boards, and processors can leverage new software availability features, in particular CPU Power Control, Alternate Pathing, and Dynamic Reconfiguration (described in Chapter 3).

Customers that have multiple Sun Enterprise X500 servers can also take advantage of component sharing between different members of the Sun Enterprise X500 server family. Since the Sun Enterprise 3500, 4500, 5500, and 6500 share the same components — system boards, processors, memory, SBus™/PCI cards, and Power/Cooling Modules — components can be easily moved from one member of the Sun Enterprise server family to another. This might be useful, for example:

- To replace a failed component in a mission-critical system
- To upgrade components without discarding the replaced components — instead migrating them to another machine, extending their useful life
- To temporarily increase the number of processors (e.g., borrowing populated CPU/Memory Boards from a smaller development machine and installing them in a production server)

Redundant Power and Cooling

One method of greatly increasing the availability of Sun Enterprise X500 servers is simply to configure them for redundant power and cooling. In this way, the server can survive a power supply or fan failure without interrupting system operations. The power system of the Sun Enterprise X500 servers implements an innovative *current share design,* which effectively ties the outputs of several smaller power supplies together, making them appear to system components as one large power supply.

Sun Enterprise X500 systems have three different types of power supplies: a 300 watt Power/Cooling Module (PCM), a 184 watt Peripheral Power Supply (PPS), and an optional 195 watt Peripheral Power Supply that is used only in the Sun Enterprise 3500 for redundant system and peripheral power. Generally, each 300 watt PCM provides power for two adjacent boards slots, supplying 5v (Volts DC), 3.3v, and 2.0v to the system boards.

Each PCM also has two fans to provide cooling for the two adjacent board slots. The fans in the power supplies receive redundant power from the system's centerplane. Even if the power supply fails, the fans continue running, providing continuous cooling to the two adjacent slots. The output of the temperature sensors on the CPU/Memory and I/O Boards is fed back to the adjacent PCM to control the speed of the fans. Hotter boards will cause the fans to spin faster. If two boards are present, the warmer of the two temperatures is used to set the fan speed.

Because of the current share design of the system, any board slot can be powered from any PCM in the system. On the 5-slot Sun Enterprise 3500, when five system boards are installed, three PCMs and the second Peripheral Power Supply provide redundant power for the system boards and the internal disk drives. When the 8-slot Sun Enterprise 4500 and Sun Enterprise 5500 servers are fully configured with 4 PCMs, and the 16-slot Sun Enterprise 6500 system is fully configured with 8 PCMs, there is still enough extra current if any one power supply fails.

Figure 4-1 illustrates the power distribution of the Sun Enterprise 4500, 5500, and 6500.

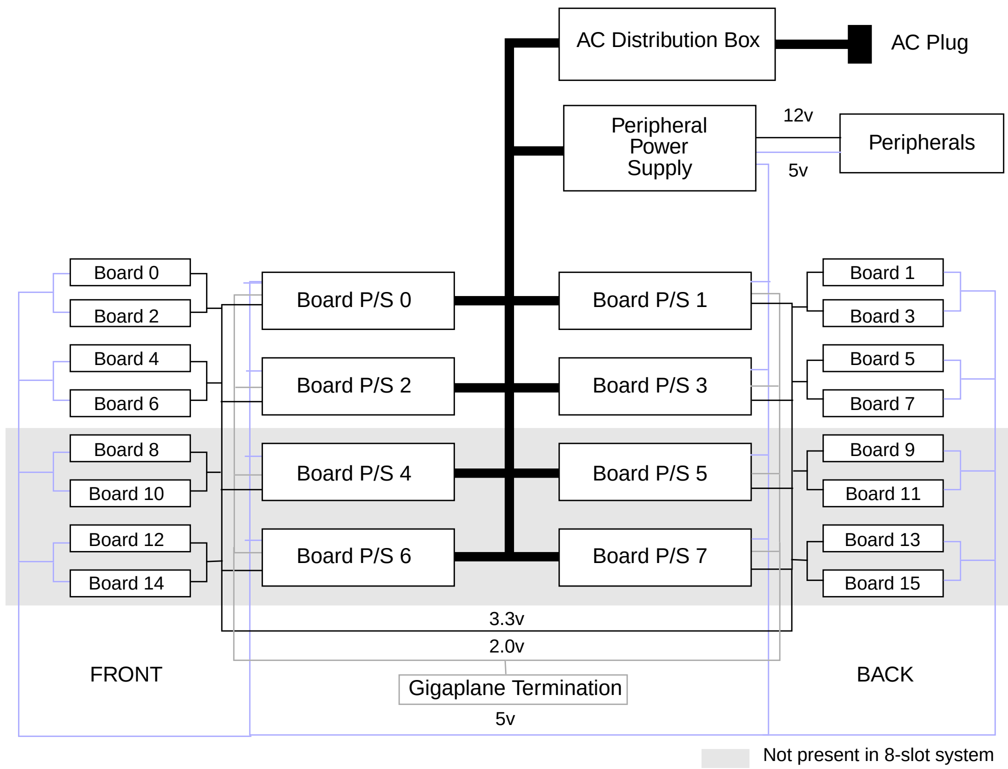

Figure 4-1 Power distribution for 8-slot and 16-slot Sun Enterprise X500 systems

Figure 4-2 illustrates the current-sharing power distribution model of the Sun Enterprise 3500.

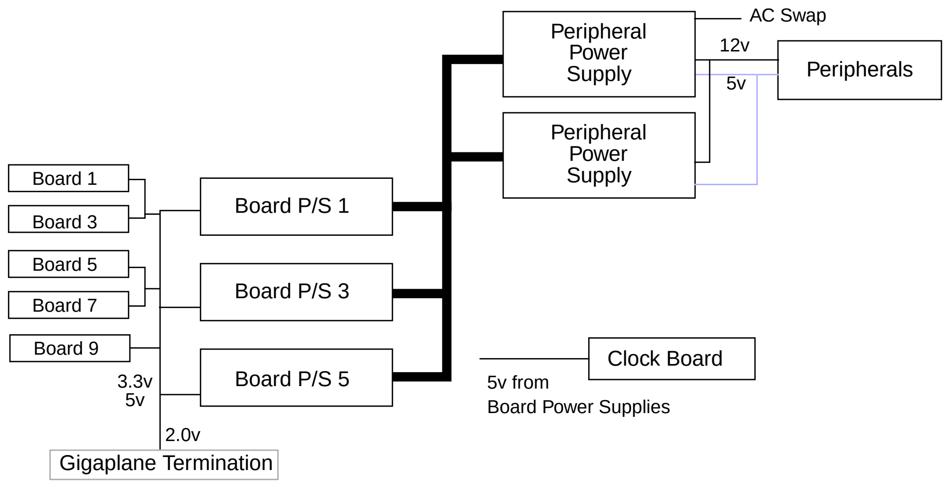

Figure 4-2 Power distribution for 5-slot Sun Enterprise 3500 system

The 184 watt PPS provides power for SCSI peripherals in the system (12v and 5v), as well as additional 5v for system boards, which supplements the 300 watt supplies. In the 5-slot Sun Enterprise 3500 system, the outputs of the two PPSs are current shared to provide redundant power to the disks in the cabinet. The peripheral supplies also provide pre-charge power for hot swap components and auxiliary power for the hardware remote console functionality on the Clock Board.

On the Sun Enterprise 4500, 5500, and 6500, there is no redundant +5v and +12v for the SCSI tray. The SCSI tray in these systems typically houses one tape drive and one CD-ROM, which are not considered critical to system operation. Disk drives in these systems are independent rack mount products with their own power supplies.

In order to take advantage of redundant power capabilities, the correct number of PCMs must be ordered with a system. This involves determining the number of supplies required for powering and cooling the system boards ordered, and then ordering an extra supply. Generally this is called an *N+1 power configuration* — one PCM is ordered for every 2 system boards, plus an extra PCM for redundancy.

Additional Cooling

In addition to the PCM fans, Sun Enterprise X500 systems include additional cooling fans to provide air flow to other areas of the cabinet. Additional fans are installed in the system cabinet to increase air flow to the rack that houses the centerplane and system boards. Multiple cabinet fans are used to protect against the failure of a single rack fan, and there is redundant power supplied to these fans as well.

The AC input box contains redundant fans that provide cooling to the Peripheral Power Supplies. Note that the 5-slot Sun Enterprise 3500 system has a fan module above the peripheral power supplies so that cooling is provided when an extra 300 watt PCM is not configured (see above, *Redundant Power and Cooling)*. This fan module ships standard with every Sun Enterprise 3500 system.

A fan is also housed in the system key switch assembly, which provides cooling for any internal SCSI devices (e.g., CD-ROM and/or tape). There is no redundancy on this fan since the components it cools are not considered critical to system operation.

Over-Temperature Sensing

The temperature sensors on the system boards provide feedback to the fans to determine fan speed, and can signal an over-temperature condition. This is usually the result of a Power/Cooling Module (PCM) problem. If an over-temperature condition is sensed on a CPU/Memory Board, the fans on the associated PCM will spin faster, trying to provide adequate cooling. If the temperature continues to rise, the system will actually power down the processors on the board. Powering down the processors eliminates the heat source and keeps the system up and running. This feature is called *CPU Power Control* (described in Chapter 3). Without this feature, the system would immediately initiate a system shutdown.

Note that CPU Power Control powers down *both* processors on a board whenever an over-temperature condition is sensed. If the over-temperature condition is reached as a result of a PCM failure, and the two boards associated with that PCM are both CPU/Memory Boards, CPU Power Control could potentially power down processors on *both boards* if both boards' temperature sensors exceed a safe operating level. For this reason, especially in configuring a system with a small number of processors, it is preferable to stagger CPU/Memory Boards with I/O Boards.

Redundancy in the Sun Enterprise 3500 Design

The Sun Enterprise 3500 is a re-design of the previous Sun Enterprise 3000 system, with enhanced performance and availability features. The redesigned Sun Enterprise 3500 cabinet provides a total of five system slots, the new 100MHz-capable Gigaplane™ system interconnect, and up to 8 internal FC-AL (Fibre-Channel Arbitrated Loop) disks, all in the same footprint as the previous Sun Enterprise 3000 design.

The Sun Enterprise 3500's internal FC-AL architecture includes two separate banks of disks, each containing four FC-AL disk drives (Figure 4-3). To provide the 100MB/second fibre channel interface for the internal disks, a FC-AL Interface board contains up to four GBIC (GigaBit Interface Convertor) modules. (These GBIC modules are identical to the ones used in the Sun™ StorEdge™ A5000, the FC-AL SBus host adapter, and on the Sun Enterprise X500 SBus I/O Board.) The four GBICs on the FC-AL Interface Board are logically independent. The GBIC modules get their power through a single connection (the FC-AL Interface Board) but the power to the board comes from the backplane, which is supported by redundant power supplies.

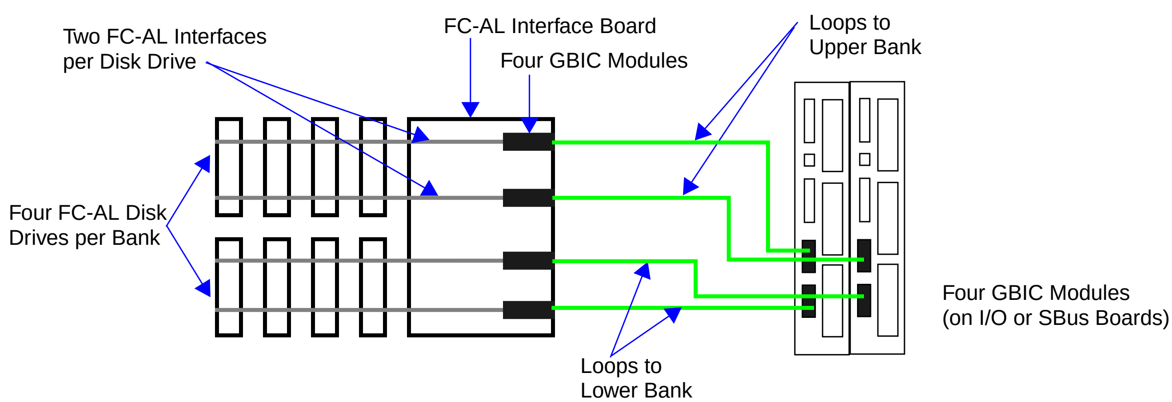

Figure 4-3 Sun Enterprise 3500 FC-AL Interface Board in dual loop configuration

Each of the two FC-AL disk banks can have one or two FC-AL loops connecting to the installed drives. A *dual loop* configuration provides redundant paths to the data, and thus a more highly available configuration. If either loop fails, there is still access to the disks via the other loop. To optimize availability, the data should also be mirrored to protect against a disk failure.

The FC-AL disks are, of course, hot swappable, so they can be easily replaced and the mirror re-established while the system is running. In a dual loop configuration, the Sun Enterprise 3500 provides a robust data storage solution.

Configuring for Alternate Pathing and Dynamic Reconfiguration

To maximize system availability and take advantage of Alternate Pathing and Dynamic Reconfiguration, it is necessary to configure an Sun Enterprise X500 system appropriately. Component redundancy is important, but so is the proper layout of components. To promote greater availability, the mirroring of data also helps to protect against a disk failure.

Redundant Paths

Configuring an Sun Enterprise X500 system with redundant I/O paths is required to take advantage of Alternate Pathing. Figure 4-4 shows a sample Sun Enterprise X500 configuration with mirrored disks in a storage array. Note that this configuration includes redundant controllers, but they are placed on the same system I/O Board — a single point of failure. This configuration would *not* be able to take advantage of the Sun Enterprise X500's Alternate Pathing and Dynamic Reconfiguration capabilities.

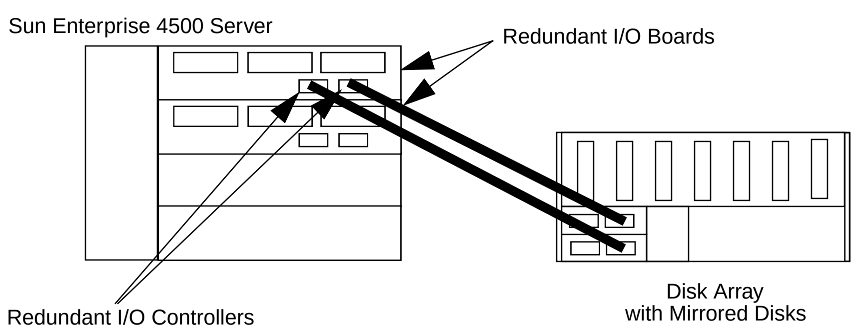

Figure 4-4 Redundant controllers on the same I/O Board

Figure 4-5 illustrates a more highly available configuration, where the two controllers are placed on separate system I/O Boards. In this configuration, *two completely separate paths exist to the data*, thus enabling the use of Alternate Pathing and Dynamic Reconfiguration for the I/O Boards.

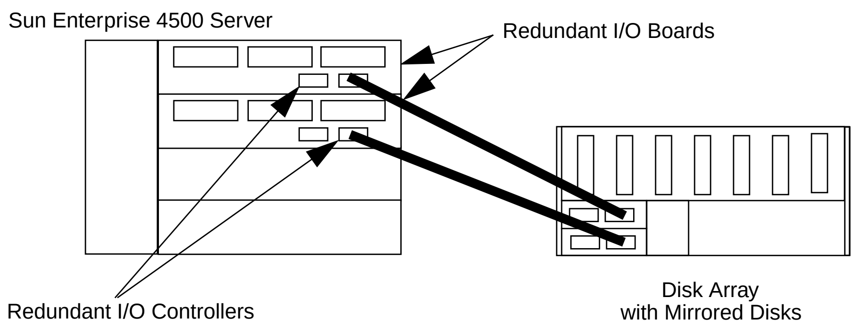

Figure 4-5 Redundant controllers placed on different I/O Boards

Recall from Chapter 3 that some device drivers — those that support only SUSPEND and RESUME operations — may be resident in a system even though DR cannot be performed on the I/O board on which they reside. Recall also that DR cannot be performed on the central I/O board (because it provides an electrical path to devices on the clock board). For this reason, it may make sense to group SUSPEND/RESUME-only devices on the central I/O board. This placement will allow DR to be performed on other I/O boards.

Even in a system configured with a single I/O Board — a configuration where AP and DR cannot be used — careful component placement can still promote a higher degree of availability. For example, each SBus I/O Board houses two separate SBus channels (Figure 4-6) controlled by separate SYSIO ASICs. One channel carries SBus traffic for the FC-AL interfaces and two of the SBus slots, while the other supports a single SBus slot, network traffic, and the on-board SCSI interface.

Configuring an SBus FC-AL interface card in SBus Slot 0 (along with using the on-board FC-AL interface) to connect to a disk array protects against the failure of a SYSIO ASIC or a SOC (FC-AL interface) chip. In addition to better availability, this configuration also balances the I/O load across the SBus channels, which may result in better performance.

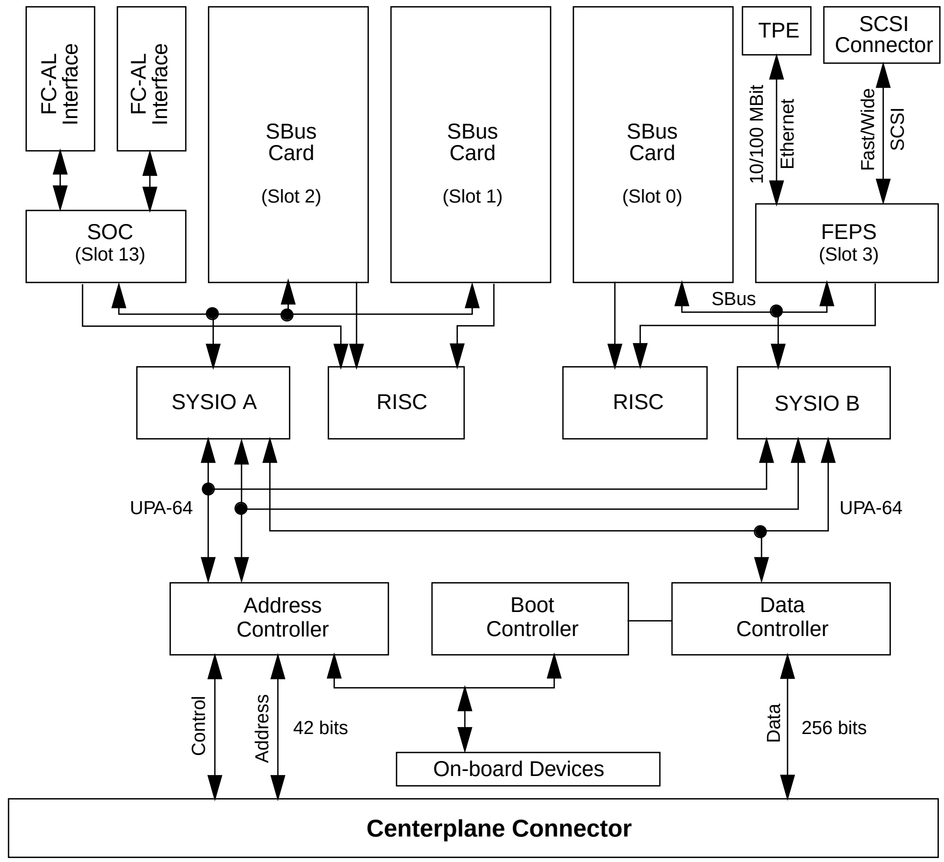

Figure 4-6 The SBus I/O Board contains two separate SBus channels

Like the SBus I/O Board, the PCI I/O Board also includes two separate 66-MHz, 64-bit PCI channels, each supporting a single PCI slot. Redundant PCI cards — each on their own PCI bus — may provide some additional availability protection and performance in a single I/O board configuration.

Of course, redundant I/O Boards eliminate a single point of failure in a system's configuration. With redundant I/O Boards and the proper layout of redundant components on those boards, a configuration can potentially support Alternate Pathing and Dynamic Reconfiguration as well.

Redundant Data

Redundant data is not required to take advantage of Alternate Pathing — as long as alternate *paths* to the data exist, the AP software enhancements can be used. However, redundant data storage is key to optimizing availability.

Generally data mirroring (RAID 1) provides the highest level of availability along with a reasonable level of performance. Striping (RAID 0) is often used in conjunction with mirroring to improve performance. Sun offers disk management software (Sun™ StorEdge™ Enterprise Volume Manager™ and Solstice™ DiskSuite™) to mirror and/or stripe data for better availability. (For a complete discussion of RAID levels, and Sun StorEdge data storage and management solutions, please consult the white papers listed under *References* at the end of this paper.)

Careful placement of controllers in relation to mirrored volumes is important in configuring for availability. Consider the example shown in Figure 4-7.

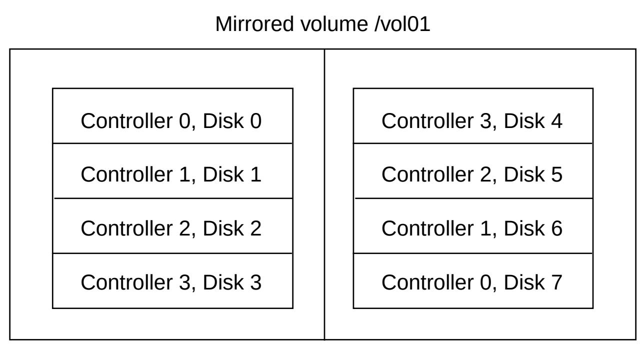

Figure 4-7 Controllers accessing both sides of a mirrored volume

In this configuration, all four controllers are assumed to be on separate I/O Boards. Disk 0 is mirrored to disk 4, disk 1 to disk 5, etc. Dynamic Reconfiguration cannot be used in this configuration, because neither side of the mirror can be quiesced in order to detach an I/O Board with a disk controller. To create a more highly available configuration, either additional controllers are necessary, or the mirror must be rearranged among the existing controllers. Redundant data, redundant paths to the data, and the layout of those paths are important considerations in configuring for optimal availability.

Alternate Path to Vital System Data

Recall from Chapter 3 that Dynamic Reconfiguration requires that the operating system quiesce all activity to a board before it can be unconfigured. In the case of an I/O Board that hosts the controller for the boot disk, there must exist an alternate path to the system data partitions. For example, if the I/O Board hosts a controller attached to a disk with the root or /usr partition, the board cannot be detached unless there is an alternate hardware path to the disk or the disk is mirrored. (For better availability, the disk should be mirrored.) The alternate path and the mirror must be hosted by other I/O Boards and controllers. Multiple swap partitions, again hosted by other I/O Boards and controllers, are also recommended.

The same logic applies to network controllers. The board that hosts a controller for a network connection cannot be detached unless an alternate path for this network exists via a controller on another board.

In summary, redundant components are necessary to eliminate single points of failure and to create more highly available configurations. Careful placement of redundant components is essential to configure for Alternate Pathing. When alternate paths are defined for network and I/O devices, the administrator can switch on-going traffic to an alternate path, enabling an I/O Board to be detached easily with Dynamic Reconfiguration and removed from a running system.

Further Planning for High Availability

Configuring the system appropriately is the first step in increasing its availability. Planning the layout of data mirrors, disk controllers, and other system components are all important considerations.

Just as important, however, is the required training and planning for what to do in the event of a failure. Since failures happen infrequently, administrators need carefully documented and tested procedures to follow. Failure scenarios should be reviewed periodically and revised as necessary.

Sun offers a comprehensive set of service, education, and consulting services to assist customers. Sun Enterprise Services even provides an availability guarantee customized to a customer's requirements as a part of its mission-critical, 7x24 service offering. Sun Education offers a full complement of training courses. Sun Enterprise Services can help with defining an availability architecture, configuring systems, and designing and implementing procedures for highly available operations.

Conclusion 5

Users continuously demand higher and higher levels of service from their Information Technology organizations — they seek better applications and systems availability in addition to improved levels of system performance. Increasingly, customers require data and applications to be available around the clock to satisfy global operations or to meet batch processing requirements. Availability has become a requirement not only for mission-critical applications, but throughout the enterprise, from department to data center.

Sun's strategy is to offer outstanding availability and manageability, in addition to performance and scalability. In the Sun™ Enterprise™ X500 servers, integrated hardware features combine with innovative software enhancements, resulting in significant improvements in system reliability, serviceability, and overall availability.

References

Sun Microsystems posts product information in the form of data sheets, specifications, and white papers on its Internet World Wide Web Home page at: *http://www.sun.com*.

Look for abstracts on these and other Sun technology white papers:

The Sun Enterprise 3500-6500 Server Family: Architecture and Implementation, Technical White Paper, Sun Microsystems, Inc., 1998.

Sun RAS Solutions for Mission-Critical Computing, A White Paper, Sun Microsystems, Inc., October 1997.

The Sun Enterprise Cluster Architecture, Technical White Paper, Sun Microsystems, Inc., October 1997.

Ultra Enterprise Cluster HA 1.3 Architecture, Technical White Paper, Sun Microsystems, Inc., 1997.

The Ultra Enterprise 10000 Server, A White Paper, Sun Microsystems, Inc., 1997.

Ultra Enterprise 10000: Dynamic System Domains, Technical White Paper, Sun Microsystems, Inc., 1997.

Ultra Enterprise 10000 Server: SunTrust Reliability, Availability, and Serviceability, Technical White Paper, Sun Microsystems, Inc., 1998.

Consolidation and Recentralization with the Ultra Enterprise 10000 Server, Technical White Paper, Sun Microsystems, Inc., 1997.

Sun's Strategy for Storage Network Management: Enterprise Storage Software Overview, A White Paper, Sun Microsystems, Inc., 1998.

Reliability, Availability, and Serviceability in the Sun StorEdge A5000 Disk Array, Technical Brief, Sun Microsystems, Inc., 1998.

The Sun StorEdge A7000 Intelligent Storage Server Architecture, Technical White Paper, Sun Microsystems, Inc., 1998.

Solstice SyMON System Monitor, Technical White Paper, Sun Microsystems, Inc., 1996.

Availability Features in the Sun Enterprise 3500 to 6500 Server Family, Technical White Paper, Sun Microsystems, Inc., 1998.

In addition to white papers, the following product documentation is available:

Sun Enterprise Servers Alternate Pathing User's Guide, Sun Microsystems, Inc., 1998.

Dynamic Reconfiguration User's Guide, Sun Microsystems, Inc., 1998.

Sun Microsystems, Incorporated
901 San Antonio Road
Palo Alto, CA 94303 USA
650 960-1300
FAX 650 969-9131
http://www.sun.com

Sales Offices

Argentina: +54-1-317-5600
Australia: +61-2-9844-5000
Austria: +43-1-60563-0
Belgium: +32-2-716-7911
Brazil: +55-11-5181-8988
Canada: +905-477-6745
Chile: +56-2-638-6364
Colombia: +571-622-1717
Commonwealth of Independent States:
+7-502-935-8411
Czech/Slovak Republics:
+42-2-205-102-33
Denmark: +45-44-89-49-89
Estonia: +372-6-308-900
Finland: +358-9-525-561
France: +33-01-30-67-50-00
Germany: +49-89-46008-0
Greece: +30-1-680-6676
Hong Kong: +852-2802-4188
Hungary: +36-1-202-4415
Iceland: +354-563-3010
India: +91-80-559-9595
Ireland: +353-1-8055-666
Israel: +972-9-956-9250
Italy: +39-39-60551
Japan: +81-3-5717-5000
Korea: +822-3469-0114
Latin America/Caribbean:
+1-650-688-9464
Latvia: +371-755-11-33
Lithuania: +370-729-8468
Luxembourg: +352-491-1331
Malaysia: +603-264-9988
Mexico: +52-5-258-6100
Netherlands: +31-33-450-1234
New Zealand: +64-4-499-2344
Norway: +47-2218-5800
People's Republic of China:
Beijing: +86-10-6849-2828
Chengdu: +86-28-678-0121
Guangzhou: +86-20-8777-9913
Shanghai: +86-21-6247-4068
Poland: +48-22-658-4535
Portugal: +351-1-412-7710
Russia: +7-502-935-8411
Singapore: +65-438-1888
South Africa: +2711-805-4305
Spain: +34-1-596-9900
Sweden: +46-8-623-90-00
Switzerland: +41-1-825-7111
Taiwan: +886-2-514-0567
Thailand: +662-636-1555
Turkey: +90-212-236 3300
United Arab Emirates:
+971-4-366-333
United Kingdom: +44-1-276-20444
United States: +1-800-821-4643
Venezuela: +58-2-286-1044
Worldwide Headquarters:
+1-650-960-1300

Printed in USA