

Sun Fire™ 3800-6800 Servers Just the Facts

Version 1.10



Copyrights

©2001 Sun Microsystems, Inc. All Rights Reserved.

Sun, Sun Microsystems, the Sun logo, Solaris, Sun Fire, SunSolve, SunSolve EarlyNotifier, Sun Enterprise, Sun StorEdge, Solaris Resource Manager, SunVTS, SunSpectrum Platinum, SunSpectrum Gold, SunSpectrum Silver, SunSpectrum Bronze and OpenVIS are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States and other countries.

All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. in the United States and other countries. Products bearing SPARC trademarks are based upon an architecture developed by Sun Microsystems, Inc.

UNIX is a registered trademark in the United States and other countries, exclusively licensed through X/Open Company, Ltd.



Table of Contents

Positioning.....	5
Introduction.....	5
Product Family Placement.....	5
Key Messages.....	6
Key Features and Benefits.....	7
Availability Date.....	8
Target Markets and Users.....	8
Selling Highlights.....	11
Market Value Proposition.....	11
Applications.....	11
Compatibility.....	11
Reliability Availability and Serviceability.....	13
Reliability.....	13
Availability.....	13
Serviceability.....	17
Auto Diagnosis and Recovery Features.....	18
UltraSPARC™ III Processor.....	19
System Architecture.....	22
Sun Fire™ 3800-6800 Server Key Facts.....	22
Technical Fact Summary.....	23
Two-Level Sun™ Fireplane Interconnect System.....	24
CPU/Memory Board.....	27
System I/O - High Performance PCI and cPCI Technology.....	29
System Controller.....	33
System Cooling.....	35
System Power.....	36
Sun StorEdge™ D240 Media Tray.....	38
Sun Fire Cabinet.....	38
Installation Data.....	41
Cabinet.....	41
Electrical and Thermal Specifications.....	41
Environmental Specifications.....	42
Regulations.....	42
Resource Management.....	43
Dynamic System Domains.....	43
Solaris Resource Manager™.....	45
Processor Sets.....	46
System Management.....	47
System Controller for System Management.....	47
Sun Management Center 3.0 Software.....	47
System Diagnostics.....	48
Automatic System Recovery.....	49
Remote Monitoring through Sun Remote Services.....	49



Software	50
Solaris™ Operating Environment.....	50
Scalability.....	50
Dynamic Reconfiguration.....	50
IP Network Multipathing.....	50
Multiplexed I/O (MPxIO).....	51
Availability Features.....	51
Global Language Support Enhancements.....	52
Application Compatibility and Investment Protection Guarantee.....	52
Capacity on Demand (COD 2.0)	52
Ordering Information	54
Standard Configurations.....	54
Configuration Rules.....	54
PCI and CompactPCI Adapters.....	56
Upgrades	57
Key Messages.....	57
Sun-to-Sun Upgrades (Trade-ins) to the Sun Fire Server Family.....	58
Server Consolidation Trade-in.....	58
System Upgrades.....	58
CPU and Memory Upgrades.....	59
Ordering Process for Chassis and Base System Upgrades.....	59
Service and Support	60
Warranty.....	61
EDU Offerings.....	62
PS Offerings.....	62
SunReady™.....	62
Sun Enterprise 3500-6500 Servers and Sun Fire 3800-6800 Servers Differences	63
Sun Fire 3800-6800 server family versus Sun Enterprise 3500-6500 Family.....	63
Glossary	65
Materials Abstract	66
Change Log	68



Sun Fire™ 3800-6800 Server Positioning

Introduction

Sun Fire™ - The First Midframe Server

Sun Microsystems, Inc.'s Sun Fire™ 3800-6800 family of servers marks the next generation in the evolution of Sun's open, binary-compatible, symmetric multi-processing (SMP) family of server products. Sun has been producing UNIX® server systems for over fifteen years and SMP server systems for nearly ten years. With Sun Fire servers, Sun continues producing a complete family of server products built around performance, maximum availability and high flexibility.

The Sun Fire servers build on the successes and the lessons learned through this long history. In particular, the Sun Fire continues the success of the previous generation of Sun servers, the Sun Enterprise™ 3500-6500 servers and the Sun Enterprise 10000 server. The Sun Enterprise family of servers has been one of the most successful lines of server products ever produced, receiving many awards and setting dozens of performance records. Sun Fire servers continue with many of the features found in the Sun Enterprise family. Additionally, features previously found only in the top-of-the-line Sun Enterprise 10000 server are provided throughout the Sun Fire midrange server family. These features include, Dynamic System Domains(DSDs), concurrent maintenance, online upgrades and full hardware redundancy. Sun Fire servers have new features such as Hot CPU Upgrades, unique to this family.

Like its predecessors, Sun Fire 3800-6800 servers are a family of general-purpose, shared-memory, SMP systems based on the UltraSPARC™ processor. And like the Sun Enterprise family, Sun Fire servers provide packaging and pricing to suit a wide variety of business environments and computing requirements. The Sun Fire 3800-6800 server family is designed to fit the needs of both the commercial and high-performance computing markets. It balances the availability and manageability features required to meet today's commercial computing needs with the demanding requirements of the high-performance computing community.

Product Family Placement

The Sun Fire 3800-6800 server family offers availability, scalability and flexibility not found in any other vendor's solution. Sun Fire 3800-6800 servers extend the current capabilities of Sun Enterprise servers to form the industry's most comprehensive server family. Sun Fire 3800-6800 servers are positioned at the heart of this family.

With the addition of features such as multiple Dynamic System Domains, the new Sun Fireplane™ interconnect system and UltraSPARC III processor architecture, the Sun Fire 3800-6800 server family brings the availability, resource management and performance previously available only to Sun Enterprise 10000 and mainframe customers to a scalable midframe server line and hence to a broader range of customers.

Key Messages

Sun Fire 3800-6800 servers offer excellent availability, resource management, performance and scalability.



- **Mainframe class features in midrange servers**

With an entry price at less than \$100,000, this family of servers will come with unparalleled availability and leading resource management features such as Dynamic System Domains, remote management via Sun™ Management Center 3.0 and Solaris 8 Operating Environments support.

- **Scalability**

In today's rapidly changing marketplace, maintaining flexibility is vital to continued success. Those who are able add capacity, performance and bandwidth with the minimum amount of complexity will be able to respond more effectively to change and to anticipate its effects.

Sun delivers the architecture on which you can build a massively scalable application environment because it allows you to scale both vertically and horizontally. Horizontal scalability enables you to increase capacity by replicating applications such as Web servers on multiple systems. Vertical scalability enables you to increase the resources within a domain to support applications that are not easily replicated, such as backend databases.

The Sun Fire servers have been designed from the ground up to scale easily at a pace that matches customers business requirements, so they can be ready for whatever tomorrow brings.

- **Availability**

The advanced availability features of the Sun Fire 3800-6800 server family support the increasing availability requirements of today's businesses. Features such as full hardware redundancy and fault-isolated domains enable the systems to continue to be available when encountering failure of any single component. The hot-swappable components, in conjunction with the Dynamic Reconfiguration (DR) feature, provide the capability for online repairs. The Hot CPU Upgrade feature provides online capacity upgrades while the users' applications continue to be available.

- **Investment Protection**

The software architecture of the Sun Fire server family is based on the industry-leading Solaris Operating Environment and the SPARC™ V9 64-bit processor architecture. The Solaris Operating Environment is one of the only operating environment that enables compatibility between hardware and software versions on a release-to-release basis. Sun is one of the few companies that offers comprehensive hardware and software investment protection.

Unique to the Sun Fire 3800-6800 server family is the mixed-speed CPU feature¹: Customers can add newer, faster CPUs at the rated speed along with existing CPUs. This provision further protects their previous investments.

One related measure in protecting customers' investments is Sun's plan to continue shipment of Sun Enterprise 6500-3500 systems for 1 year after the introduction of the Sun Fire 3800-6800 systems. The UAP Plus Program provides customers with investment protection through trade-in value protection and technology refresh options.

Key Features and Benefits

Feature

- Full hardware redundancy - unique capability

Benefit

- Very high application availability in the event of hardware failure. If properly configured, the system can survive a failure that involves a single hardware component. Only Sun Fire servers offer redundant interconnects, clocks and System Controllers.

¹ Mixed Speed CPU will be available with UltraSPARC III Cu 900Mhz and above processors



Feature

- Hot CPU upgrades² - unique capability
- Fault-tolerant, redundant power supplies, fans and system clock - unique capability
- Fault-isolated Dynamic System Domains - unique capability
- Dynamic Reconfiguration (DR)
- UltraSPARC III processors
- Redundant, reconfigurable Sun Fireplane™ interconnect
- Common components
- Support for mixed speed CPUs⁴
- Enhanced Auto Recovery (System Controller Failover)
- SSH (Secure Shell)

Benefit

- Upgrade to faster speed CPUs while the Solaris Operating Environment and applications continue to be available.
- Continuous availability in the event of a component failure.
- The flexibility of running multiple, isolated Solaris Operating Environment instances and applications on one server.
- Hot swap resources in and out of a domain through Dynamic Reconfiguration without impacting availability.
- Ability to add or remove hardware components while applications and Solaris Operating Environment continue to run. DR can be used to repair, upgrade, or move system resources between Dynamic System Domains.
- Enhanced application performance with full binary compatibility for all applications running on existing Sun systems.
- Provides for lower latency, greater memory and I/O bandwidth. Enables Sun Fire servers to scale up to 24 UltraSPARC III processors.
- Investment protection and improved serviceability across Sun Fire 3800-6800 server family.
- Investment protection by avoiding the need to upgrade existing CPUs when additional CPU capacity is required.
- The System Controller (SC) Failover capability enables the main SC to failover to the spare SC automatically without operator intervention. The spare SC assumes the role of the main and takes over all SC responsibilities.
- The Sun Fire 3800-6800 server line enables implementation of a secure communication channel for network based administration of the system controllers. The midrange server SC offers SSH server functionality, and the SSH server responds to SSHv2 client requests.

² This feature is available on UltraSPARC III Cu 900 Mhz and above only. All CPUs on a CPU/Memory board must run at the same speed. UltraSPARC III 750 Mhz and 900 MHz may not be mixed in the same system.



Feature

- Scalability

Benefit

Like its predecessors, the Sun Fire servers is a family of large-scale, shared memory, symmetric multiprocessor (SMP) systems based on the UltraSPARC microprocessor.

The all-new Sun Fireplane interconnect system provides superior memory and I/O bandwidth, helping to ensure that Sun Fire servers deliver balanced and predictable performance under the most demanding loads. As resources are added to the systems, the Sun Fireplane interconnect system scales, providing high-speed, low latency access to CPU, memory and I/O devices to help ensure consistent high performance, even in the largest system configurations.

- Solaris 8 Operating Environment

- The industry leading UNIX software environment offers industry leading availability and scalability with support for over 12,000 applications, 64-bit support and unmatched features for network computing.

Availability Date

The Sun Fire 3800-6800 servers began shipping in May, 2001.

Target Markets and Users

The Sun Fire 3800-6800 server family is a general purpose server suitable for all industries and business needs. Sun will specifically target the following strategic markets - manufacturing, finance, telecommunications, government, health care, retail, design automation and oil and gas. Each server in the Sun Fire family has different characteristics and capabilities and can therefore be deployed in different ways, targeted to the needs of more specific markets and users.

Sun Fire 3800 Server Positioning

The Sun Fire 3800 server is Sun's entry-level midrange server. It offers unprecedented availability, performance and flexibility along with mainframe features such as Dynamic System Domains and full hardware redundancy, previously only available in the Sun's Enterprise 10000 server or in mainframe-class servers.

With up to eight CPU's, 64 GB memory and 12 CompactPCI (cPCI) slots for I/O, the Sun Fire 3800 server is ideal for customers who need application servers with tremendous processing power and network throughput, in addition to high reliability, availability and serviceability. The Sun Fire 3800 server is also ideal for price-sensitive customers who require mainframe features in a small footprint or who want the investment protection offered by an expandable family of servers.

The Sun Fire 3800 server is designed to be rack-mounted in combination with storage or with other Sun servers, including other Sun Fire 3800 servers. This server is a logical choice as an upgrade for customers currently using a Sun Enterprise 420R, 450, 3500, or 4500 server and looking for high availability and better resource management and performance.



Sun Fire 4800 Server Positioning

The Sun Fire 4800 server is a versatile server with exceptional value and scalability for companies requiring an affordable, highly available business server with tremendous computing power, the ability to scale system performance and capacity as their needs grow and the versatility of Dynamic System Domains.

The Sun Fire 4800 server is ideal for customers who need an enterprise-class application server with high reliability, availability and serviceability. Typical Sun Fire 4800 server customers may already own a Sun Enterprise 4500 or 5500 system.

The Sun Fire 4800 server is recommended for running mission-critical applications that support thousands of users or other large workloads. These applications include large departmental databases, customer-management applications, decision-support applications or HPC (High Performance Computing) workloads.

Sun Fire 4810 Server Positioning

Designed for medium-to-large ISPs, ASPs, Internet data centers, telecommunications, military and industrial applications, the Sun Fire 4810 server offers the same features and capabilities as the Sun Fire 4800 server, but with 100% front access to all components and a shallower rack depth. Sun Fire 4810 servers fit in industry standard 19" with 30" depth racks.

Sun Fire 6800 Server Positioning

The Sun Fire 6800 server is ideal for customers who wish to add the flexibility and robustness offered by the Sun Enterprise 10000 server to their large enterprise or Internet data center applications. The customer receives the benefits of improved scalability and compatibility, allowing applications to grow to tens of thousands of users, to multiple terabytes (TB) of data storage and up to four Dynamic System Domains. Typical Sun Fire 6800 server customers may already own Sun Enterprise 6000, 6500, or 10000 systems.

The Sun Fire 6800 server offers twice the CPU and memory expandability of the Sun Fire 4800 and 4810 servers. The Sun Fire 6800 server is an ideal choice when the Sun Fire 4800 server does not offer enough expandability or when customers require up to four Dynamic System Domains.

The Sun Fire 6800 server is recommended over the Sun Enterprise 10000 server when:

- There is a requirement for UltraSPARC III processor technology, including the power of faster CPUs
- Applications run on Solaris 8 Operating Environment
- The customer wants the latest technology
- A maximum of four domains and 24 or fewer CPUs are required



Positioning Summary

Market/User	System Choice(s)	Key Features to Highlight
Service provider or mission-critical IT provider requiring maximum availability at lowest cost	Sun Fire 3800	<ul style="list-style-type: none"> • Full hardware redundancy • Entry level systems with upgrade path • Small footprint and rack-mountable • 1 or 2 Dynamic System Domains
Customers who want mainframe-like features (flexibility and management tools) and high performance on a midrange server	Sun Fire 4800 Sun Fire 6800	<ul style="list-style-type: none"> • Centralizes management to cut administration costs • Multiple Dynamic System Domains • Leading open-system performance and I/O throughput • Enhanced Auto Recovery (SC Failover)
Users with mission-critical applications, e.g. E-commerce, 24 x 7 core business systems	Sun Fire 3800-6800 Sun Enterprise 10000	<ul style="list-style-type: none"> • High Reliability, Availability and Serviceability for maximum availability • Full hardware redundancy and Hot CPU upgrades • Enhanced Auto Recovery (SC Failover)
Internet Service Providers, Application Service Providers, military and industrial customers needing front-access and high availability	Sun Fire 4810	<ul style="list-style-type: none"> • All front-access Field-Replaceable Units (FRUs) • Full hardware redundancy • Shallow rack footprint (24") - fits in 30" deep rack (19" wide x 30" deep) • Hot-swappable components
Customers looking to lower total-cost-of-ownership through server consolidation	Sun Fire 3800-6800 Sun Enterprise 10000	<ul style="list-style-type: none"> • Multiple Dynamic System Domains • Centralized management through Sun Management Center 3.0 • Full hardware redundancy to meet strict Service Level Agreements
Technical and high-performance computing customers requiring market-leading performance and attractive prices	Sun Fire 3800-6800	<ul style="list-style-type: none"> • High-performance UltraSPARC III CPUs • High-bandwidth Sun Fireplane interconnect system • Low-latency Sun Fireplane interconnect system • Large physical memory • High I/O throughput, support for 66MHz, 64-bit PCI adapters
High performance computing customers with distributed application requirements	Sun Fire 3800-6800	<ul style="list-style-type: none"> • High-Performance UltraSPARC III CPUs • High I/O throughput • Large physical memory space • Support for high performance interconnects
Commercial sites that require high availability, more than four domains, proven scalability, or can not move to Solaris 8 Operating Environment. These customers may be late adopters.	Sun Enterprise 10000	<ul style="list-style-type: none"> • Up to 16 Dynamic System Domains • Interdomain networking • Automatic Dynamic Reconfiguration • Scalability up to 64 CPUs • Can use Solaris 2.6 Operating Environment or Solaris 7 Operating Environment



Selling Highlights

Market Value Proposition

This new generation of midframe servers from Sun offers features for the next generation of network computing. Sun and its customers have identified that features such as full hardware redundancy, Hot CPU Upgrades, Dynamic System Domains and Dynamic Reconfiguration are needed for mission-critical solutions. The Sun Fire 3800-6800 server family offers these mainframe-class features.

This product line also continues Sun's focus on protecting our customers' investments in technology, utilizing the same Solaris Operating Environment, 64-bit superscalar UltraSPARC processor architecture, Sun Management Center console and many other advanced features. The advantages of evolving these features is that it retains the common application programming interface (API), application binary interface, operational practices and management practices to ensure that current customers leverage their existing knowledge of and investment in Sun while taking advantage of the latest technology required to expand their business.

The Sun Fire 3800-6800 servers offer some of the most advanced platforms for technical computing based on the UltraSPARC III CPU and Sun Fireplane interconnect system.

Applications

Based on years of experience in scaling and supporting the SMP architecture, the Sun Fire 3800-6800 servers are designed for deploying all of today's mission-critical applications, from ERP to eCommerce to HPC. The scalability and flexibility are excellent features for applications that are deployed in a multi-tier environment, have indeterminate potential growth, or need extra processing capacity to satisfy monthly, weekly, or other regular excess transaction load.

The availability features of the product line are ideally suited to applications that require very high availability. Support for Sun Cluster 3.0 software is expected to be available later in 2001, which will provide for even greater availability requirements. However when all components are configured redundantly, a single Sun Fire 3800-6800 system offers outstanding availability often exceeding the requirements of our customers.

Compatibility

The Sun Fire 3800-6800 servers run the Solaris 8 Operating Environment on which thousands of applications are available. The UltraSPARC III 750 systems require Solaris 8 4/01 or later OS. The UltraSPARC III Cu 900 systems require Solaris 8 10/01 or later OS.

Compatibility with the Solaris Operating Environment brings with it the ability to run any existing Solaris application developed for the SPARC 32-bit or 64-bit processor architecture. The same application programming interfaces (APIs) and the same application binary interface (ABI) in previous versions of Solaris Operating Environment exist in Solaris 8 Operating Environment running on the Sun Fire 3800-6800 servers. Customers can also take advantage of the Solaris Application Guarantee program to enable a smooth migration from previous versions of Solaris or earlier versions of the SPARC processor architecture. *See the Software section later in this document for details.*



All of the following Sun StorEdge™ products are supported on the Sun Fire 3800-6800 servers:

- Sun StorEdge UniPack for F4800 Deskside
- Sun StorEdge MultiPack for F4800 Deskside
- Sun StorEdge D240 Media Tray
- Sun StorEdge D1000 Array
- Sun StorEdge A1000 Array
- Sun StorEdge A5000 and A5200 Array
- Sun StorEdge A3000 and A3500 Array
- Sun StorEdge T3 Array
- Sun StorEdge 9900

Sun StorEdge A3500 array has reached end-of-life status and is no longer available for ordering.

A full range of PCI and CompactPCI I/O adapters are supported with the Sun Fire 3800-4800 servers. They are listed in the Ordering Information section.

PCI and CompactPCI(cPCI) Adapters

I/O connectivity is available for the Sun Fire 3800-6800 via a comprehensive range of PCI and cPCI adapters. The Sun Fire 3800 server supports only cPCI adapters. The 4800 - 6800 support both PCI and cPCI adapters. They are listed in the Ordering Information section.



Reliability, Availability and Serviceability

Reliability, availability and serviceability (RAS) are three aspects of a system's design that contribute to continuous operation and the reduction of system downtime.

The Sun Fire 3800-6800 server family brings together the finest aspects of Highly Available engineering, extending the feature set offered by the Sun Enterprise 10000 server.

Reliability

Reliability refers to a system's ability to operate continuously without failures and to maintain data integrity. The Sun Fire 3800-6800 server family includes the following features that enhance reliability, many of which will be detailed in the System Architecture section of this document.

- Extensive environmental monitoring, which helps ensure that the system is not allowed to operate outside its specifications.
- End-to-end ECC error detection and correction on all data paths within the system to detect and correct single-bit errors. This feature maintains data integrity without the need to halt system operation in the case of single-bit errors. ECC will also detect double bit errors and log them to the System Controller.
- Enhanced ECC to protect against naturally occurring radiation. Up to four hard failures on adjacent bits from naturally occurring radiation can be corrected. (See System Architecture section for details).
- Parity detection on all address path segments for improved data integrity. This is in addition to ECC on all data paths.
- Memory scrubbing to enable the memory contents to be frequently refreshed. This feature reduces the chance of single-bit errors.
- Extensive component and interconnect Power On Systems Test (POST) and other tests performed prior to system boot to help ensure that faulty components are excluded from the system configuration.
- Passive centerplane (Sun Fire 4800-6800 servers) helps provide high mean time between failures (MTBF). Because the centerplane is passive, (that is, it contains no chips on board), it is less likely to fail.
- Boards are locked in place and I/O connectors are secure, which prevents system errors due to loose connections. Boards also have different form factors, minimizing the chance of inserting the wrong board in a centerplane slot.

Availability

System availability refers to the percentage of time that a system remains accessible and usable. Sun Fire 3800-6800 servers offer many advanced availability features. Some of these features are detailed in the System Architecture section of this document.

- Full Hardware Redundancy
 - Redundant CPUs



- memory banks
- CPU/memory boards
- I/O assemblies
- I/O adapters (if configured)
- System Controllers
- System clock with automatic failover
- Fireplane interconnect system switches
- AC power sources, facilitated by the Redundant Transfer Switch
- Power supplies and intelligent power switching mechanism to failover to remaining power modules
- Dual power grids option
- Fully Fault-Isolated Environment
 - Fireplane Interconnect can be configured into two segments. Each segment is isolated from failures in the other crossbar segment. This is in addition to the fault isolation at domain level, as described below.
 - Each crossbar segment can host one or two Dynamic System Domains. Each domain is isolated from failures that originate in other domains, including both hardware and software faults from:
 - Application or Solaris Operating Environment failures
 - Errors originating in CPUs, memory modules, CPU/Memory boards, I/O assemblies, PCI cards or cPCI cards.

The following diagram, at both domain or crossbar segment level, lists potential failures and the isolated environments that normally can be achieved on a carefully configured Sun Fire 6800 system by following the practices outlined in the rest of this Just the Facts. Similar fault isolated environments can also be configured on a Sun Fire 3800, 4800 or 4810 systems.




Figure 1 - Fully fault-isolated environments in a Sun Fire 6800 system

- Dynamic Reconfiguration



- Hot Pluggable and Hot Swappable Components

Hot Plug is the ability to add, remove or replace a component while the system is running. Hot Swap is the ability to enable the use of that component while the system is running and without a reboot. Hot Swap of a CPU/Memory board or an I/O Assembly requires the support of Dynamic Reconfiguration.

Many components on the Sun Fire 3800-6800 system can be hot-pluggable or hot-swappable as shown in the table below:

	Hot Swappable	Hot Pluggable
Fan	Yes	Yes
CPU/Memory Board	Yes	Yes
Fireplane Switch	-	Yes
Power Supplies	Yes	Yes
I/O Assemblies	Yes	Yes
cPCI I/O Adapters	Yes	Yes

Figure 2 - Hot swappable or hot pluggable components

- Mixed speed CPUs³

By separating the clock signals and device arbitration on this family of servers, the clock signal for CPUs within the system are generated by multiplying the 150 MHz signal from the System Controller on a board-by-board basis. The benefit to the customer is that they have the flexibility to add processors of different speeds online as long as they do not mix processors of differing speeds within one system board.

- The diagram below illustrates how processors of different speeds can be intermixed and deployed on a Sun Fire 3800-6800 server.




Figure 3 - Implementing mixed speed CPUs

³ This feature is available on UltraSPARC III Cu 900 Mhz and above only. UltraSPARC III 750 Mhz and 900 MHz may not be mixed in the same system. .



- There are some restrictions to observe when using mixed-speed CPUs:
 - Only UltraSPARC III Cu 900MHz and above processors support mixed speeds in a domain/system.
 - In a system running UltraSPARC III 750 MHz processors, different speed processors are not supported.
- Hot CPU Upgrades. This feature facilitates the online upgrading of the CPUs while the rest of the system continues to be available. This feature is supported partially by the Dynamic Reconfiguration feature and requires that all CPUs involved be UltraSPARC III Cu 900MHz or above.

The diagram below illustrates the three-step upgrade process to upgrade all CPUs on two CPU/Memory boards from a speed greater than 750 MHz to even faster speed. During this online upgrade process, the application remains available.




Figure 4 - Hot CPU Upgrade

- IP multi-pathing to provide automatic network path failover

When two I/O paths can be defined to connect a network to a Sun Fire server through the presence of two I/O adapter cards, the IP multi-pathing (IPMP) feature of Solaris 8 Operating Environment can be employed to map out a dual-path network connection with automatic network failover capability.

The automatic network failover capability helps ensure that in the event of a failure of a single I/O card, the server will be able to continue communicating over the network.
- Parallel device probing

In the event of a system reboot, if your configuration comprises more than one CPU, then the CPUs are used to probe for devices concurrently, significantly reducing the time it takes for the operating environment to identify and configure attached devices. This feature helps to shorten down time.
- What would be the impact from various failures? The table below summarizes the impacts into four categories and lists the major causes for each category. There are very few failures that would bring down a complete system, a crossbar segment or all domains.






Figure 5 - Normally expected recoverability from a single failure

Serviceability

Serviceability relates to the time it takes to restore a system to service following a system failure.

Some of the serviceability features of the Sun Fire 3800-6800 server family include:

- Hot swappable CPU/memory boards, I/O assemblies, CompactPCI adapters, Redundant Transfer Switch modules, power supplies and cooling units, which provide for on-line replacement
- Dynamic Reconfiguration and IP multi-pathing, which allow for CPU/memory boards, I/O assemblies or CompactPCI adapters to be taken off-line for service without interruption to Solaris Operating Environment or application
- Virtual Key Switch and Virtual Console features on System Controller, which allows the system to be powered on or off and rebooted remotely; with optional diagnostics
- Mechanical Serviceability
 - Connectors keyed so that boards cannot be inserted incorrectly
 - No jumpers are required for configuration of the Sun Fire 3800-6800 servers
 - All FRUs (Field-Replaceable Units) are designed for quick and easy replacement, with a #2 Philips screwdriver.
 - Safe electrical voltages (48 VDC and below) used throughout all system boards
 - All FRUs identified with LED service indicators for positive indication of whether a FRU can be removed
 - Electronic serial numbers on all FRUs enabling improved component tracking.
- Front access of all components (on Sun Fire 4810 server only)
- Physical view of system from Sun Management Center software
Also aids location of failed components, helping speed servicing and avoid mistakes



- Sun Validation Test Suite software (SunVTS™) allows Administrators to perform system diagnostics.
- Sun Remote Services provides remote monitoring, diagnostics and service dispatching.
- A new feature in the Sun Fire 3800-6800 server family, the Serial EEPROM is a part of all system FRUs. It consists of a small amount of read-only memory containing information to identify the FRU. This feature enhances serviceability.

Auto Diagnosis and Recovery Features

- Auto Diagnosis and Recovery provides new levels of reliability, availability, and serviceability for the Sun Fire 3800-6800 servers to help customers minimize both planned and unplanned downtime and significantly reduce business interruptions through faster time to service.
- Auto Diagnosis with Messaging: Auto Diagnosis detects the first instance of an error and records the Component Health Status (CHS) of the faulty or suspect components. Auto Diagnosis provides easy-to-understand error reports that identify faulty or suspect components to a single-Field Replaceable (FRU) or multi-FRUs. The server's Power On Self Test (POST) then uses the CHS information to deconfigure the faulty component out of the system. Diagnostic information is reported through the platform and domain console event message or the loghost output.
- Component Health Status: Deconfigures faulty components (CPU/memory boards, CPUs, L2_SRAM modules, DIMMs, I/O boards, and Fireplane switches) after a hard failure. CHS identifies ("blacklists") faulty components due to a fault or suspected fault. The reconfigured system is automatically rebooted. The fault is isolated from the new configuration, significantly reducing the chances of a similar failure.
- CHS and diagnostic information are persistently stored on a component, which prevents the reoccurrence of a fault even if the component is moved to a different location. Preventing the reoccurrence of a fault improves the availability characteristic of Sun Fire 3800-6800 servers. As the diagnosis information is contained inside the component, service and repair of these systems becomes faster and more efficient. Valuable service information is collected and reported to help Sun detect hardware errors and improve product quality.
- Auto Restoration: Working in conjunction with the Auto Diagnosis engine and Component Health Status, Auto Restoration automatically restores the system with the fault isolated in the event of a fatal error as a result of a faulty component. If a FRU or component is disabled because of its CHS, immediate replacement is not necessary because the domain is restored with the fault isolated.
- Automatic Recovery of Hung Domains: System detects and recovers from hung domains. A domain hang occurs when a domain does not respond to user commands or is not reachable via the network.
- Recovery from Repeated Domain Panics: Domain panics can be caused by both software and hardware. To prevent hardware faults from causing panic reboot loops, the SC Firmware has been enhanced to run POST diagnostics at more detailed diagnostic levels. On the first panic the domain reboots and writes a core file. The core file can be used to analyze the problem. However, if further panics occur within a short time period, it is desirable to run POST automatically at a higher level as part of domain restoration. POST diagnostics verify the status of the hardware and could identify and isolate faulty components (if any). After identifying faulty components, POST updates their CHS appropriately.
- Solaris OE Enhancements: Kernel Updates for Solaris 8 OE and Solaris 9 OE on UltraSPARC III systems enhance the correctable error (CE) DIMM and L2_SRAM module handling. Multiple



Correctable Errors on accessing a DIMM or L2_SRAM module indicate a higher probability of experiencing an uncorrectable error (UE). To prevent a fatal UE, memory pages are retired and CPUs are automatically off-lined. The availability of domain increases, because the Solaris OE does not access pages or L2_SRAM modules that have an increased failure probability.

UltraSPARC III Processor

The UltraSPARC III processor is part of a third generation of UltraSPARC processor-based products. In addition to using a new processor technology, the UltraSPARC III processor provides a higher clock frequency, reduced on-chip latencies, support for greater amounts of level-one and level-two cache, an integrated external memory controller and address snooping to provide cache-coherency in multi-CPU configurations. Other new features include support for increased reliability through enhanced error detection and correction. This new generation of products provides software compatibility with existing UltraSPARC processor-based systems.

UltraSPARC III processors offer the following features:

- Four-way associative, on-chip 64-KB data and 32-KB instruction cache, with up to 8 MB of ECC-protected external level 2 cache through an integrated memory controller
- Integrated DRAM controller with support for up to 8 GB of memory, able to transfer data up to 2.4 GB/second
- SPARC Version 9 processor architecture compliant
- Binary compatible with all existing SPARC processor applications
- Enhanced OpenVIS™ Instruction Set to support advanced multimedia capabilities
- 64-bit address pointers that provide transparent compatibility with 32-bit addressing
- 64-KB, 4-way, set-associative data cache
- 32-KB, 4-way, set-associative instruction cache steers up to four instructions per cycle to six execution pipes
- Integrated second-level cache controller supports 8-MB caches. Sustained throughput of one load per cycle and 3.2 GB/second processor-cache bandwidth

UltraSPARC III processors are binary compatible with all Sun SPARC processor technology based systems, continuing a characteristic of previous generation products. The UltraSPARC III processor provides the highest integer and floating point performance in the SPARC processor family and is able to address the most computationally demanding applications.



The following table compares UltraSPARC III processor with previous generations of Sun's 64-bit CPUs:

	UltraSPARC I	UltraSPARC II	UltraSPARC III	UltraSPARC III Cu
Transistors	5.2 million	5.4 million	29 million (12m cache, 4m logic)	29 million (12m cache, 4m logic)
Processor Technology	0.5 μ , 4 metal layers	0.35 μ , 5 metal layers	0.18 μ , 6 metal layers	0.15 μ , 6 metal layers
Performance	8.3/11.4 SPECint/fp95 (200 MHz, 2 MB e-cache)	15/22 SPECint/fp95 (400 MHz, 4 MB e-cache)	395/421 SPECint/fp2000 Peak (750 MHz, 8 MB e-cache)	467/482* SPECint/fp2000 Peak (900 MHz, 8 MB e-cache) *Estimate
Data Cache Size, Associativity	16 KB, Direct	16 KB, Direct	64 KB, 4-way	64 KB, 4-way
Instruction Cache Size, Associativity	16 KB, 2-way	16 KB, 2-way	32 KB, 4-way	32 KB, 4-way

Cache Coherency

All of Sun's systems today use processor caches to speed access to data and reduce traffic through the system interconnect. Data that has been recently used, or whose impending use is anticipated, is retrieved and kept local to the processor(s) that will need it. In a shared memory system, making use of multiple processors, the task of maintaining consistency in the system caches requires assistance from the system interconnect.

All of Sun's multiprocessor systems incorporate a technique known as *snooping* to keep the caches up-to-date and to maintain coherency. In a snooping system, the addresses of all transactions on the system interconnect are monitored for cache update transaction data.

An enhancement of the UltraSPARC III processor is that cache snooping is now integrated directly in the CPU, whereas previous SPARC processors used external cache controllers. This feature integrates the CPU more tightly with the interconnect and is a critical feature for midframe UltraSPARC III processor systems because it enables cache coherency to scale linearly as CPUs are added to the system.

Benefits

Customers benefit from the UltraSPARC III processor in the following ways:

Reliability and Availability

- Extensive error correction and isolation features including a new *UpTime* bus
- ECC (Error Correction) extended to L2 Cache for increased reliability



Scalability

- Near-linear scalability in multiprocessing systems to support hundreds of processors
- On-chip memory controller enables the processor to map and manage very large, system-wide memory images.
- Integrated memory controller and cache snooping provides for an architecture that can scale to more than 100+ processors and to more than one TeraByte of RAM.

Investment Protection

- Binary compatibility for applications written for previous versions of UltraSPARC processors and Solaris Operating Environment technology-based systems with no recompiling required
- Synergistic design process involving UltraSPARC processor, Solaris Operating Environment and Sun systems products teams
- Sun long term commitment to UltraSPARC processor architecture for future generations of products

Performance

- Best-in-class performance metrics when compared to competing 64-bit server-class processors
- "Design for Bandwidth" engineering philosophy emphasizing low latency (integrated L2 cache), maximum data throughput and support for massive cache memories
- Initial versions at 750 MHz clock speeds with a published roadmap of versions exceeding 1.5 GHz.



System Architecture

Sun Fire 3800-6800 Server Key Facts

The following is a configuration summary for the Sun Fire 3800-6800 servers.

	Sun Fire 3800	Sun Fire 4800/4810	Sun Fire 6800
CPUs	2 to 8	2 to 12	2 to 24
Memory (max. GB)	64	96	192
CPU/Memory Boards	1-2	1-3	1-6
8-slot PCI I/O assemblies	-	2	4
6-slot cPCI I/O assemblies	2	-	-
4-slot cPCI I/O assemblies⁴	-	2	4
Max. number of cPCI slots	12	8	16
Max. number of PCI slots	-	16	32
Fireplane Interconnect Boards	in centerplane	2	4
Dynamic System Domains	1-2	1-2	1-4
Number of Power Cords	3	3	2 or 4 ⁵

Figure 6 - Outline of Sun Fire 3800-6800 servers

Each of the systems in the Sun Fire 3800-6800 server family includes the following features:

- Two-level hierarchical Sun Fireplane interconnect systems, providing the high throughput required by UltraSPARC processors, high performance memory and I/O subsystems
- Hot-swappable CPU/memory boards, I/O assemblies, power supplies and cooling modules
- High performance 66 MHz PCI and 33Mhz PCI/cPCI I/O
- Extensive environmental sensing, monitored by software, including Sun Management Center software.
- Dynamic System Domains
- Support for these features (and others) from Solaris 8 4/01 Operating Environment or later and Sun Management Center 3.0 software or later

⁴ This feature is scheduled for delivery in Q4CY01

⁵ Sun Fire 6800 server requires 4 power cords when dual power grids are installed.



Technical Fact Summary

- **High-performance Sun Fireplane interconnect**
 - The Fireplane interconnect system provides improved performance, improved scalability, lower latency and three times the system bandwidth of previous midrange servers.
 - Provides up to two Dynamic System Domains on Sun Fire 3800, 4800 and 4810 servers; up to four Dynamic System Domains on the Sun Fire 6800 server.
- **High-performance memory subsystem**
 - Up to 192 GB (on the Sun Fire 6800 server); A maximum of 8 GB memory per configured processor.
 - Supports 1 GB, 2 GB, or 4 GB per memory bank and up to two banks per processor.
 - Up to 16-way memory interleaving within each CPU/memory board.
- **High-performance I/O**
 - PCI I/O bus offering four independent PCI buses per I/O assembly and 66-MHz PCI support
 - cPCI for higher availability and serviceability
- **Hot-swappable CPU/Memory boards, I/O Assemblies and cPCI cards**
 - Combined with Dynamic Reconfiguration, IP Multipathing (IPMP) and Multiplex I/O (MPxIO) provides for automatic failovers, on-line upgrades or repairs.
- **Multiple Dynamic System Domains**
 - Now available in midframe systems
 - Run multiple Solaris Operating Environment instances securely with fault isolation.
- **Power supplies**
 - N+1 fully redundant power supplies
 - Power supplies are hot swappable
 - Dual independent power sources, switched via logic-controlled Redundant Transfer Switch.
 - The Sun Fire 6800 server additionally offers dual power grids feature.
- **System enclosure**
 - The Sun Fire 3800, 4800 and 4810 systems can be mounted in industry standard 19" racks; the Sun Fire 4800 server is also available as a desktide tower.
 - The Sun Fire 4800 server is configurable in a Sun Fire cabinet or in a standard 19-inch wide, 36-inch deep commercial rack. The Sun Fire 4810 server can be mounted in a 19-inch wide, 30-inch deep rack. The Sun Fire 3800 server can be mounted in a 40-inch deep, 19-inch wide, commercial rack.
 - The Sun StorEdge D240 media tray is available with the Sun Fire servers for compact, hot pluggable boot disks, DVD/CDs and tape drives, and can be mounted in the same rack as the server.



Two-Level Sun Fireplane Interconnect System

The Sun Fireplane interconnect system is a two-level interconnect that provides very low latency with high address and data bandwidth to enable 24-way scalability.

The Sun Fireplane interconnect system operates at 150 MHz. Address snoop bandwidth is 9.6 GB/s and the address portion of the interconnect provides for 150 million snoops per second. Memory latency is very low - between 180 and 240 nanoseconds. The higher number reflects accessing data on a different CPU/Memory board by a CPU.

Hierarchical Interconnect Structure

In the Sun Fire 3800-6800 server family, the Sun Fireplane level 1 (L1) interconnect exists within the CPU/Memory board and the I/O assembly. The L1 interface for both the address and data interconnects is connected to a level 2 (L2) interface. The L2 interface is responsible for directing incoming transactions back to the correct L1 target interconnect, which is responsible for directing the signal to the particular target device.




Figure 7 - Logical view of Sun Fire 3800-6800 system

The diagram above is a high-level logical view of the two-level interconnect and some of the major sub-components used in the Sun Fire 3800-6800 server family. The L2 Interconnect (Sun Fireplane interconnect system) is discussed later. Note that the Sun Fireplane switches are built into the centerplane of the Sun Fire 3800 server.



Redundancy

The centerplane in the Sun Fire 3800-6800 server is designed as a highly-available redundant crossbar. The Sun Fire 3800-4810 centerplane has two Sun Fireplane switches. The Sun Fire 6800 server has four Sun Fireplane switches. Although a single Sun Fireplane switch (or a pair of these switches in the Sun Fire 6800 server) can support the full set of system resources, the number of Sun Fireplane switches are doubled on each of the Sun Fire 3800-6800 server to provide redundancy and improved availability.

The centerplane provides a switched crossbar interconnection with separate address and data buses to CPU/memory or I/O Controller boards. The centerplane data path is fully ECC protected and the address is parity protected throughout.

The Sun Fireplane interconnect system is unique in the industry. The Sun Fireplane Interconnect boards can be configured either as a single crossbar segment or as two independent crossbar segments. When configured as two crossbar segments the Sun Fireplane provides for fault isolation between segments, thus increasing availability.

The failure of a Sun Fireplane Interconnect in a segment affects only that segment and the domains associated with that segment. The other segments and domains continue unaffected. This provides a level of availability typically not matched by other vendors. It is also an improvement over the Sun Enterprise 10000 server where a centerplane error would affect all domains in the system.

Sun Fireplane Interconnect Boards (Sun Fire 4800-6800 servers)

The Sun Fire 4800, 4810 and 6800 servers use Sun Fireplane Interconnect boards. The L2 interface (Sun Fireplane interface) is implemented by a pair of boards. In the Sun Fire 4800-4810 servers, each board contains logic for Sun Fireplane interfaces for both address and data interconnects. Transactions are bit-sliced across the two boards. The Sun Fireplane Interconnect boards are hot pluggable and fault resilient. In the event of a failure that is not recoverable, one of the Sun Fireplane boards will continue to operate following a manual system reboot, although at half the bandwidth of the full system.

In the Sun Fire 6800 server, a second pair of Sun Fireplane Interconnect boards are used. These four boards are treated as two groups of two boards each. Within each group, transactions are bit-sliced across two Sun Fireplane Interconnect boards. In the case of a failure, the group that contains the failing board will fail. However, the other group can survive the failure and continue to be available, if properly configured.

The Sun Fire 3800 server uses an active centerplane approach to minimize the physical size of the system. In this case, the centerplane itself holds the Sun Fireplane interface logic. Even though the Sun Fire 3800 switch is implemented on the centerplane, the interconnect acts as if it is two separate boards. If one part of the interconnect fails, the system will continue to run after a manual reboot.

High-Performance Design

To provide the bandwidth required by today's high performance computer systems and to avoid the shortcomings of interconnect designs that are based on bus technology that can become overloaded at high workloads where many devices are simultaneously accessing the interconnect, the Sun Fireplane crossbar interconnect is designed as a packet-switched interconnect and is built to provide very high sustained bandwidth.

The interconnect operates at 150 MHz and supports 1 snoop per cycle. This equals a coherency bandwidth of 150 million address snoops per second, which is three times the current address of the Enterprise midrange servers. The cache line size (and hence the maximum transfer size on the Sun Fireplane data crossbar) is 64 bytes. This equates to an available data bandwidth of 9.6 GB/s⁶.

⁶ Where 1 GB = 1 x 10⁹ bytes



The design of the interconnect leads to very high rates of utilization of the available bandwidth, so the performance of workloads that move large volumes of data are not restricted by the system interconnect.

The Sun Fireplane interconnect system has separate address and data paths. The arbitration occurs in parallel with separate address broadcasts and data transfers so it does not consume the available bandwidth. The address and data paths of the interconnect are separate, so an address broadcast does not consume the bandwidth available for transferring data. This results in an average latency of 180 nanoseconds for accessing memory on the same CPU/Memory board, or 240 nanoseconds for accessing memory on different board by any CPU in the system. These are very low interconnect latency numbers.

	XDBus	Gigaplane™	Sun Fireplane
ample system	SPARCcenter 2000	Sun Enterprise 3500-6500	Sun Fire 3800-6800
Architecture	Interleaved, dual packet-switched buses	Packet-switched buses	Packet-switched crossbar
Operating speed	50 MHz	100 MHz ⁷	150 MHz
Snoop bandwidth	N/A	50 million addresses/sec	150 million addresses/sec
Cache line data width	8 bytes	32 bytes	64 bytes
Sustained bandwidth	625 MB/sec	3.2 GB/sec	9.6 GB/sec
Hot plug system boards	No	Yes	Yes
Dynamic System Domains	No	No	Yes

Figure 8 - Bus verses Centerplane

The *packet switched* nature of the Sun Fireplane interconnect system means that requests for data do not have to be followed immediately by the data transfer in order to satisfy the request. Also, data transfers can occur *out of step*, with the Sun Fireplane interconnect system supporting a pipeline of up to fifteen outstanding requests. This means the gap between any address request and the data transfer to fulfill the request can be used to complete another unrelated request instead of incurring idle cycles on the bus.

This flexibility allows systems designed around the Sun Fireplane interconnect system to utilize virtually all of the 9.6 GB/sec available bandwidth, while keeping latencies on individual transfers low - providing an excellent platform for memory-intensive computation or high-volume I/O.

Performance Characterization:

System Bandwidth - the Sun Way or the Aggregate Way?

In describing the system bandwidth, Sun believes in quoting sustained performance numbers. However, some of our competitors have been quoting aggregate numbers. That is, they add each interface bandwidth together for a total system bandwidth. This is a bit like saying a car can travel at 400 miles per hour because each of the tires revolves at 100 miles per hour. We can play that game too.

The following table highlights the various bandwidth claims that can be made:

⁷ 84 MHz maximum for the Sun Enterprise 6500



	Sun Fire 3800	Sun Fire 4800/4810	Sun Fire 6800
Sustained			
Each CPU/memory board	4.8 GB/s	4.8 GB/s	4.8 GB/s
Number of CPU/memory boards	2	3	6
Each I/O assembly	1.2 GB/s	1.2 GB/s	1.2 GB/s
Number of I/O assemblies	2	2	4
Sustained system bandwidth	9.6 GB/s	9.6 GB/s	9.6 GB/s
Aggregate			
CPU/Memory bandwidth	9.6 x 2 = 19.2 GB/s	9.6 x 3 = 28.8 GB/s	9.6 x 6 = 57.6 GB/s
I/O bandwidth	2.4 x 2 = 4.8 GB/s	2.4 x 2 = 4.8 GB/s	2.4 x 4 = 9.6 GB/s
Aggregated system bandwidth	9.6 x 2 + 2.4 x 2 = 24.0 GB/s	9.6 x 3 + 2.4 x 2 = 33.6 GB/s	9.6 x 6 + 2.4 x 4 = 67.2 GB/s

Figure 9 - Various Bandwidth Characterization

CPU/Memory Board

The fundamental building block for all Sun Fire 3800-6800 systems is the CPU/Memory board. The Sun Fire 3800-6800 CPU/Memory board holds processors and memory for the system. Each board can contain two or four UltraSPARC III processors and up to 32 GB of memory (32GB is for a board with 4 CPUs, 16GB for a board with 2 CPUs). The UltraSPARC III processor, by virtue of its integrated memory controller, is responsible for memory management of local memory and for running diagnostics on the CPU/Memory board.




Figure 10 - Sun Fire 3800-6800 CPU/memory board (processor/e-cache cover has been removed)

Memory banks are interleaved within each board to improve performance. Due to the implementation of Dynamic Reconfiguration, memory interleaving between boards in the system is not supported.

The same CPU/memory board is used throughout the Sun Fire 3800-6800 server family, so boards initially deployed in one system may be moved to another as computing needs change. Processors and



memory on a board can be upgraded. Upgrades from 2 CPU boards to 4 CPU boards are available (*See upgrade section*).

The memory is protected with ECC to correct all single-bit errors and to detect all double-bit errors. Furthermore the Enhanced ECC feature is implemented to protect against naturally occurring radiation. Those multi-bit errors, induced by naturally occurring radiation, usually occur on adjacent bits and would result in a hard failure/outage with normal ECC implementation. In Sun's unique implementation, we assign error checking for every four adjacent bits to four different sets of ECC circuits. Each of the ECC circuits can detect a single-bit error. This means that failures on up to 4 adjacent bits can be corrected since each error is a single-bit failure within the assigned ECC circuits.




Figure 11 - Enhanced ECC to protect against naturally occurring radiation. This feature can correct up to 4 single-bit failures from a single hit from naturally occurring radiation

System I/O - High Performance PCI and cPCI Technology

Sun continues its support of open standards in the area of I/O connectivity. The Sun Fire 4800-6800 servers support the widely used PCI standard: The standard PCI form factor will be used initially and the CompactPCI (cPCI) form factor will be available in Q4CY01. The Sun Fire 3800 supports cPCI exclusively. Sun has continued the use of dedicated I/O assemblies, separate from the CPU and memory boards. A new feature is no slot tradeoffs, which means that the customer does not have to choose between I/O or CPU/Memory resources. It is possible to configure the Sun Fire servers with maximum CPUs/Mem and I/O.

Each type of I/O connectivity - PCI and cPCI - uses an I/O assembly specific to that type of I/O card. Each assembly contains slots for only one type of I/O card; however, different types of I/O assemblies can be installed in the same system chassis. There are no built-in connectors on any of the I/O assemblies; any I/O connections must be made via I/O cards installed in I/O assemblies.

Each I/O assembly is capable of 2.4 GB/s throughput, exceeding the capacity of current PCI adapters and assuring capacity to support higher-speed, future generations of PCI I/O adapter cards.

8-slot PCI I/O Assembly (for Sun Fire 4800, 4810 and 6800 servers)

The 8-slot PCI I/O assembly is designed to support both large capacity and high bandwidth. Two independent PCI controllers are provided on each PCI I/O assembly. Each of these controllers has two



separate PCI buses. Within each PCI controller, one bus supports three 33-MHz slots operating at either 64 bits or 32 bits and the other supports a single extended 66 MHz PCI slot, also operating at either 64 bits or 32 bits. This configuration provides a total of eight PCI slots per I/O assembly - six 33 MHz slots and two 66 MHz slots.




Figure 12 - 8-slot PCI I/O assembly

Physical Slot	Card size	PCI controller, bus	Operating speed, voltage
0	short	0, B	33 MHz, 5V
1	short	0, B	33 MHz, 5V
2	full length or short	0, B	33 MHz, 5V
3	full length or short	0, A	66 or 33 MHz, 3.3V
4	full length or short	1, B	33 MHz, 5V
5	full length or short	1, B	33 MHz, 5V
6	full length or short	1, B	33 MHz, 5V
7	full length or short	1, A	66 or 33 MHz, 3.3V

- *Note: Due to the placement of the voltage regulator, full-length cards cannot be used in slots 0 and 1.*






Figure 13 - Layout of Buses in an 8-slot PCI I/O assembly

6-slot cPCI I/O Assembly (Sun Fire 3800 Server Only)

The Sun Fire 3800 server has its own specific I/O assembly, as the mounting provided for in the Sun Fire 3800 system is different than that in the Sun Fire 4800-6800 servers. The resulting space allows for an I/O assembly comprising 6 cPCI slots.

cPCI is a different form factor than PCI. The cPCI cards are planar and can be inserted from the front of the I/O assembly without removing the entire assembly from the system chassis. Since the I/O assembly does not have to be removed to insert or remove cPCI cards, the cPCI cards are hot-swappable, with Dynamic Reconfiguration, in Sun Fire 3800-6800 systems.




Figure 14 - 6-slot CompactPCI I/O Assembly

Each cPCI assembly has dual PCI controllers, each with dual PCI buses. On each controller, one bus supports a single 66 MHz cPCI slot and the second bus supports two 33 MHz cPCI slots. A total of six cPCI cards are supported per cPCI I/O assembly. Both 32-bit and 64-bit cards are supported in all slots.






Figure 15 - Layout of Buses in a 6-slot PCI I/O assembly

Physical Slot	Card size	PCI controller, bus	Slot speed, voltage
0	3U cPCI card only	0, A	66 or 33 MHz, 3.3V
1	3U cPCI card only	1, A	66 or 33 MHz, 3.3V
2	3U cPCI card only	0, B	33 MHz, 5V
3	3U cPCI card only	0, B	33 MHz, 5V
4	3U cPCI card only	1, B	33 MHz, 5V
5	3U cPCI card only	1, B	33 MHz, 5V

4-slot cPCI I/O Assembly (for Sun Fire 4800, 4810 and 6800 servers)⁸

The 4-slot cPCI I/O assembly is almost identical to the 6-slot cPCI I/O assembly. This I/O assembly is designed for the Sun Fire 4800, 4810 and 6800 servers and has the same dimensions (and is designed to fit into the same bays) as the 8-slot PCI I/O assembly. Because cPCI adapters are physically larger at their edge than conventional PCI adapters, there is only enough space for four cPCI adapters, as opposed to the 8-slot PCI I/O assembly. As in the 6-slot cPCI I/O assembly, the adapter cards are hot-swappable. The cPCI I/O assembly supports fewer I/O slots than the PCI assembly because the cPCI cards are larger and longer than PCI cards.

⁸ The 4-slot CompactPCI (and hence, support for cPCI on the Sun Fire 4800, 4810 and 6800 servers) is scheduled to be available in Q4CY01.




Physical Slot	Card size	PCI controller, bus	Slot speed, voltage
0			66 or 33 MHz, 3.3V
1			66 or 33 MHz, 3.3V
2			33 MHz, 5V
3			33 MHz, 5V

Figure 16 - 4-slot CompactPCI I/O assembly

Functionally, the 4-slot cPCI I/O assembly is different from the 6-slot cPCI I/O assembly in that it just supports only four cPCI cards. There are still two PCI controllers, each with two buses. Each controller supports one 66 MHz card on one bus and one 33 MHz card on the other bus. Both 32-bit and 64-bit cards are supported in all slots.

System Controller

The Sun Fire 3800-6800 System Controller (SC) is responsible for the overall control of the system. A separate embedded system, the system controller is responsible for managing and monitoring the system. The system controller provides both the main system clock and the administrative console to the system. Configuration of system components and Dynamic System Domains is done via the System Controller. The console can be accessed through either a network or a serial connection. Dual system controllers can be configured for increased availability in the Sun Fire 3800, 4800 and 4810 servers. The Sun Fire 6800 server comes standard with dual system controllers.




Figure 17 - System Controller board (Sun Fire 4800-6800 servers)

The SC communicates with all of the different boards of the system via the Console Bus. This bus is dedicated for control and management operations between the SC and the system boards. Communication external to the Sun Fire 3800-6800 server is done either via a serial connection or via an Ethernet connection. These external interfaces to the System Controller support direct serial, telnet and SNMP connections.

The Sun Fire 3800-6800 System Controller is similar to a combination of the Sun Enterprise 10000 System Service Processor, Control Board and Centerplane Support Board on one assembly.



Each System Controller includes one RS-232/423 serial interface and one 10/100baseT Ethernet interface. All System Controller features are available through either a Graphical User Interface (provided through Sun Management Center) or a command line interface. The System Controller is the same on the 4800, 4810 and 6800 servers but different from the System Controller on the 3800, which has a longer form factor.

System Controller Functions

The system controller is responsible for many different functions that all have an impact on system RAS, system management and resource management.

- **Virtual System Clock**

The Sun Fire system controller is responsible for generating system clock signals. The clock signal from the master system controller is distributed to the system boards. The system clock will automatically failover to the secondary system controller in the event of a system controller failure with no system downtime or interruption to any domain.

The system controller is also responsible for providing the Time-of-Day clock to all domains within the system. Each domain is able to maintain its own local skew to the Time-of-Day clock.

- **Virtual System Console**

The system controller provides the system console interface to each domain. Connections are made to the system controller and a selection is made as to which domain to connect. A connection to the console of a particular domain may be placed in *advise mode*. This allows other read-only connections to the console of that domain to view all console activity.

- **Virtual Key Switch**

Domains do not have individual system key switches; the system console implements a virtual key switch for each domain. This allows the system administrator to secure the domain in the same way a physical key would be used. The key switch supports the same four settings used on all Sun servers: Off, Normal operation, System Diagnostics and Secure Operation.

- **Power Control**

The system controller is used to control the main 48V power supplies within the system and is used to turn individual boards and field replaceable units on and off.

- **Environmental Monitoring and Reporting**

The system contains sensors for temperature, voltage and current. The system controller polls these sensors periodically and makes the environmental data available. An over-temperature event will cause alarms to be sent to the Sun Management Center software as well as to the Solaris Operating Environment. If necessary, the system controller may shut down various components to prevent damage.

- **Error Management**

The system controller acts on errors reported to it via system components. Using messages from different components along the error path, the system controller attempts to identify the particular component or components generating the error. The system controller is then responsible for marking the failed component so that it is removed from the system. This also prevents a bad component from being incorporated into a domain and causing a fault on that domain.

- **Serial EEPROM Support**

The system controller interfaces with the Serial EEPROM for each FRU.



- **Hardware Configuration Management**

Hardware configuration management refers to several different areas. First, the system controller is responsible for Automatic System Recovery (ASR). ASR process occurs after the failure of a domain and if the reboot-on-error flag is set. The default needs be altered to use ASR.

In the ASR process, the system controller checks each part making up a domain for errors by using POST (Power On Self Test). If a part fails POST, that part is removed or blacklisted from the system being booted. During the ASR process, all interconnects between boards and Fireplane switches are also tested. When encountering a faulty interconnect, ASR process will cause the CPU/Memory or I/O board be blacklisted. Because the MTBF of Fireplane switches is very high, this is usually the correct choice.

If the reboot-on-error flag is set, ASR will automatically reboot the domain(s) with all remaining resources that have passed the test.

The second area of hardware configuration management is the area of Dynamic Reconfiguration, configuring crossbar segments and domain creation. The system controller is responsible for performing the necessary configuration within the Sun Fireplane interconnect systems in order to identify which boards are used to make up each crossbar segment and each domain. This can be done when a domain is created or when resources are added deleted from a domain. In either case, the system interconnect must be advised as to which boards make up each domain so that it can isolate its activity to one domain at a time.

The last area of hardware configuration management is the creation of interconnect segments. The System Controller is used to assign Sun Fireplane switches to segments.

- **POST Management**

Like all Sun servers, the Sun Fire 3800-6800 servers run a series of tests when they are powered on. These tests, known as POST (Power On Self Test), are more extensive and more complete in the Sun Fire 3800-6800 server family than in previous generations of servers. One fundamental difference in the way POST is done is the use of the System Controller to initiate, sequence and schedule the tests. The System Controller is able to schedule POST testing of domains and boards in parallel as far as possible, helping to reduce the time needed to test a system and minimize downtime in the event of a component failure. This speeds up MTTR (Mean Time To Repair) and increases total availability.

- **Interface to Sun Management Center**

The System Controller runs a Sun Management Center agent, which allows for the forwarding of environmental data, alarms and configuration status to monitoring software, such as Sun Management Center. The interface also allows Dynamic Reconfiguration operations to be performed via Sun Management Center 3.0.

For more information, see the Sun Management Center section later in this document.

- **Enhanced security**

The design of the System Controller allows the roles of system administrator and domain administrator to be redefined. No longer is it necessary to have root access to perform any administrative tasks. By separating the role of system and domain administrator, the management of any domain can be assigned to individuals with fewer security privileges to enhance the overall system security.

- **Enhanced Auto Recovery (System Controller Failover)**

The Sun Fire 3800-6800 server line can be configured with two system controllers for high availability. In a high-availability system controller (SC) configuration, one SC serves as the main SC, which manages all the system resources, while the other SC serves as a spare. When certain conditions cause the main SC to fail, a failover from the main SC to the spare is triggered



automatically, without operator intervention. The spare SC assumes the role of the main and takes over all system controller responsibilities.

System Cooling

Sun Fire 3800-6800 systems support fully redundant cooling. If any of the cooling modules should fail, the remaining cooling modules (if they have been configured for redundancy) are able to compensate by increasing their speed to maintain safe operating temperatures. The Sun Fire 3800, 4800 and 4810 servers offer an optional redundancy kit to provide redundant fans, power supplies and System Controller. The Sun Fire 6800 server base package is fully redundant and does not require the redundancy kit.

The system controller is responsible for monitoring the temperature of the system and of the surrounding air. Two-speed fans are used in the Sun Fire 3800-6800 systems. Temperature sensors are placed throughout the system on the different system boards or embedded in the ASICs to measure the temperature not only of processors but also of other ASICs and components, as well as the ambient temperature. The system controller monitors this data. The system controller can adjust fan speed to maintain safe operating temperatures in response to ambient changes.

All cooling of the Sun Fire 3800-6800 system is from front to back. This removes the problem of maintaining sufficient airflow that arises when locating rack systems next to one another. In order to maintain sufficient cooling, the airflow through the system is physically routed and controlled. This routing also makes it possible to mount systems in third-party racks.

All cooling modules are hot swappable and may be replaced without interrupting service.

System Power

The Sun Fire 3800-6800 server family can be configured with N+1 redundant DC bulk power supplies, which are hot swappable. The DC bulk power supplies provide power for all of the Sun Fire 3800-6800 systems' internal requirements, on a load-shared 48VDC power bus.

The Sun Fire 3800-6800 server family and the Sun Fire Cabinet use a sophisticated Redundant Transfer Unit (RTU) that can choose between power sources for providing power to the entire System Cabinet. AC power is supplied to the systems through up to four independent 30-amp single-phase Redundant Transfer Switch (RTS) modules (one or two per RTU). Each RTU provides power to two or three AC 48VDC power converters.

The dual AC connections into the RTU should be supplied via separate circuit breakers and can be located on isolated power grids if a high level of availability is required. Where separate power sources are not available, an Uninterruptible Power Supply (UPS) unit can be connected to one RTS module to provide power in the event of a power failure. Each RTS module is hot swappable and fitted with service LEDs for increased availability and serviceability.

RTUs have no single point of failure. All the failover logic is in the redundant RTS modules.

The RTU is effectively a very fast switch with microprocessor control and decision making programming to take an incoming feed from either one power source or the other. It monitors the health of incoming power and can switch between a failing feed and a good feed before the system would experience a *brown out*. This feature creates a reliable, single, AC source, which can also be used by peripherals with single AC input cables. Exceptional redundancy and real-time checking has been built into the RTU to meet stringent safety requirements.



The RTU communicates with the System Controller Board and the Frame Manager to provide information on the status of AC power. The System Controller also has control of the RTU to facilitate service procedures.




Figure 18 - Sun Fire 4800-6800 Server Power Distribution

The Sun Fire 6800 server is configured with two RTUs, giving the system a total of four RTS modules. The RTS modules are hot swappable if it is necessary to replace a failed RTS unit.

	Sun Fire 3800	Sun Fire 4800/4810	Sun Fire 6800
Number of RTU	-	-	2
Number of RTS	-	-	4
Number of Power Cords	3	3	2 or 4 ⁹

RTU only comes with the Sun Fire Cabinet. Also the information on the number of power cords and total system power required are listed in the Installation Data section.

⁹ Sun Fire 6800 server requires 4 power cords when dual power grid is installed.



Sun StorEdge D240 Media Tray

The Sun Fire 3800-6800 server family supports the Sun StorEdge D240 Media Tray for booting or small storage needs. The Sun StorEdge D240 Media Tray is a 2U-high, rack-mounted unit into which two UltraWide SCSI3 disk drive and one DVD drive and one DDS4 tape drive can be fitted. The DVD and DDS4 devices can be replaced by extra disk drives if required.

The Sun StorEdge D240 Media Tray also includes redundant hot swappable power and cooling to facilitate continuous operation.

One Media Tray provides mirrored boot disks, as well as a DVD drive for software installation and a tape drive for low duty backup and archival. These comprise the peripherals typically needed by a system (or a domain).

The Media Tray's SCSI backplane can be split to provide boot disks for two domains via separate SCSI buses. Alternatively, this feature can be used to support mirrored boot disks for a single domain. The Sun StorEdge D240 Media Tray requires a PCI Ultra SCSI adapter to be configured in a PCI or cPCI slot for each domain it is connected.

A tabletop version of the Sun StorEdge D240 Media Tray is also available for use with the deskside tower version of the Sun Fire 4800 server.

Sun Fire Cabinet

The Sun Fire 3800, 4800 and 4810 servers can be mounted in a Sun Fire Cabinet which has been specifically designed for this new family of systems, plus any rack mounted Sun storage device that is connected to the systems.

Each Sun Fire Cabinet contains:

- A Frame Manager¹⁰
- One or two Redundant Transfer Units
- Two rack fan trays
- A patch panel

The Frame Manager is an LCD display that is located in the top right-hand corner of the Sun Fire Cabinet. It manages and monitors fans and RTU.

The Sun Fire Cabinet contains one or two AC Redundant Transfer Units. Each RTU contains one or two RTS modules.

Fan trays are part of the Sun Fire Cabinet. The cabinet can operate on one or two fan trays, located at the top of the cabinet.

A single patch panel is located in the front of the cabinet on the bottom of the rack chassis. The patch panel provides an interface between the rack's Frame Manager and the system controller(s).

The Sun Fire Cabinet will also house storage or peripheral devices. But because Sun Fire Cabinet is designed primarily to support Sun Fire servers with front to back cooling, side to side cooling devices are not supported.

¹⁰ Frame Manager functional availability is TBD.



System Packaging

	Height	Rack Units	Space remaining for other systems/ peripherals ¹¹	Max. systems per Sun Fire Cabinet
Sun Fire 3800 server	15.00 in, 38.1 cm	8.5 U	21 U	3
Sun Fire 4800 server	30.0 in, 76.2 cm	17.5 U	12 U	1
Sun Fire 4810 server	36.75 in, 93.4 cm	21 U	9 U	1
Sun Fire 6800 server	75 in, 190.5 cm	28 U	5 U	1
Sun StorEdge D240 Media Tray	3.25 in, 8.25 cm	2 U		

The following are possible rack configurations:

- One Sun Fire 6800 server, plus two Sun StorEdge D240 media trays, leaving 3 U free
- One Sun Fire 4810 server, plus two Sun StorEdge D240 media trays, leaving 7 U free
- One Sun Fire 4800 server, plus two Sun StorEdge D240 media trays, leaving 10 U free
- Three Sun Fire 3800 servers, plus three Sun StorEdge D240 media trays, leaving no free space (or 0.5 U to be precise)

The last configuration of three Sun Fire 3800 servers plus one media tray for each server requires the power from two RTU's mounted in the cabinet. The Sun Fire 3800 server when installed in a Sun Fire Cabinet requires a rack-extension kit(X4347A).

Using Third-Party Racks

Each of the Sun Fire 3800, 4800 and 4810 systems plus the Sun StorEdge D240 Media Tray are designed to be mounted in a standard 19-inch-wide commercial rack which conforms to EIA standard EIA-310-D¹². The Sun Fire 4810 server is designed to be mounted in a 19-inch-wide, 30-inch-deep rack. The physical dimensions for mounting these servers in a commercial rack are:

Figure 19 - System Dimensions

	Width	Depth	Height
Sun Fire 3800	17.72 in, 450 mm	37.75 in, 959 mm	15.00 in, 381 mm
Sun Fire 4800	17.56 in, 446 mm	28.50 in, 724 mm	30.0 in, 762 mm
Sun Fire 4810	17.56 in, 446 mm	23.30 in, 592 mm	36.75 in, 934 mm
Sun Fire 6800	N/A	N/A	N/A

¹¹ After the installation of 1 x system, 2 x RTU and 1 x Sun StorEdge D240 Media Tray

¹² Sun supplied rack mount kits are designed for use in the Sun Fire Cabinet and are not adjustable. Check with the manufacturer of the rack for support constraints. For support guidelines when mounting servers in a 3rd party rack, refer to "Thermal Guidelines for mounting Sun Fire 4810/4800/3800 systems in a non-Sun cabinet" document.



3rd Party Rack Considerations:

The RTU is not supported in third-party racks and there is no support for dual power grids when Sun Fire 3800, 4800 or 4810 servers are installed in a third-party rack.

Sun does not provide support for third-party racks, but if the rack and its installation meet all relevant site planning requirements (e.g. sufficient airflow for front to back cooling), the use of racks other than the Sun Fire Cabinet is permitted.

The Sun Fire 6800 server may not be installed in a third party rack due to its size.



Installation Data

Cabinet

SYSTEM	Width	Height	Depth	Weight
Sun Fire 3800, rack-mount	17.72 in, 45.0 cm	15.00 in, 38.1 cm	34.80 in, 88.4 cm	200 lbs, 90.7 kg
Sun Fire 4800, deskside tower	18.90 in, 48.1 cm	32.40 in, 82.3 cm	34.2 in, 86.9 cm	289 lbs, 131.1 kg
Sun Fire 4800, rack-mount	17.56 in, 44.6 cm	30.00 in, 76.2 cm	28.50 in, 72.4 cm	289 lbs, 131.1 kg
Sun Fire 4810, rack-mount	17.56 in, 44.6 cm	36.75 in, 93.4 cm	23.62 in, 60.0 cm	305 lbs, 138.3 kg
Sun Fire 6800, incl. rack	24.00 in, 61.0 cm	75.00 in, 190.5 cm	53.00 in, 134.6 cm	1200 lbs, 544.3 kg
Sun Fire Cabinet (with doors)	24.00 in, 61.0 cm	75.00 in, 190.5 cm	37in, 94.0 cm	325 lbs, 147 kg
Sun StorEdge D240 Media Tray	17.56 in, 44.6 cm	3.25 in, 8.25 cm	23.0 in, 58.4 cm	50 lbs, 23 kg

Electrical and Thermal Specifications

Sun Fire 3800 Server

The Sun Fire 3800 server can accommodate 100VAC, 120 VAC, or 200-240 VAC as power sources. Please choose data from only one column for planning.

	100VAC	120VAC	200-240VAC
Total Current Required (maximum configuration)	24A	19.7A	11.4A
Frequency	47-63 Hz		
Maximum power	2,280 W	2,240 W	2,170 W
Volt-AMP	2,400 VA	2,360 VA	2,280 VA
Heat load	7,775 BTUs/hr	7,645 BTUs/hr	7,385 BTUs/hr
Power factor	0.95 (with Sun products)		
Connector type - US	3 x NEMA 5-15R		3 x NEMA 6-15R
Connector type - International	-		3 x 10A, single phase IEC320
Number of Power Cords	3	3	3



Sun Fire 4800, 4810 and 6800 Servers

	Sun Fire 4800	Sun Fire 4810	Sun Fire 6800
Voltage range	200-240 VAC		
Total Current Required (maximum configuration)	16A	16.4A	34A with 2 RTS
Frequency	47-63 Hz		
Maximum power	3,040 W	3,120 W	6,460 W
Volt-AMP	3,200 VA	3,280 VA	6,800 VA
Heat load	10,370 BTUs/hr	10,625 BTUs/hr	22,030 BTUs/hr
Power factor	0.95 (with Sun products)		
Connector type - US	3 x NEMA 6-15R		4 x NEMA L6-30R
Connector type - International	3 x 10A, single phase IEC320		4 x 32 A, IEC309
Number of Power Cords	3		4

Environmental Specifications

Operating Environment	
<ul style="list-style-type: none"> Humidity Temperature Altitude Vibration Shock Declared acoustics 	20% to 80% relative, noncondensing, 27C max wet bulb 5 to 35° C (41 to 95° F) at sea level 0 to 3,000 meters (0 to 10,000 feet) 0.15 G's, 5 to 500Hz, swept-sine 3 G's, 11 ms, half-sine 69 dB(A)
Non-operating Environment	
<ul style="list-style-type: none"> Humidity Temperature Altitude Vibration Threshold 	5-93% relative, noncondensing -20 to 60° C (-4 to 140° F) 0 to 12,000 meters (0 to 40,000 feet) 1.5mm single amplitude/0.6 G's, 5 to 500Hz, swept-sine, vertical 0.75mm single amplitude/0.3 G's, 5 to 500Hz, swept-sine, horizontal 0.75 meter/sec

Regulations

Safety	UL 1950, CSA C22.2 950, TUB EN60950, CB Scheme with all Country Deviations
Emissions	FCC Class A, ICES-003 Class A, EN55022 Class A, VCCI Class A, BSMI Class A, EN61000-3-2 & EN61000-3-3
Immunity	EN55024



Resource Management

Dynamic System Domains

Dynamic System Domains (or domains for short) are a concept introduced to the open systems computing world by Sun Microsystems with the Sun Enterprise 10000 server. With domains, a single computer system can be partitioned into separate "computers," or separate virtual servers, each with its own instance of the Solaris Operating Environment and its own I/O devices. This is a powerful concept that has been used with great success in the mainframe world for a number of years. The success of the Sun E10000 server has shown that domains are equally powerful for open systems computing.

The minimum configuration for a domain requires at least one CPU/Memory board for processors and memory and at least one I/O assembly for connections to the network and to private boot devices. These minimum requirements lead to two domains being supported in the Sun Fire 3800, 4800 and 4810 servers and four domains being supported in the Sun Fire 6800 system.

The Sun Fire 3800-6800 family of servers is configured at the factory with only one domain. There is no reconfiguration required upon first use. To deploy more than one domain will require reconfiguration through the System Controller.

In order to maximize the availability potential of a Sun Fire 3800-6800 system when creating more than one domain, we need to discuss the crossbar segment and domains more thoroughly.

Sun Fireplane Interconnect System

The Sun Fireplane interconnect system provides a switched crossbar interconnection establishing an ECC-protected data path between major system components (CPUs, memory, I/O controllers). The Sun Fireplane interconnect system can be configured either as a single segment or as two segments.

A segment refers to all or part of the Sun Fireplane interconnect system. Segments are a way of reconfiguring the centerplane or Sun Fireplane interconnect system into two parts. This is equivalent to replacing the single centerplane board with two separate boards. The interconnect system can be configured either as one big segment or as two smaller segments. Because the interconnect system is implemented as two hot-pluggable Sun Fireplane switches on a Sun Fire 4800 or 4810 server (and is implemented as four Sun Fireplane switches on a Sun Fire 6800), segmenting is essentially the division of these Sun Fireplane switches into two groups and the management of the pool of resources as two separate resource pools.

The failure of a Sun Fireplane Interconnect in a segment affects only that segment and the domains associated with that segment. The other segments and domains continue unaffected. This provides a level of redundancy virtually unmatched by other vendors. It is also an improvement over the Sun Enterprise 10000 server where a centerplane error would affect all domains in the system.

Domains

In each segment, resources such as CPU/Memory boards and I/O assemblies can be divided along the FRU boundary into one or two groups of resources. This process creates domains on a Sun Fire 3800-6800 server. Other than the sharing of the Sun Fireplane switches, this type of division is still quite thorough and would provide a highly isolated, secured environment for workloads with such requirements.



When a pair of Sun Fireplane Interconnect boards is configured to support domains, access to the Sun Fireplane Interconnect board alternates between the domains on each clock cycle. When a domain is active on the Sun Fireplane switch, the ports connecting to boards not configured as a part of that domain are disabled.

In the Sun Fire 6800 server, the Sun Fireplane Switches operate in pairs and a single segment on a Sun Fire 6800 server is made up of two Sun Fireplane Switches. In the case of a segmented system, the links between the two pairs of Sun Fireplane Interconnect boards in the Sun Fire 6800 server are disabled. Rather than splitting each transaction across both pairs of Sun Fireplane switches, only a single pair is accessible to each segment. Within a segment, the Sun Fireplane switches may be configured to support a pair of domains. This allows the Sun Fire 6800 server to support up to four domains.

Choose Crossbar Segment or Domains to Maximize Availability

Segments provide complete hardware-level isolation from errors occurring in other segment because the links between the segments are disconnected. Domains provide error isolation by configuring the ports to turn on/off the links on the Sun Fireplane Interconnect boards. Domains provide adequate fault isolation for failures in applications, O/S, CPUs, memory, CPU/Memory boards, I/O assemblies and I/O cards. In summary, customers should first implement two segments whenever two or more domains are needed on a Sun Fire 3800-6800 system in order to improve the availability. This is because two segments isolate each other from the failures of any Sun Fireplane switch, a capability not offered by domains.

A special case of isolation occurs when a Sun Fire 6800 system is segmented exactly along its spine. In this case, since each segment resides within a physical half of the chassis, the power supplies between the two segments are isolated as well. For this reason, Sun highly recommends that customers implement segments along the power boundary on a Sun Fire 6800 server to enhance availability.

A Sun Fire 6800 server is an ideal platform on which to implement a two-node High-Availability (HA) cluster within a box. By carefully observing the power boundary, a two-node HA cluster can be configured to achieve domains fault isolation, power plane isolation and power sourcing isolation (if the dual power-grids feature is in use).

The following table highlights the best configurations for maximizing the overall system availability against the number of domains required on a Sun Fire 3800, 4800 4810, or 6800 system.

Sun Fire Model	Sun Fire 3800, 4800, or 4810		Sun Fire 6800		
Number of domains needed	1	2	1	2	3 or 4
Number of segments to implement	1	2	1	2	2
Number of Sun Fireplane switches per segment	2	1	4	2	2
Number of domains per segment	1 or 2	1	1 or 2	1 or 2	1 or 2
Total maximum number of domains permitted	2	2	2	4	4
To implement separated power planes	N/A	N/A	Not needed	Yes	Yes

Figure 20 - Recommendations on configuring for maximum availability



Recovery from Crossbar Segment Outages

During an unlikely event of incurring a non-recoverable Sun Fireplane Interconnect failure, the whole segment and the domains it controls will be lost. The process of restoring a failed domain depends on several variables:

- If there is only one segment defined in the system (there are redundant Sun Fireplane switches available), the customer can reconfigure the crossbar segment by eliminating the failed Sun Fireplane Interconnect board. The customer can manually reboot one or two domains as needed. Although the system data bandwidth will be cut in half, forming a functioning crossbar segment requires only one Sun Fireplane Interconnect on a Sun Fire 3800-4800 system, or two switches on a Sun Fire 6800 system.
- If there are two crossbar segments defined on the system, the customer should replace the failed Sun Fireplane switch. Following this replacement, the customer may manually reboot all failed domains.
- Alternatively, the customer can bring up a failed domain without fixing the failed Sun Fireplane Interconnect first, by following these steps:
 - Ensure that there is room for one more domain on the surviving crossbar segment.
 - Configure a new domain in the surviving segment to include all the resources needed. Mostly likely, these are resources from the failed segment and domain.
 - Reboot the failed domain on the surviving segment.
 - Because the failed domain will have a new identifier (i.e. It was domain A, now it is domain C or D), it must have a new MAC address and hostid.

Solaris Resource Manager™ Software

When there is a need to divide system resources such as processor or memory resources in a domain, Solaris Resource Manager™ software is an ideal tool to use.

Based on a fair-share concept of scheduling resources among several competing applications, Solaris Resource Manager software is a control extension to the Solaris Operating Environment. It is comprised of the following kernel components:

- a fair-share scheduler module
- a user-interface module
- a resource-limits control module

Together, these components provide control of the processor usage and virtual memory space consumption.

Solaris Resource Manager software has been successfully deployed by many Sun Enterprise 3500-6500 and 10000 customers to implement application consolidation and is expected to play an important role in assisting Sun Fire 3800-6800 server customers to achieve capacity granularity effectively.



Processor Sets

The Solaris Operating Environment allows processors at different speeds to coexist within a single Sun Fire system. It is becoming very important to be able to assign a specific processor or group of processors to a particular application or process thread. Processor Sets is a mechanism for dividing the set of processors in an MP system into disjoint subsets, where threads will not generally migrate between processors in different subsets. These subsets provide a flexible way for managing processor allocation in multiprocessors.

The partitioning functionality is available to the user as processor sets - groups of processors to which processes and Lightweight Processes (LWPs) may be exclusively bound. This exclusive binding prevents other processes in the system from using those processors. When different speed CPUs are running in the same domain, processor sets can be used to assign certain jobs to the faster CPUs and other jobs to the slower CPUs.



System Management

System Controller for System Management

The general capabilities of the System Controller are discussed in the System Architecture section found earlier in this document.

In the area of System Management, the System Controller is responsible for a number of significant aspects of the Sun Fire 3800-6800 system:

- Virtual consoles are provided by the System Console
- Power On Self Test (POST) and provision of OpenBoot firmware
- Capture and forwarding of environmental data
- Configuration of the Sun Fireplane interconnect system (management of Segments, Dynamic System Domains, Dynamic Reconfiguration operations, etc.)

Sun Management Center 3.0 Software

Sun Management Center Software is an open, extensible, standards-based system management solution that facilitates active, enterprise-wide management of Sun system products and their subsystems and components.

The Sun Fire 3800-6800 servers employ Sun Management Center Software to provide the GUI (Graphical User Interface) for managing the system. Whereas an available command line interface provides full functionality, Sun Management Center 3.0 software is capable of providing the vast majority of such functions through the GUI.

Sun Management Center Software is also the key technology for delivering Sun Remote Services (SRS); therefore, Sun highly recommends the use of Sun Management Center Software.


Sun Management Center Software interfaced with a networked System Controller gives the administrator a high level of system and resource management. The Sun Management Center software provides features such as Dynamic System Domain management, as well as a GUI interface to Dynamic Reconfiguration and other system configuration commands. Additionally, it provides photo-realistic images of hardware information, environmental monitoring and propagation of *alarms* to associated devices.

Sun Management Center Software uses a three-tiered, client-server architecture to provide a high level of scalability. The three tiers in the Sun Management Center Software architecture consist of:

- The console layer
- The server layer
- The agent layer

By utilizing an object-centric, dynamic, intelligent design, Sun Management Center Software enables processing to be done at the source where the managed object is located. By distributing management intelligence, Sun Management Center Software enhances the reliability, availability and serviceability of the enterprise network.





Sun Management Center Software employs autonomous agent technology, a technique in which agents are not dependent on other software components. Based on SNMP technology, these agents collect and process data locally and can act on data to send SNMP traps, run processes and so on, even if connection to the manager is severed. These intelligent agents can monitor Dynamic System Domains as well as initiate alarms, notification or specific actions based on collected data or messages through customizable rules and thresholds.

Because no one management package can meet every demand of the enterprise, additional modules can be added dynamically to Sun Management Center Software agents from the console without disruption to the management system. This capability provides administrators with a flexible and extensible application that can be tailored to form a comprehensive enterprise management solution that adapts to their changing needs.

Third-party enterprise-management tools - such as Tivoli TES, Unicenter TNG and OpenView VantagePoint Operations - are available for application monitoring and management, as well as managing heterogeneous environments. Sun Management Center 3.0 Software is designed for monitoring the Solaris Operating Environment and Sun Fire servers. Sun Management Center Software can be integrated with third-party enterprise management tools to enable superior management of Sun servers within a heterogeneous compute environment.

System Diagnostics

The System Controller, in conjunction with Sun Management Center Software, is responsible for boot time and dynamic reconfiguration hardware diagnostics. The System Controller is also responsible for the management and reporting of ongoing environmental and other diagnostic information. Additionally, SunVTS software provides functionality for testing general hardware subsystems.



Automatic System Recovery

In the event of a hardware failure which causes the Solaris Operating Environment to reboot, the System Controller will attempt to identify which unit(s) failed. The extensive use of parity checking between subsystems of the Sun Fire 3800-6800 server family means that errors can be isolated to a single serviceable part for later replacement by service personnel.

If a failure causes a domain to go down, the System Controller will attempt to reboot the Solaris Operating Environment on the domain(s) affected, if the reboot-on-error flag is set. (The default for this flag is not set. Therefore it is necessary to set the flag to use ASR). Before this is initiated, extensive testing of subsystems (processors, memory modules, ASICs, etc.) is performed to establish if there are any other parts of the system that may have failed or may be failing. During the ASR process, all interconnects between boards and Fireplane switches are also tested. When encountering a faulty interconnect, the ASR process will cause the connected CPU/Memory or I/O board be blacklisted. Once a configuration of valid system components is established for each domain, the System Controller will boot that domain. If the exclusion of failed parts means that a domain no longer comprises a valid hardware configuration, the domain will remain down.

During this recovery process, the System Controller records the failing components into a file and *blacklists* the failed CPU and memory as not available for use. Blacklisting is the mechanism where faulty components are taken off-line either by POST as described above or manually via the SC.

This Automatic System Recovery procedure helps ensure that the Solaris Operating Environment will not attempt to boot using components that may be faulty, which, if attempted, could lead to the corruption of user data.

Manual intervention is required to reboot all affected domains if the failure is in a Sun Fireplane switch. The exact procedure for manual recovery would be dependent on the configuration of the Sun Fireplane switch.

Remote Monitoring through Sun Remote Services (SRS)

Sun Remote Services is a suite of remote management services designed to improve mission-critical support and systems availability. The core offering - SRS Event Monitoring and Management Service - uses tools, technologies and expertise to enable proactive problem detection and prompt resolution of customers' system events. System experts in Sun service centers analyze information gathered by SRS. Additional SRS services will be added in the future to further support system management and availability needs.

SRS's continuous system monitoring utilizes an intelligent, agent-based architecture to monitor key variables on a Sun Fire 3800-6800 system remotely. System health and utilization information is regularly collected, filtered and compiled using an advanced management toolset. When a problem is detected, an alarm is automatically generated to notify Sun's service organization and initiate the resolution process. Sun's advanced management tools also provide key diagnostic information to help pinpoint and resolve a problem quickly.

The key benefits of the SRS services include:

- Proactive notification of Sun Service in the event of problems
- Enhanced technical expertise of customers' internal staff with Sun specialists and advanced technology
- Faster problem resolution through early alerts and notification
- Reduced workload of IT staff by meeting day-to-day monitoring needs



- Automated technology to provide better monitoring capabilities with around-the-clock coverage
- Fast, direct access to Sun experts and design engineers to help identify and resolve problems more quickly
- Important information for future planning and redistribution of resources based on the monitored information and reports

Remote systems monitoring and remote predictive failure reporting are features of the mission-critical or business-critical support. *Please refer to the Service and Support section for details.*



Software

Solaris™ Operating Environment

The Solaris Operating Environment is one of the key features of the Sun Fire 3800-6800 server family. The advanced features it offers in the areas of scalability, availability, reliability, security and management are fundamental to the value of Sun Fire servers as platforms for today's users and applications.

Scalability

Solaris 8 Operating Environment is a proven platform for deploying applications that require maximum scalability. Built around a Symmetric Multi-Processor architecture - where system tasks are multi-threaded to scale with user workload - the Solaris Operating Environment has more than seven year's experience of scaling systems up to 64 processors. Solaris 8 Operating Environment became Generally Available (GA) in February 2000 and all key application software packages have been tested for scaling to 64 processors.

Dynamic Reconfiguration

Dynamic Reconfiguration (DR) is the software feature built into the Solaris Operating Environment that enables the hardware configuration to be altered while the operating environment and user processes (applications) continue to run. The capability to physically add and remove components from the system is a feature of the Sun Fireplane interconnect system and of the Sun Fire 3800-6800 system design, where the interconnect allows hardware changes without requiring a reset and the board connections are made in a way that does not stress power distribution within the system.

The DR feature allows a machine to be serviced without bringing it down. DR is used to deconfigure the resources on a faulty system board from a running system so that it may be hot swapped from the system. The repaired board or replacement may then be hot inserted into the machine. After testing the board, DR configures the resources on the board into the running system. DR also makes hardware upgrades possible without bringing the system down.

DR enhances Dynamic System Domains by allowing components to be actively added to or removed from live domains. This feature allows the resources available on system boards to be shared between domains. Also available will be command-line utilities to perform DR operations, as well as a GUI for DR operations, provided as part of Sun Management Center Software (similar to the functionality that exists for the Sun Enterprise 3000-6000 servers and the Sun 3500-6500 servers).

A new feature to be included is automatic, hot insertion of system boards. The System Controller maintains a listing of which system boards can be used in each domain on the system, including which boards can alternate between multiple domains. If a board is placed into a slot that is allocated to a particular domain, the system can be configured to automatically perform a DR operation to add the board as soon as it has been tested.

IP Network Multipathing

IP Network Multipathing is a new feature in Solaris 8 Operating Environment providing for failover and IP-link aggregation. It has a number of advantages over Alternate Pathing and improves the availability, serviceability and performance of a Sun Fire 3800-6800 server:



- **Failure Detection:** Ability to detect when a network adapter has failed and to switch automatically (failover) the network access to an alternate network adapter. This process assumes that the customer has configured an alternate network adapter.
- **Repair Detection:** Ability to detect when a network adapter that failed previously has been repaired and to switch back automatically (failback) the network access to an alternate network adapter. This assumes that the customer has enabled failbacks.
- **Outbound Load Spreading:** Outbound network packets are spread across multiple network adapters without affecting the ordering of packets in order to achieve higher throughput. Load spreading occurs only when the network traffic is flowing to multiple destinations using multiple connections.

For further information, refer to *IP Network Multipathing Administration Guide*.

Alternate Pathing will not be supported on Sun Fire 3800-6800 servers.

Multiplexed I/O (MPxIO)

To increase both availability and I/O bandwidth to the storage pools, including failover capability on storage devices that are attached to a Sun Fire 3800-6800 server with multiple host controller interfaces, the Solaris Operating Environment has an enhanced core I/O framework to represent and manage devices that are accessible through multiple host controller interfaces.

MPxIO is available as a patch to Solaris 8 7/01 Operating Environment from Network Storage and will support FC-AL adapters. Major features provided include:

- Support of automatic failover to route I/O requests through alternate active paths on transport failures
- Support of manual switchover to enable Dynamic Reconfiguration
- Load balancing - the ability to route I/O requests through different paths for better utilization of host controller resources and for improved I/O performance

Availability Features

The Solaris 8 Operating Environment supports a number of software components that increase overall availability as well as resource management.

- Sun Cluster 3.0 software extends the Solaris Operating Environment with the Sun Cluster framework, enabling the use of core Solaris Operating Environment services such as devices, file systems and networks, seamlessly across a tightly coupled cluster while maintaining full Solaris Operating Environment compatibility for existing applications.

Sun Cluster 3.0 software is a platform that provides High Availability (HA) to everyday Solaris Operating Environment applications through continuous network and data availability. Applications that are written to use the simple Sun Cluster 3.0 software API can accomplish even higher availability as well as increased scalability.



Sun Cluster 3.0 software significantly differentiates Sun's Cluster product from competitors' products by integrating clustering functionality into the Solaris Operating Environment and adding advanced features such as Global Devices, Global File System, Global IP Services and Scalable Services. These advanced features provide a managed application environment, helping to enable Sun to leapfrog the competition with one of the industry's best platforms for highly available and scalable network, data and application services.

- Solaris Resource Manager™ software provides fine-grained control of system resources, helping to ensure a consistent level of service to users, groups and applications.
- Solaris Bandwidth Manager enhances the customer's ability to control and provision IP traffic priorities and bandwidth, helping to enable network resource availability.

Global Language Support Enhancements

The Solaris Operating Environment offers an internationalized framework that makes it easy for companies and individual users to work effectively in a multi-lingual world.

The Solaris 8 Operating Environment provides support for over 90 locales, a new, intuitive interface for installing languages, expanded Unicode support and improved data interoperability utilities.

For more information, see the *International Language Environments Guide*:
<http://docs.sun.com:80/ab2/coll.45.13/I18NDG>

Application Compatibility and Investment Protection Guaranteed

The Solaris Application Guarantee program will help ensure that existing applications written to the Solaris Application Binary Interface (ABI) will run without modification on the Solaris 8 Operating Environment. For SPARC processor platform customers, this means that existing SPARC processor binaries will run unmodified on Sun's UltraSPARC III processor based systems.

This combination of hardware and software compatibility offers industry leading investment protection for Sun's customers as they upgrade. The Solaris Operating Environment Application Guarantee reflects the confidence Sun has in the compatibility of applications from one release of Solaris Operating Environment to the next and is designed to make re-qualification and porting a thing of the past.

- *NOTE: The Solaris Application Guarantee is a limited-time offer. Test Suite software results must be completed and the registration form submitted by June 30, 2002. Failures of registered applications must be reported to Sun by September 30, 2002.*

Capacity on Demand (COD 2.0)

The Capacity on Demand option is for both the SF 3800-6800 and Sun Fire 12K/15K servers. COD 2.0 provides rapid access to extra capacity for unpredictable peaks in demand. There is no limit on adding COD Uniboards within the configuration guidelines of the server. The maximum number of instant access resources available on Sun Fire 3800-6800 servers is four CPUs.

With COD 2.0, SF 3800-6800 servers can be configured with any number of COD CPU resources depending on customer requirements. When the COD CPUs are needed, customers activate CPUs installed in a Sun Fire 3800-6800 server by purchasing additional right (RTU) licenses. This gives customers rapid access to additional capacity so that they can more quickly respond to sudden increases in demand. With COD, customers have systems that accommodate growth, but only pay for the capacity they use, better matching capital outlays to resource utilization. The COD 2.0 option helps enable customers to configure new or existing Sun Fire systems with additional "capacity on demand" processor and memory resources at lower up-front costs.



The Capacity on Demand (COD) option provides additional processing resources that you pay for when you use them. Through the COD option, you purchase and install unlicensed COD CPU/Memory boards in your system. Each COD CPU/Memory board contains four CPUs, which are considered as available processing resources. However, you do not have the right to use these COD CPUs until you also purchase the right-to-use (RTU) licenses for them. The purchase of a COD RTU license entitles you to receive a license key, which helps enable the appropriate number of COD processors. The COD 2.0 software that is included in all Sun Fire server systems self-monitors the COD resources and virtually eliminates the need for the customer to provide system usage reports or system monitoring. The COD 2.0 software for the Sun Fire 3800-6800 is the 5.14.0 firmware version.

The Sun Fire COD uniboards include three different configured options - The first COD uniboard option contains 4 unlicensed COD CPUs running at 1050MHz; 4x8MB Ecache;8 banks of 2GB memory options (16GB total for the board). The second COD uniboard contains 4 unlicensed COD CPUs running at 1.2 GHz; 4x8MB Ecache;8 banks of 2GB memory options (16GB total for the board). The third COD uniboard contains 4 unlicensed COD CPUs running at 1.2 Ghz; 8 banks of 4GB memory options (32 x 1GB DIMMs, 32GB total for the board). Also, x-option COD uniboards are also available with 1050MHz and 1.2 GHz processors. One Sun Fire COD 2.0 RTU is required for each COD CPU to be activated.

Sun Fire Capacity on Demand customers will be required to abide by specific legal terms and conditions, which can be found on the Terms and Conditions addendum. A copy of the contract addendum can be found off the following web link (select COD 2.0):

<http://sunlegal.central/Legal/GlobalAccts/ContractsPUBLIC.html>

Because the memory is directly controlled by the CPU, the physical memory banks associated with a particular CPU will be disabled if that CPU is disabled. It is possible to have a single CPU/Memory board with a combination of enabled/disabled CPUs. Only the memory associated with the enabled CPUs would be accessible. If a customer experiences a CPU failure on a COD system they can use a COD CPU as a hot spare as long as the total number of active CPUs in the system are licensed. You can purchase COD RTU licenses at any time from your Sun sales representative or reseller. You can obtain a license key (for the COD resources purchased) from the Sun License Center.

The Service Maintenance contract is structured so that you pay only 20% for the COD board, and an additional 20% for each licensed CPU on the COD board.

The COD 2.0 uniboards will be offered with each new iteration of CPU/Memory uniboard speeds going forward in the future for Sun Fire 12K/15K, and Sun Fire 3800-6800 servers.



Ordering Information

Standard Configurations give the customer the best lead time with the best discount. These configurations have limited flexibility and allow the installation of building block options in the factory.

There are many reasons why Sun has chosen to provide standard configurations. Primarily, they reduce lead time to customers, improve system availability, increase the quality of systems, simplify the sales order process and offer attractive pricing.

NOTE: Standard Configurations are the only method available to configure and to order Sun Fire servers.

Standard Configurations

Standard configurations start with a base package and are built using building blocks. A building block is a fixed set of components that are pre-built and tested to enable Sun Operations to assemble a system in a short time and provide high quality. These bundles include:

- Base Package
- CPU/Memory Board Bundles
- Redundancy Kit (F3800, F4800, F4810)
- cPCI and PCI I/O cards
- Sun Fire Cabinet (one system per cabinet)
- Media Tray
- Pretested Storage Devices

At a minimum, a Base Package, I/O assembly with 1 NIC card that supports 10/100-Mb Ethernet and CPU/Memory Board bundle are required.

Configuration Rules

The following section describes the rules for a minimum configuration. For minimum and maximum quantities of components supported, refer to Figure 7 Outline of Sun Fire 3800-6800.

Sun Fire 3800 Server

The Sun Fire 3800 server base package contains 1 CPU/Memory board bundle with either 2-CPU/2-GB Memory or 4-CPU/4-GB Memory, 2 cPCI I/O assemblies, 2 Power Supplies, 3 Fan Trays and 1 System Controller. This system does not require a customer to select a CPU/Memory board in a minimum configuration. The customer must select the following as a minimum system:

- 1 Base Package
- 1 cPCI I/O card that supports 10/100-Mb Ethernet
- 2 Power Cords



Sun Fire 4800 and Sun Fire 4810 Servers

The Sun Fire 4800 and Sun Fire 4810 server base packages contain: 2 Power Supplies, 2 Fan Trays and 1 System Controller. A minimum configuration must contain the following:

- 1 Base Package
- 1 CPU/Memory Board Bundle
- 2 PCI I/O Assemblies or 2 cPCI I/O Assemblies
- 1 PCI or cPCI I/O card that supports 10/100-Mb Ethernet
- 2 Power Cords

Sun Fire 6800 Server

The Sun Fire 6800 server base package contains: 6 Power Supplies, 4 Fan Trays and 2 System Controllers, 2 RTUs with 4 RTS units. This base package comes with full redundancy built in. A minimum configuration must contain the following:

- 1 Base Package
- 1 CPU/Memory Board Bundle
- 4 PCI I/O Assemblies or 4 cPCI I/O Assemblies
- 1 PCI or cPCI I/O card that supports 10/100-Mb Ethernet
- 4 Power Cords



PCI and CompactPCI Adapters

I/O connectivity is available for the Sun Fire 3800-6800 via a comprehensive range of PCI and CompactPCI adapters. They are:

Adapter	Format	Part number	Type	Speed
Storage Cards				
Single Channel UltraSCSI	cPCI	1232A	3U, 3.3/5V, 64bit	33 MHz
Dual FC-AL 100MB/s - Optical	cPCI	6748A	3U, 3.3/5V, 64bit	66 MHz
Dual Ultra Differential SCSI	cPCI	6749A	3U, 3.3/5V, 64bit	33 MHz
Single Channel Ultra SCSI SE	PCI	1032A	short, 3.3/5V, 32bit	33 MHz
Dual FC-AL 100MB/s - Optical	PCI	6727A	short, 3.3/5V, 64bit	66 MHz
Dual Channel SCSI - Ultra SE	PCI	6540A	long, 3.3/5V, 32bit	33 MHz
Dual Differential Ultra SCSI	PCI	6541A	long, 3.3/5V, 32bit	33 MHz
Single FC Network Adapter 100MB/s	PCI	6799A	short, 3.3/5V, 64bit	66 MHz
Network Cards				
Combo 10/100 ethernet (shown as a storage card also)	cPCI	1232A	3U, 3.3/5V, 64bit	33 MHz
QuadFast Ethernet (QFE)	cPCI	1234A	3U, 3.3/5V, 64bit	33 MHz
SunATM 155A/MFiber 3U	cPCI	1266A	3U, 3.3/5V, 64bit	33 MHz
SunATM 622A/MFiber 3U	cPCI	1268A	3U, 3.3/5V, 64bit	33 MHz
Gigabit Ethernet	CPCI	1261A	3U, 3.3/5V, 64bit	66 MHz
10 /100 Ethernet combo card (shows as a storage card also)	PCI	1032A	short, 3.3/5V, 32bit	33 MHz
Fast ethernet	PCI	1033A	short, 3.3/5V, 32bit	33 MHz
QFE	PCI	1034A	short, 3.3/5V, 64bit	33 MHz
HIPPI 1.0	PCI	X1071A	short, 5V, 32bit	66 MHz
Gigabit Ethernet	PCI	1141A	short, 3.3/5V, 64bit	66 MHz
High Speed Serial Interface Adapt.	PCI	1155A	short, 5V, 32bit	33 MHz
SunATM 155 MMF 4.0	PCI	1157A	short, 3.3/5V, 64bit	33 MHz
SunATM 155 UTP 4.0	PCI	1158A	short, 3.3/5V, 64bit	33 MHz
SunATM 622 /Mfiber Adapter 4.0	PCI	1159A	short, 3.3/5V, 64bit	66 MHz
Serial Asynchronous Interface Adapt.	PCI	2156A	short, 3.3/5V, 32bit	33 MHz



Upgrades

Key Messages

Sun is committed to offering upgrade programs as a way to make it easy for customers to move to the latest technology and to stay current while lowering their total cost of ownership. Hardware upgrades are critical to maintaining and growing our business. Customers are reassured knowing the Sun assets they buy have upgrade value.

What is the Upgrade Program?

Sun values our installed base customers. That is why we offer trade-in value for older Sun servers and upgrade paths to the latest server models. Sun will give a percentage discount off the list price of the new system being purchased based upon the type and age of the equipment being returned.

Servers can be traded in for servers, storage for storage and desktops for desktops. No substitutions (i.e. storage for servers) are permitted. We have created a set of upgrade paths that define the amount of trade-in value, as a percentage, that will be given depending on what is being traded in and what is being purchased.

Why this is good for our customers

The upgrade programs allow customers to write down or write off the equipment being traded in, which helps to offset the cost of the new equipment. Customers can remain competitive at a reduced total cost of ownership (TCO). Our upgrade programs are very flexible: we allow full system upgrades and/or component upgrades, (CPU, memory, disk, monitors).

Why this is good for Sun

These programs help protect our installed base from competitive takeouts. It is easier to go into customer sites and offer to take old equipment out of their data centers rather than new equipment. If we do not upgrade our installed base, then surely our competition will. The programs also help Sun to increase market share by allowing competitive trade-ins. The best upgrade is a competitive takeout. When we take out competitive equipment and replace it with Sun, we gain a customer at our competition's expense.

Our upgrade program summarized

Sun's upgrade programs offer many advantages for the customer. In particular, the programs:

- are comprehensive. Any hardware that can be purchased new is also available as an upgrade.
- offers the trade-in value at the time of purchase, not after the equipment is returned
- shows customers, on the quote, the trade-in value of their equipment. This allows customers to write down the amount of the trade in from their books.
- pays the cost of transporting the returned equipment back to Sun (customers only need to call the 800 number)
- allows up to 60 days after receipt of the new equipment for the residual equipment to be returned

Sun offers one of the most comprehensive, systematic and flexible upgrade programs currently available in the industry.



Sun-to-Sun & Competitive Upgrades (Trade-ins) to the Sun Fire Server Family

The following chart illustrates the available upgrade paths to the Sun Fire 3800-6800 server family. Note that the only components that can be kept in upgrading from an Sun Enterprise 3500-6500 server to a Sun Fire 3800-6800 server are PCI adapters.

For more information on upgrade paths and allowances, go to <http://www.sun.com/ibb/upgrades>

Upgrade to:	Sun Fire 3800 Server	Sun Fire 4800/4810 Servers	Sun Fire 6800 Server
Upgrade from:			
Sun Enterprise series	Sun Enterprise 3000-4000 server Sun Enterprise 3500-4500 server	Sun Enterprise 4000-5000 server Sun Enterprise 4500-5500 server	Sun Enterprise 4000-6000 server Sun Enterprise 4500-6500 server
Older SPARC processor systems	Can be upgraded using Sun's Server Consolidation Trade-In Program		
Competitive servers	Qualified systems		

Server Consolidation Trade-in

Sun's Server Consolidation Trade-In Program is an upgrade program that allows customers to consolidate applications from many servers to one (or few) by trading in multiple older systems from Sun or other vendors in return for an allowance toward a single (or few) new system(s). This is achieved by assigning each system to be traded in a point value from a Server Trade-In Point Matrix, with the points being totaled to arrive at the allowance that will be given.

The Sun Fire 3800-6800 servers are ideal platforms for this program. They are ideal as a Server Consolidation platform because of features such as Dynamic System Domains, high reliability, availability and serviceability, as well as outstanding performance and scalability.

More details on the Server Consolidation Trade-In Program can be found in the *Worldwide Configuration and Ordering Guide*:

System Upgrades

Full system upgrades to Sun Fire systems are available. Refer to table above for upgrade path information. Complete details on the Sun Upgrade Allowance Program(UAP) are available at the IBB web site (<http://ibb.eng.upgrades>).



CPU and Memory Upgrades

Upgrades to processors in the Sun Fire 3800-6800 server family are undertaken via a CPU/memory board replacement. Replacement of CPU modules on the CPU/memory board may not be undertaken on site. A customer who wants to add or upgrade their CPUs must order a CPU/memory board containing 2 or 4 CPUs, then return a board from their system to Sun after the system boards have been swapped. The advantage of this method is higher quality and reduced downtime.

Customer upgrading a domain comprising more than one CPU/memory board, where each board contains CPUs and memory, can complete the upgrade without incurring any downtime as long as the existing CPUs are UltraSPARC III Cu 900MHz or above processors. This means that CPU upgrades from UltraSPARC III 750MHz to UltraSPARC III Cu 900MHz will incur some system downtime.

Memory modules will be available as field installable options. If no CPUs are being added or upgraded, the board replacement method is not necessary.

Please refer to the Upgrade Section of the Worldwide Configuration Guide for details on CPU and memory upgrades.

I/O Upgrades

cPCI to PCI and PCI to cPCI upgrades are available. Please refer to the Upgrade Section of the Worldwide Configuration Guide for details.

Ordering Process for Chassis and Base System Upgrades

Please contact your ESun representative to obtain the upgrade procedures and assistance on how to order such upgrades.



Service and Support

The SunSpectrumSM program is an innovative and flexible program that allows customers to choose the level of service best suited to their needs, ranging from mission critical support for maximum solution availability to backup assistance for self-support customers. The SunSpectrum program provides a simple pricing structure in which a single fee covers support for an entire system, including related hardware, peripherals, the Solaris Operating Environment, software and telephone support for Sun software packages. The majority of Sun's customers today take advantage of the SunSpectrum program, underscoring the excellent value that it represents. Customers should check with their local Sun representatives for program and feature availability in their areas.

FEATURE	SUNSPECTRUM PLATINUM SM Mission-critical Support	SUNSPECTRUM GOLD SM Business-critical Support	SUNSPECTRUM SILVER SM Systems Support	SUNSPECTRUM BRONZE SM Self Support
Coverage / Response Time (cont.)				
Customer defined priority setting	Yes	Yes	Yes	Option
• Urgent (phone/on-site)	Live transfer/ 2 hour	Live transfer/ 4 hour	Live transfer/ 4 hour	4 hour / N/A
• Serious (phone/on-site)	Live transfer/ 4 hour	2 hour/next day	2 hour/next day	4 hour / N/A
• Not critical (phone/on-site)	Live transfer/ customer convenience	4 hour/ customer convenience	4 hour/ customer convenience	4 hour / N/A
2-hour on-site response	Yes	Option	Option	No
Additional contacts	Option	Option	Option	Option
Premier Support Features				
Mission critical support team	Yes	For urgent problems	No	No
Sun Vendor Integration Program (SunVIP TM)	Yes	Yes	No	No
Software patch management assistance	Yes	No	No	No
Field change order (FCO) management assistance	Yes	No	No	No
Hardware Support Delivery				
Replacement hardware parts	On-site technician	On-site technician	On-site technician	Courier
Two day parts delivery	N/A	N/A	N/A	Yes
Overnight parts delivery	N/A	N/A	N/A	Option
Same-day parts delivery	Yes	Yes	Yes	Option
Remote Systems Diagnostics				
Remote dial-in analysis	Yes	Yes	Yes	Yes
Sun Remote Services	Yes	Yes	No	No
Remote predictive failure reporting	Yes	Yes	No	No



FEATURE	SUNSPECTRUM PLATINUM SM Mission-critical Support	SUNSPECTRUM GOLD SM Business-critical Support	SUNSPECTRUM SILVER SM Systems Support	SUNSPECTRUM BRONZE SM Self Support
Internet and CD-ROM Support Tools				
SunSolve™ license	Yes	Yes	Yes	Yes
SunSolve EarlyNotifier™ Service	Yes	Yes	Yes	Yes

Warranty

For Sun 3800, 4800 and 4810 servers:

- On-site hardware service/8am-5pm M-F
- 8am-8pm Phone Support/M-F
- 4 hour average response time
- Warranty Duration/1 year
- Customer defined priority
- Knowledge Database Access
- Solaris Operating Environment patches and maintenance releases
- Software upgrade info service
- Software Installation and Media Support/90 days
- Initial installation (up to two domains) is included in the price of the system

For Sun Fire 6800 servers

- On-site hardware service/7x24
- 7x24 Phone Support
- 4 hour average response time
- Warranty Duration/1 year
- Remote systems monitoring
- Multivendor interoperability assistance
- Access to mission critical support process
- Customer defined priority
- Knowledge Database Access
- Solaris Operating Environment patches and maintenance releases
- Software upgrade info service
- Account Management
- Account Reviews



- Software Installation and Media Support/90 days
- Initial installation (up to two domains) is included in the price of the system

Educational Offerings

Sun offers two hands-on courses for customers who want to become familiar with the Sun Fire 3800-6800 hardware and software environments:

- **SM-340: Sun Fire Workgroup/Enterprise Server Maintenance:** This course provides customers with the information and skills necessary to install, configure and maintain the Sun Fire 280R, 3800, 4800, 4810 and 6800 Workgroup-to-Midframe servers. Additionally, customers will learn to diagnose, remove, replace and troubleshoot hardware, software and firmware issues on these servers. Those who would benefit might include field service engineers, Enterprise system network engineers and anyone responsible for maintaining workgroup-to-midframe networked database servers.
- **ES-420: Sun Fire Workgroup/Enterprise Server Administration:** This course provides customers with the information and skills necessary to support the Sun Fire 3800, 4800, 4810 and 6800 midframe servers. Additionally, customers will learn to diagnose and troubleshoot hardware and software issues. Enterprise system administrators, senior service providers, operations managers, Information Technology (IT) planners and anyone responsible for supporting workgroup-to-midframe networked database servers would benefit from this course.

Professional Services Offerings

Sun Fire Mission Critical Solution Services

- **Sun Fire Application Readiness Service**
The Sun Fire Application Readiness Service contains the essential minimum implementation and project management services required to help smoothly implement and configure new Sun Fire installations to the point where the application is ready to load. This service is designed to help fulfill the requirements of Sun's Availability and Quality Policy.
- **Sun Fire Data Center Readiness Service**
The Sun Fire Data Center Readiness Service addresses the particular needs of customers who want to leverage Sun's best practices and experiences to get to market with the fastest implementation approach to running a mission-critical data center environment. This service provides the essential minimum implementation and project management services to help fulfill the requirements of Sun's Availability and Quality Policy. This service includes design and testing phases and provides setup and best-practice guidance for essential operational and procedural issues of Sun Fire systems. This service addresses many issues of configuration, testing, process and documentation not covered in the basic Sun Fire System Application Readiness Service.

Sun ReadySM Service

SunReady services is a holistic phased approach for creating, delivering and sustaining mission-critical production environments based on the customer's service levels to end users. SunReady service is available for Sun Fire 3800-6800 server customers.



Sun Enterprise 3500-6500 Server and Sun Fire 3800-6800 Server Differences

Sun Fire 3800-6800 server family versus Sun Enterprise 3500-6500 Server Family

The Sun Fire 3800-6800 systems are Sun's new midframe servers and offer configuration capacities that are similar to the existing Sun Enterprise 3500-6500 server family. When comparing the two families, the following chart summarizes the new, extended and changed features of the Sun Fire 3800-6800 servers.

	E3500	E4500	E5500	E6500	3800	4800	4810	6800
# of CPUs	8	14	14	30	8	12	12	24
Max Memory	16 GB	28 GB	28 GB	60 GB	64 GB	96 GB	96 GB	192 GB
# of Sbus	12	21	21	45	-	-	-	-
# of PCI	8	14	14	30	-	16	16	32
# of cPCI	-	-	-	-	12	8	8	16
# of Domains	1				Up to 2			Up to 4
Sustained System Bandwidth	3.2 GB/s			2.688 GB/s	9.6 GB/s			
System Interconnect	Bus				Crossbar			
System Clock	Variable (83-100 MHz)			Variable (80-84 MHz)	Fixed (150 MHz)			
Data Transfer Size	32 Bytes				64 Bytes			
O/S	Solaris 2.5.1, 2.6, 7, 8 Operating Environments				Solaris 8 Operating Environment			
Framebuffer Support	Yes				No			
Slot Tradeoff	Yes				No			
System Controller	No				Yes			
Airflow	Side to Side				Front to Back			
Centerplane	Active				Active	Passive		
Power Requirement	100-120 or 220-240 VAC		220-240 VAC		100-120 or 220-240 VAC	220-240 VAC		
Power Cords	1				3			2 or 4 ¹³

¹³ When the Dual Power Grid feature is used, the number of power cords required is increased to 4.



	E3500	E4500	E5500	E6500	3800	4800	4810	6800
Deskside	Yes	Yes	-	-	No	Yes	No	-
Rackmount	No	Yes	-	-	Yes	Yes	Yes	-
Availability Comparison								
Full H/W Redundancy	No				Yes			
Redundant Interconnect	No				Yes			
Redundant Clock	No				Yes, with auto failover			
Enhanced Auto Recovery (SC Failover)	No				Yes			
Dynamic System Domains	No				Yes			
Mixed CPU Speed Support	No				UltraSPARC III Cu 900MHz and above processor			
Hot CPU Upgrade	No				Yes			
Memory ECC	ECC				ECC			
E-cache	Parity				ECC			
Redundant Power & Cooling	Yes				Yes			
Red. Power Feeds	No				Yes, with Redundant Transfer Switch			
Separate Power Planes	No				No			Yes
Management Comparison								
Management Device	No				Embedded System Controller			
Primary Management Interface, SunMC Agent Deployment	Solaris Operating Environment tools via locally attached monitor and keyboard				<ul style="list-style-type: none"> - Sun Management Center Console - Terminal server 			
	One agent				<ul style="list-style-type: none"> - One agent per domain - Proxy module under agent loaded on separate system 			



Glossary

CompactPCI (cPCI)	The standard industry PCI card with hot swap capability.
CPU/Memory board	The board containing up to four UltraSPARC III processors, E cache modules and four memory banks.
DIMM	Dual In-line Memory Module. A small card with DRAM chips on it, used as main memory on the CPU/Memory board.
Domain	Dynamic System Domains - a single instance of Solaris Operating Environment running on a dynamically configurable subset of the hardware resources of a system.
Dynamic Reconfiguration	Ability to change the configuration of a running system by bringing components online or offline without disruption of system operation.
E cache Module	External (level 2) processor cache memory module.
Sun Fireplane Interconnect Board	A board that acts as a point-to-point connection and connects multiple CPU/Memory boards and I/O assemblies.
FRU	Field-replaceable unit or replacement part.
Hot Plug	Allows for the removal of failed components and insertion of replacements without first powering down the system..
Hot Swap	Allows newly-inserted devices to be added to the running system without interruption.
I/O Assembly	A card cage that mounts into the chassis of all of the Sun Fire 3800-6800 servers. The assembly hosts 8 PCI cards or 4 or 6 cPCI cards. The 6-slot cPCI assembly is available only for the Sun Fire 3800 server.
IP link	A communication facility or medium over which nodes can communicate at the link layer. The link layer is the layer immediately below IPv4/IPv6. IP Network Multipathing provides for links that comprise more than one physical wire between a host and a switch.
POST	Power-on self test.
RTS	Redundant Transfer Switch. A very fast switch with microprocessor control and decision-making programming that can switch between two incoming power sources to ensure the system is supplied with good power.
Segment	A logical separation of domains on a machine by pooling domains under control of different Sun Fireplane switch. On a Sun Fire 6800 server, such separation can be physical as well because the power supplies can be segregated as well.
System Controller Board	A board containing a MicroSPARC processor, which oversees operation of the system and provides clocks and the console bus.



Materials Abstract

Collateral	Description	Purpose	Distribution	Token # or COMAC Order #
References				
– <i>Sun Intro</i>	Introduction E-mail	Sales Tool	SunWIN	131754
– <i>SunStore Website</i>	store.sun.com	Configuration details		
Product Literature				
– <i>Sun Fire 3800-6800 Business Focus Whitepaper</i>	Technical paper describing business challenges and how Sun Fire technology can be leveraged to address those business issues and drive competitive advantage	Technical Sales Tool	SunWIN, Reseller Web	131492
– <i>Sun Fire 3800-6800 Technical Overview Whitepaper</i>	Technical tour of Sun Fire product innovations including Dynamic System Domains, Dynamic Reconfiguration and Sun Fireplane interconnect system.	Technical Sales Tool	SunWIN, Reseller Web	131493
– <i>Sun Fire 3800-6800 Availability Whitepaper</i>	Technical paper discussing Sun Fire availability features, the underlying technology and the benefit the customer derives from those features	Technical Sales Tool	SunWIN, Reseller Web	131497
– <i>Sun Fire 3800 Server Datasheet</i>		Technical Sales Tool	SunWIN JavaCart	125853 DE1265-0
– <i>Sun Fire 4800 Server Datasheet</i>		Technical Sales Tool	SunWIN JavaCart	125855 DE1266-0
– <i>Sun Fire 4810 Server Datasheet</i>		Technical Sales Tool	SunWIN JavaCart	125857 DE1267-0
– <i>Sun Fire 6800 Server Datasheet</i>		Technical Sales Tool	SunWIN JavaCart	125859 DE1268-0
– <i>Sun Fire Midrange Server Family Datasheet</i>		Technical Sales Tool	SunWIN JavaCart	131459 DE1424-0
– <i>Sun Fire Midrange Server Family Brochure</i>		Sales Tool	SunWIN JavaCart	69480 BE607-3
– <i>Gold Pitch Customer Presentation</i>	Launch messages for customers	SalesTool	SunWIN	132047
– <i>Business Presentation</i>		SalesTool	SunWIN	123056
– <i>Technical Presentation</i>	This is a midlevel presentation with technical details on the entire Sun Fire server family. This should be given after the customer has been given the general presentation.	Technical Sales Tool	SunWIN	134035
Training				



Collateral	Description	Purpose	Distribution	Token # or COMAC Order #
<ul style="list-style-type: none"> - <i>Web url</i> - <i>SunExperts CD</i> 	webtrain.central	Training Training	JavaCart	WE323-0
Success Stories				
<ul style="list-style-type: none"> - <i>Available in second half of 2001</i> 		Sales Tool		
Demos				
<ul style="list-style-type: none"> - <i>Tour URL</i> 	Virtual reality tour of the Sun Fire systems www.sun.com/servers/midrange/interactivetour/	Technical Sales Tool	CD available via JavaCart	WE311-0
Videos				
<ul style="list-style-type: none"> - <i>Reference (Cust, ISVs)</i> - <i>Demo</i> 		Sales Tool		ME2562-0
External Web Sites				
<ul style="list-style-type: none"> - <i>Sun Documentation</i> 	http://docs.sun.com			
Internal Web Sites	http://systems.corp/			
<ul style="list-style-type: none"> - <i>Sun Fire 3800-6800 product pages</i> 	http://systems.corp/products/servers/datacenter/ http://onestop.central/hw/sunfire-midframe.shtml			
<ul style="list-style-type: none"> - <i>Worldwide Configuration and Ordering Guide</i> 	http://mysales.central/public/configGuide/			
<ul style="list-style-type: none"> - <i>Upgrade Information</i> 	http://ibb.eng/upgrades			



Change Log

August, 2001

Mixed speed CPU plan updated
Rack configuration information updated
Materials Abstract updated
Service and Support pages updated
I/O Matrix updated
Roadmap updated

November, 2001

Add ULTRASparc III Cu
Add StorEdge 9900
Update Electrical/Thermal specifications
Update Warranty section

January, 2002

Dynamic Reconfiguration updated
Dynamic System Domains updated

June, 2002

Enhanced Auto Recovery (SC Failover) updated

November, 2002

COD 2.0 updated

February, 2003

Updated COD 2.0 section

April, 2003

Add Auto Diagnosis and Recovery Features

November, 2003

Updated COD 2.0 section

February, 2004

Add SSH feature

